

# 000 KNOWLEDGESMITH: UNCOVERING KNOWLEDGE 001 UPDATING IN LLMs WITH MODEL EDITING AND UN- 002 LEARNING 003

006 **Anonymous authors**

007 Paper under double-blind review

## 011 ABSTRACT

013 Knowledge editing and machine unlearning are two popular approaches for large  
014 language models (LLMs) to stay up-to-date. However, the knowledge updating  
015 mechanism of LLMs remains largely unexplored due to insufficient, isolated, and  
016 small-scale evaluation. For instance, are LLMs similar to humans in modifying  
017 certain knowledge? What differs editing and unlearning as training data increases?  
018 This paper proposes KnowledgeSmith, a unified framework to systematically un-  
019 derstand the updating mechanism of LLMs. We first cast editing and unlearning  
020 as instances of one constrained optimization problem. Then, we propose an au-  
021 tomatic dataset generator that provides structured interventions across multiple  
022 graph levels and data scales, enabling controlled studies of how different mod-  
023 ification strategies propagate through model knowledge. Extensive experiments  
024 demonstrate nuanced insights over knowledge propagation, plasticity scaling, con-  
025 sistency, and robustness. For instance, our results show that LLMs do not exhibit  
026 similar updating as humans for different levels of knowledge, and there exists  
027 consistency-capacity trade-off. We hope our findings can offer suggestions to the  
028 design of more reliable and scalable strategies.

## 029 1 INTRODUCTION

030 Human knowledge is not stored as isolated facts but as a vast, interconnected web (Liu et al., 2024).  
031 From early encyclopedias to modern knowledge graphs, we represent knowledge as structured re-  
032 lations (Yang et al., 2025): concepts (nodes) linked by semantic or causal connections (edges).  
033 This networked organization enables humans to reason flexibly (Mark et al., 2020), update beliefs  
034 (Paulheim, 2016) when new evidence arises, and propagate changes across related domains (Flouris  
035 et al., 2008). For instance, when scientists revised the classification of Pluto from a planet to a  
036 dwarf planet, the update did not merely alter one fact but cascaded through textbooks, curricula, and  
037 related scientific explanations.

038 Do Large language models (LLMs) exhibit similar properties? Zhang et al. (2024) showed that they  
039 store and retrieve information at scale, generating answers that span diverse domains; Yet, unlike  
040 human knowledge graphs, the internal structure of LLM knowledge remains opaque (Zhang et al.,  
041 2023). Fine-tuning can overwrite large swaths of parameters but is resource-intensive and imprecise  
042 (Balne et al., 2024; Gekhman et al., 2024), often introducing instability or hallucinations (Khan  
043 et al., 2025; Ovadia et al., 2024). Researchers have recently shifted attention toward knowledge  
044 editing (Wei et al., 2024; Markowitz et al., 2025; Wang et al., 2024) and unlearning (Yao et al.,  
045 2024; Pawelczyk et al., 2024; Hong et al., 2024), where editing offers targeted modifications and  
046 unlearning aims to broadly remove specific information. Both are valuable, yet they are typically  
047 studied in isolation and without grounding in structured knowledge representations.

048 How to understand the knowledge updating mechanism in LLMs? Recent efforts show that editing  
049 techniques can be adapted for forgetting by redirecting or suppressing knowledge representations  
050 (Li et al., 2025b; Jung et al., 2025), while unlearning methods sometimes resemble coarse-grained  
051 editing at the dataset level (Guo et al., 2019). Other works investigate continual or compositional  
052 settings, where localized edits may interfere with broader forgetting objectives or vice versa (Gupta  
053 et al., 2024; Chen et al., 2024). A parallel strand examines the tension between specificity and gen-  
054 eralization: editing often prioritizes precision but risks side effects, whereas unlearning emphasizes  
055 removal but may fail to incorporate new or corrected knowledge (Yao et al., 2023a).

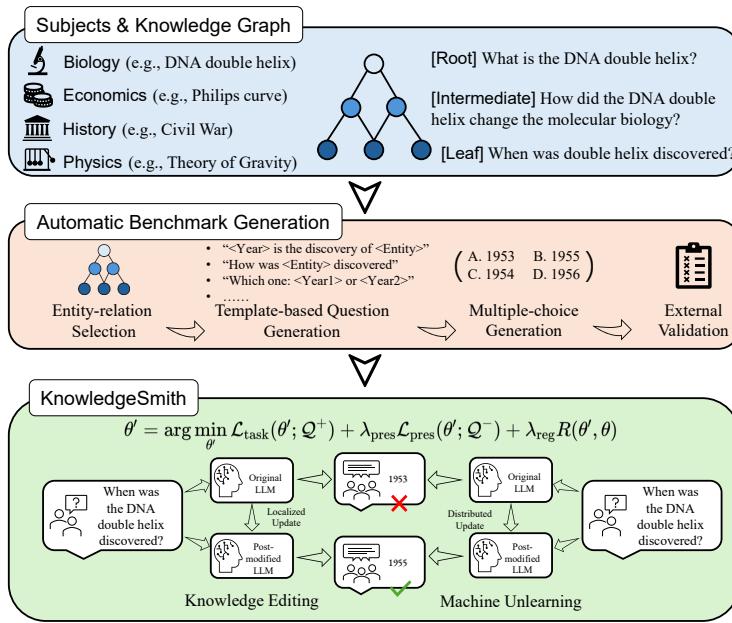


Figure 1: KnowledgeSmith pipeline. Starting from static KG, we generate dynamic probes at root, intermediate, and leaf levels, enabling evaluation of direct and propagated effects.

Despite recent progress, there are still three critical challenges. First, most evaluations target isolated facts, neglecting the structured and interconnected nature of real-world knowledge (Thede et al., 2025). For example, if we update the fact that “Lyon is the capital of France” instead of Paris, a coherent system should also adjust related knowledge such as “the Eiffel Tower is located in France’s capital,” which otherwise becomes inconsistent. Second, the role of data scale in editing vs. unlearning remains unclear, with small data often sufficing for edits but not for forgetting (Zhong et al., 2023; Meng et al., 2022a). Third, there is no unified framework to jointly understand editing and unlearning, leaving their trade-offs in propagation, stability, and generalization unclear.

In this paper, we introduce **KnowledgeSmith** (Figure 1), a unified framework to understand the knowledge updating mechanisms in LLMs.<sup>1</sup> Theoretically, our framework casts editing and unlearning as complementary forms of constrained optimization. Empirically, building on the intuition that human knowledge is naturally structured as knowledge graphs (KGs), our framework can automatically transform any existing KG-related dataset into a benchmark for knowledge intervention evaluation, enabling systematic and scalable assessment without the need for hand-crafted test sets. For instance, more insights can be gained through interventions across hierarchical levels (root, intermediate, leaf) and data scales (from single instances to millions). Then, we conduct an extensive evaluation of editing and unlearning on different LLM families to explore knowledge propagation, scaling laws, representation shifts, and robustness under stress tests. Our key findings are:

- Propagation Asymmetry and Plasticity Limits:** Editing can over-spread(unintentionally altering related nodes), especially at higher nodes, while unlearning mostly under-spreads(forgetting failing to propagate beyond the target node). Hierarchical branch structure imposes intrinsic ceilings on update effectiveness, with higher or more central nodes limiting achievable knowledge modifications(§5.2.1,§5.2.2).
- Consistency–Capacity Tradeoff and Subject-Dependent Update:** Increasing data can trigger consistency collapse, where local updates contradict other knowledge; editing prioritizes local enforcement, unlearning preserves broader consistency. Some domains, like history, resist updates more than others, highlighting the need for subject-aware evaluation (§5.2.3,§5.2.4).
- Model Robustness:** Editing improves in-domain accuracy but harms OOD and adversarial stability, while unlearning preserves global robustness at the cost of weaker local gains(§5.3).
- Method-level Trade-offs:** Editing balances integration and preservation with strong low-data efficiency, unlearning is conservative but stable, while LoRA fine-tuning is unstable and prone to drift, making it unreliable for continual updates (§5.4).

<sup>1</sup>Other approaches can also update knowledge in LLMs; we focus on editing and unlearning in this paper.

108     5. **Unified Failure Modes and Stress Testing:** By observing model behavior on open-ended  
 109     questions, we identify six main failure modes and find that unlearning preserves general task  
 110     integrity better, whereas editing is more aggressive but effective in low-data regimes (§5.5).

111     **Contributions.** (1) We introduce KnowledgeSmith as a unified framework to understand knowledge  
 112     updating in LLMs with editing and unlearning. (2) We present automatic data generation pipeline  
 113     for LLM evaluation with scalable KG-structured interventions. (3) Our experiments demonstrate  
 114     several insightful findings towards LLM knowledge updating that could inspire future research.

## 116     2 RELATED WORK

118     Other than fine-tuning which is expensive and requires large amount of training data, knowledge  
 119     editing and machine unlearning are two popular and effective approaches to update LLMs' knowl-  
 120     edge. Knowledge editing modifies LLMs' internal parameters to update its predictions on specific  
 121     factual associations while ideally preserving unrelated knowledge (Yao et al., 2023b; Cao et al.,  
 122     2021; Sinitzin et al., 2020). Existing approaches include gradient-based fine-tuning (Sinitzin et al.,  
 123     2020; Zhu et al., 2020), localized weight modifications such as ROME (Meng et al., 2022a), MEMIT  
 124     (Meng et al., 2022b), and SERAC (Mitchell et al., 2021), and memory-augmented methods that  
 125     externalize edits (Mitchell et al., 2022). However, most prior evaluations are restricted to small  
 126     benchmarks (Levy et al., 2017; Meng et al., 2022a) and do not examine how edits propagate through  
 127     structured knowledge dependencies.

128     On the other hand, motivated by ethical, legal, or safety considerations, machine unlearning seeks  
 129     to selectively erase information linked to a dataset, (Izzo et al., 2021; Thudi et al., 2022; Xu et al.,  
 130     2025). Methods include retraining-based approaches (Ginart et al., 2019), negative-gradient fine-  
 131     tuning (Thudi et al., 2022), regularization-based constraints (Golatkar et al., 2020), and approximate  
 132     removal via influence functions or Fisher-weighted updates (Guo et al., 2019; Baumhauer et al.,  
 133     2022). Yet, unlearning has largely been studied in isolation from editing, without systematic com-  
 134     parisons or evaluation in structured knowledge contexts.

135     In short, existing research highlights strong methodological advances but leaves two key gaps: (1)  
 136     editing and unlearning are often treated as disjoint problems despite their conceptual overlap, and  
 137     (2) evaluations rely on narrow datasets that fail to capture scaling behavior or structured propagation  
 138     effects. Our work tries to establish a unified view of them and present an extensive analysis towards  
 139     understanding LLM knowledge updating.

## 141     3 KNOWLEDGESMITH

143     In this section, we propose **KnowledgeSmith**, a unified framework to view editing and unlearning  
 144     as complementary interventions.

### 146     3.1 PROBLEM DEFINITION

148     Let  $f_\theta$  denote a language model parameterized by  $\theta$ , defining a conditional distribution  $p_\theta(y | x)$   
 149     over output  $y$  given input  $x$ . We study targeted interventions that modify or remove specific knowl-  
 150     edge while preserving the model's general behavior.

151     An *update request* is given by an item  $e$  (e.g., a factual triple, a prompt-response pair, or a small  
 152     dataset), optionally accompanied by a scope  $c$  that defines locality or related probes. For example,  
 153     if  $e$  is the fact “Paris is the capital of France”,  $c$  could include all prompts asking about European  
 154     capitals such as “What is the capital of France?” or “Name the capital of European countries” while  
 155     excluding unrelated prompts like “Who is the president of the United States?”, ensuring that only  
 156     related knowledge is affected while leaving unrelated knowledge untouched. Applying an update  
 157     operator  $\mathcal{T}$  (e.g., editing or unlearning) yields updated parameters:

$$158 \quad \theta' = \mathcal{T}(\theta; e, c), \quad \Delta = \theta' - \theta, \quad (1)$$

159     where  $\Delta$  is the parameter update.

160     The objective is therefore to update the targeted knowledge while preserving unrelated knowledge.  
 161     To facilitate analysis, we define two probe sets: (1) *Positive probes*  $\mathcal{Q}^+$  are inputs where the model's

162 predictions should change; and (2) *Preservation probes*  $\mathcal{Q}^-$  are inputs where predictions should  
 163 remain unchanged. Formally, for an input  $x$ , denote  $p_\theta(\cdot | x)$  and  $p_{\theta'}(\cdot | x)$  as the output distribution  
 164 of the model before and after KnowledgeSmith intervention, respectively, we have:

$$\begin{aligned} 165 \quad d(p_{\theta'}(\cdot | x), q_{\text{target}}(\cdot | x)) &\leq \eta^+, \quad \forall x \in \mathcal{Q}^+, \\ 166 \quad d(p_{\theta'}(\cdot | x), p_\theta(\cdot | x)) &\leq \varepsilon, \quad \forall x \in \mathcal{Q}^-, \end{aligned} \quad (2)$$

168 where  $d(\cdot, \cdot)$  is a divergence or distance measure between distributions (e.g., KL divergence, cross-  
 169 entropy, or  $\ell_2$  distance over logits),  $q_{\text{target}}(\cdot | x)$  is the desired post-intervention distribution on  
 170 positive probes, the constant  $\eta^+$  specifies a tolerance threshold for successful edits, reflecting that  
 171 editing algorithms may only approximate the target distribution rather than match it exactly, and  $\varepsilon$   
 172 is a stability threshold controlling how much drift is allowed on  $\mathcal{Q}^-$ .

### 173 3.2 A UNIFIED FRAMEWORK FOR ANALYZING EDITING AND UNLEARNING

175 While Equation (2) formalizes the objectives using tolerance thresholds  $\eta^+$  and  $\varepsilon$ , in practice we  
 176 implement these constraints by relaxing them into loss terms over probes. Specifically,  $\mathcal{L}_{\text{task}}(\theta'; \mathcal{Q}^+)$   
 177 penalizes deviations from the target distribution on  $\mathcal{Q}^+$ ,  $\mathcal{L}_{\text{pres}}(\theta'; \mathcal{Q}^-)$  penalizes drift on  $\mathcal{Q}^-$ , and  
 178  $R(\theta', \theta)$  regularizes the overall update. Thus, both model editing and unlearning can be cast as a  
 179 constrained optimization over model parameters:

$$180 \quad \theta' = \arg \min_{\theta'} \mathcal{L}_{\text{task}}(\theta'; \mathcal{Q}^+) + \lambda_{\text{pres}} \mathcal{L}_{\text{pres}}(\theta'; \mathcal{Q}^-) + \lambda_{\text{reg}} R(\theta', \theta), \quad (3)$$

182 where  $\mathcal{L}_{\text{task}}$  enforces the desired behavior on  $\mathcal{Q}^+$ ,  $\mathcal{L}_{\text{pres}}$  penalizes drift on  $\mathcal{Q}^-$ , and  $R(\theta', \theta)$  regular-  
 183 izes the update (e.g.,  $\|\Delta\|_2^2$  (Ng, 2004), Fisher norm (Gu et al., 2012), or others (Hu et al., 2022)).

184 **Editing as targeted alignment.** Knowledge editing can be viewed as minimizing  $\mathcal{L}_{\text{task}}$  toward a  
 185 distribution  $q_{\text{target}}$  that encodes corrected knowledge. For example, ROME (Meng et al., 2022a) and  
 186 MEMIT (Meng et al., 2022b) locate and modify specific MLP weights to enforce new facts, while  
 187 MEND (Mitchell et al., 2021) trains an auxiliary retriever–classifier to redirect predictions on edited  
 188 queries. Other approaches apply gradient-based updates on  $\mathcal{Q}^+$  while regularizing drift, such as  
 189 GRACE (Hartvigsen et al., 2023). Even parameter-efficient methods like LoRA-based editing (Hu  
 190 et al., 2022; Zheng et al., 2023) fit this form, with  $R(\theta', \theta)$  enforcing low-rank adaptation.

191 **Unlearning as neutral alignment.** Unlearning corresponds to the same objective but with  
 192  $q_{\text{target}}$  chosen as a *neutral distribution*  $q_{\text{neutral}}$  that suppresses unwanted associations. This captures  
 193 approaches that erase knowledge through gradient descent (Thudi et al., 2022), influence-  
 194 function-based forgetting (Golatkar et al., 2020; Guo et al., 2019), or certified removal in convex  
 195 models (Ginart et al., 2019). Recent work on unlearning in deep networks (Jagielski et al., 2022)  
 196 also fits: their objectives penalize predictive alignment with sensitive data while constraining per-  
 197 formance on  $\mathcal{Q}^-$ , exactly corresponding to the  $\mathcal{L}_{\text{pres}}$  and  $R(\theta', \theta)$  terms above.

198 **A unifying lens.** In this view, the distinction between editing and unlearning reduces to the choice  
 199 of  $q_{\text{target}}$ : Editing:  $q_{\text{target}}$  encodes a factual correction (e.g., “Paris is the capital of Germany”). Un-  
 200 learning:  $q_{\text{target}}$  is neutral, erasing prior associations (e.g., “Paris is the capital of [MASK]”). This  
 201 framework subsumes methods across the spectrum: localized weight modifications (Meng et al.,  
 202 2022b;a), memory-based editors (Mitchell et al., 2021), parameter-efficient adaptations (Hu et al.,  
 203 2022; Zheng et al., 2023), influence-based forgetting (Golatkar et al., 2020), and certified removal  
 204 (Ginart et al., 2019). Despite methodological differences, all can be interpreted as solving the same  
 205 constrained optimization problem with different instantiations of  $\mathcal{L}_{\text{task}}$ ,  $\mathcal{L}_{\text{pres}}$ , and  $R(\theta', \theta)$ .

206 Our formulation provides a principled and generalized lens for analyzing parameter modifications in  
 207 LLMs, enabling fair comparison of editing and unlearning on their trade-offs in plasticity, stability,  
 208 and generalization. However, to rigorously measure these effects in practice, we need benchmarks  
 209 that capture hierarchical dependencies, e.g., local versus global changes, and multilevel propaga-  
 210 tion of updates, which are largely missing from existing datasets. This motivates our automated  
 211 benchmark construction in the following.

## 212 4 CONSTRUCTING EVALUATION BENCHMARK

213 Existing benchmarks (Meng et al., 2022a; Levy et al., 2017) for knowledge intervention evaluation  
 214 suffer from two major limitations. First, they are largely *static*, testing only isolated facts without

216 accounting for how updates might affect related knowledge. Second, they fail to capture *dependencies across facts*, which are crucial for understanding how changes propagate through the model and  
 217 for revealing trade-offs between editing and unlearning.  
 218

219 We leverage *knowledge graphs (KGs)* to address these gaps, which dynamically encode hierarchical  
 220 and relational dependencies among facts. Anchoring probes in a curated KG enables us to generate  
 221 both local edits and their downstream consequences, transforming a single KG into a dynamic  
 222 benchmark. Specifically, by targeting *root*, *intermediate*, and *leaf* nodes, our framework systematically  
 223 tests how interventions propagate across multiple levels of dependency, thus providing a rigorous  
 224 way to evaluate whether models can coherently update, forget, or preserve knowledge while  
 225 maintaining global consistency. Concretely speaking, our data generation method can automatically  
 226 transform any existing knowledge-related benchmarks such as MMLU (Hendrycks et al., 2021) into  
 227 new ones, providing domain coverage and a standardized multiple-choice QA format for easy eval-  
 228 uation. Our pipeline consists of three stages (Figure 1), ensuring both quality and flexibility:  
 229

- 230 1. **Entity–Relation Selection:** We begin by prompting GPT-4o to generate a KG where enti-  
 231 ties and relations are organized hierarchically. The model is then asked to categorize nodes  
 232 into three levels: *root* (broad, domain-level concepts), *intermediate* (mid-level categories or  
 233 subtopics), and *leaf* (specific entities or instances). Sampling nodes from all three categories  
 234 preserves the KG’s hierarchical structure, ensuring evaluation goes beyond isolated facts to  
 235 capture how edits or deletions propagate across different levels of related knowledge.
- 236 2. **Template-Based Question Generation:** Multiple question forms are generated for each triple,  
 237 varying in directness and context. All templates are manually verified for grammaticality and  
 238 factual alignment, preserving unambiguous mapping back to the KG. Six categories of probes  
 239 are constructed (direct, reverse, conflict, multi-hop, comparison and contextual), each tied to a  
 240 different aspect of model behavior under intervention.
- 241 3. **Multiple-Choice Construction:** Each probe is cast as a four-choice QA item, consistent with  
 242 the MMLU-inspired format, ensuring that evaluation reflects true knowledge states rather than  
 243 guesswork. Entity substitution and paraphrasing yield over one million samples across do-  
 244 mains. All items are validated against the KG, with manual spot checks for quality assurance.

245 **Connection to KG-Based Evaluation.** Our generation pipeline is organized around two comple-  
 246 mentary families of probes: (1) *Positive probes*  $Q^+$ , which directly test the edited or redirected  
 247 knowledge, including its hierarchical propagation across root, intermediate, and leaf nodes. (2)  
 248 *Preservation probes*  $Q^-$ , which ensure that unrelated or out-of-scope knowledge remains intact,  
 249 guarding against collateral damage.

250 To operationalize these two families, we instantiate six probe types. *Direct probes* ( $Q^+$ ) test whether  
 251 the target fact itself is recalled or updated at different hierarchical levels. *Reverse probes* ( $Q^+$ ) ex-  
 252 amine whether knowledge updates preserve relation directionality. *Conflict probes* ( $Q^+ / Q^-$ ) expose  
 253 residual beliefs and adversarial robustness by checking for contradictions after intervention. *Multi-*  
 254 *hop probes* ( $Q^+$ ) evaluate whether interventions correctly propagate through chained relations in the  
 255 KG. *Comparison probes* ( $Q^+$ ) assess whether the updated knowledge is consistently preferred when  
 256 contrasted with alternatives or distractors. Finally, *Contextual probes* ( $Q^-$ ) test whether unrelated  
 257 in-domain or OOD knowledge remains preserved in naturalistic settings. This design aligns directly  
 258 with our experimental analyses: By explicitly embedding these probe types into the KG’s hierar-  
 259 chical structure, the benchmark enables analyses that go beyond isolated fact checking, revealing  
 260 whether interventions cascade consistently across levels of related knowledge.

261 **Generated Benchmark Dataset.** Our method allows flexible data generation across domains. In  
 262 this paper, we instantiate the benchmark in four domains: economics, physics, history, and biology.<sup>2</sup> Each domain  
 263 yields paired *pre-edit* and *post-edit* datasets that preserve entities but differ in factual content. Probes  
 264 span root, intermediate, and leaf nodes, with conflict, propagation, comparative, and reverse variants,  
 265 and include multiple paraphrased realizations. For each branch within every domain, we generate  
 266 10,000 samples each for editing and unlearning, plus 100 evaluation probe sets, leading to 360,000  
 267 training samples in total. This design creates a benchmark that is both large-scale and structurally  
 268 sensitive, allowing systematic evaluation of edits and unlearning not just at the point of intervention  
 269 but throughout the knowledge hierarchy. Dataset examples are in Appendix A.

<sup>2</sup>These subjects span both STEM and humanities, offering a representative testbed. Our pipeline is directly extensible to other domains such as law and medicine.

270 

## 5 EXPERIMENTS

271 

### 5.1 SETUP

273 **Models.** Our evaluation covers 6 families of LLMs with 1B to 123B parameters, leading to a total of  
 274 13 models: LLaMA-3 (1B, 3B, 8B, 70B) (Meta, 2024), Qwen-3 (1.7B, 14B, 32B) (Team, 2025b),  
 275 QwQ-32B (Team, 2025a), Mistral (24B, 123B) (Jiang et al., 2023), Gemma (2B, 7B) (Team, 2024),  
 276 and DeepSeek-R1-0528-Qwen3-8B (DeepSeek-AI, 2025). This broad coverage enables us to study  
 277 whether scaling behaviors and editing/unlearning performance generalize across architectures.  
 278

279 **Implementation Details.** We adopted AlphaEdit (Fang et al., 2025) and ReLearn (Xu et al., 2025).<sup>3</sup>  
 280 AlphaEdit is a state-of-the-art editor that has been shown to outperform prior methods such as  
 281 MEMIT(Meng et al., 2022b) and ROME(Meng et al., 2022a) in editing tasks, while ReLearn repre-  
 282 sents a leading approach to unlearning. Importantly, our framework is method-agnostic and directly  
 283 extensible to other baselines, making it straightforward to integrate additional methods. Unlike tra-  
 284 ditional unlearning approaches where the retain set corresponds to the original knowledge, in our  
 285 redirection-based setup the retain set is defined as the post-updated knowledge, ensuring that the  
 286 model preserves the rewritten fact rather than reverting to its prior belief. This redirection-based  
 287 formulation aligns better with real-world scenarios where knowledge is updated rather than erased.  
 288 Editing and unlearning were applied separately to leaf, intermediate, and root nodes of the knowl-  
 289 edge graph, with training data sizes ranging from 1 to 10,000 samples. This setup allowed us to  
 290 systematically analyze the effect of both hierarchy depth and data scale on the success of editing  
 291 and unlearning. For evaluation, since each knowledge probing question is multiple-choice, we re-  
 292 port accuracy as the proportion of questions for which the model selects the correct choice. This  
 293 metric directly reflects the model’s correctness in retrieving or updating the intended knowledge.  
 294

295 

### 5.2 COMPARATIVE ANALYSIS OF EDITING AND UNLEARNING

296 

#### 5.2.1 PROPAGATION ASYMMETRY: OVER- VS. UNDER-SPREADING

297 Human learners expect hierarchical consistency: updating a root concept should cas-  
 298 cade to its descendants, while modifying a leaf should remain localized. We evaluate this in  
 299 LLMs by applying editing or unlearning at  
 300 three hierarchy levels (root, intermediate, leaf)  
 301 and measuring performance on both targeted  
 302 and structurally related nodes. We quantify  
 303 these effects using *direct* vs. *multi-hop accu-*  
 304 *racy* (Figure 2) as a proxy for propagation metrics:  
 305 the *Collateral Change Ratio* (*CCR*) captures  
 306 over-spreading for editing, and the *Residual Retention* (*RR*) captures under-spreading for unlearning  
 (For the complete definitions of CCR and RR, see Appendix B).

307 Our results reveal a clear asymmetry: **editing tends to over-spread**, unintentionally altering related  
 308 nodes, especially in lower hierarchy levels, whereas **unlearning often under-spreads**, failing to  
 309 propagate forgetting beyond the target. These simple, interpretable metrics allow us to visualize  
 310 propagation behavior across hierarchical branches.

311 

#### 5.2.2 PLASTICITY SCALING AND BRANCH-DEPENDENT LIMITS

312 Plasticity captures how readily a model can update knowledge in response to limited training data,  
 313 balancing the optimization of  $\mathcal{L}_{\text{task}}$  on positive probes  $\mathcal{Q}^+$  against preservation constraints  $\mathcal{L}_{\text{pres}}$  on  
 314  $\mathcal{Q}^-$ . We extend this notion to *plasticity scaling*, examining systematically how model size, data  
 315 scale, and hierarchical branch jointly influence the effectiveness of editing and unlearning.

316 Our main observations are as follows. First, as shown in Figures 3a and 3b, **smaller models exhibit**  
 317 **higher immediate plasticity**, rapidly adapting to few-shot interventions and achieving strong in-  
 318 domain performance on  $\mathcal{Q}^+$ , but their changes are often unstable, leading to degraded preservation  
 319 on  $\mathcal{Q}^-$ . **Larger models require more data to register updates, reflecting lower short-term plas-**  
 320 **ticity**, yet once modified they maintain stronger out-of-domain consistency, indicating more reliable  
 321 preservation. Second, **branch-dependent upper bounds**. As shown in Figure 3c, different hier-

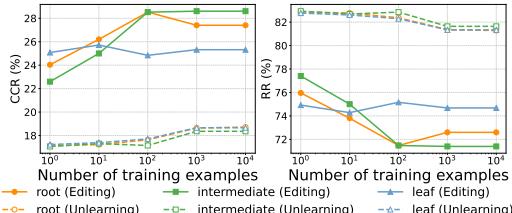


Figure 2: Propagation asymmetry metrics.

323 <sup>3</sup>We also conduct experiments on some other methods, which show similar performance. Reported at  
 Appendix E.

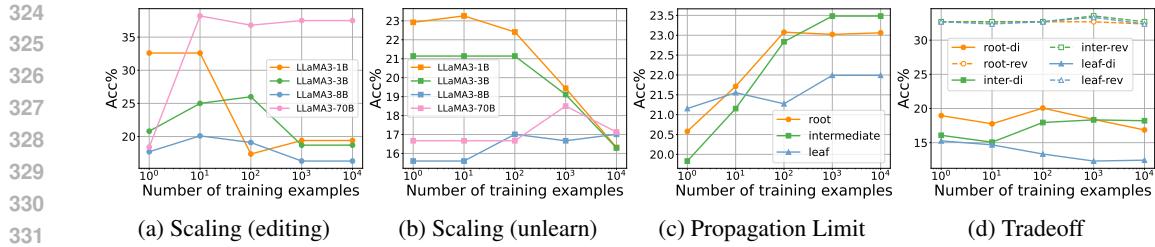


Figure 3: Plasticity scaling of the LLaMA3 family under (a) editing and (b) unlearning. (c) Propagation limits across three branches. (d) Consistency capacity tradeoff.

archival branches exhibit distinct ceilings for achievable accuracy. Root-level edits/unlearning face a lower ceiling due to structural complexity and the need for coherent propagation across descendants. Intermediate-level branches achieve moderate ceilings. Leaf-level edits/unlearning can reach near-perfect in-domain accuracy with fewer examples, reflecting minimal propagation constraints. This reveals the effectiveness of updates is not uniform across the hierarchy: higher or more central nodes constrain achievable plasticity, while lower nodes allow maximal update with limited data.

### 5.2.3 CONSISTENCY–CAPACITY TRADE-OFF

Most prior work (Zhong et al., 2023; Park et al., 2025; Shi et al., 2024; Li et al., 2025a) primarily assess whether the target fact is updated successfully, without probing inverse relations. To our knowledge, no prior work explicitly quantifies this type of cross-relation or hierarchical consistency. In this work, we define consistency as the model’s ability to maintain logical coherence across related knowledge after an intervention. Specifically, we test consistency by probing both the direct relation (e.g., “Paris is the capital of France”) and the inverse or complementary relation (e.g., “France has capital Paris”), as well as across hierarchical or semantically related branches. A consistent update should correctly modify the target knowledge while preserving these related facts.

We uncover a new phenomenon: **consistency collapses once data scale surpasses the model capacity**. We term this the *consistency–capacity trade-off*, observed both in relation–inverse relation pairs (e.g., *capital-of* vs. *has-capital*) and across hierarchical branches. As shown in Figure 3d, direct probes initially respond to interventions but plateau or degrade as training scale grows, whereas reverse probes remain stably high, indicating preservation of contradictory knowledge. The divergence defines a *consistency collapse point*, occurring earlier in lower branches (intermediate, leaf) than root. Editing typically achieves stronger local updates but triggers earlier global inconsistency; unlearning preserves broader consistency but rarely removes the targeted knowledge completely.

**Representation and Efficiency.** Table 1 shows the analysis of internal representations via Centered Kernel Alignment (CKA) (Kornblith et al., 2019), KL divergence, L2 distance and Fisher score (Zhang et al., 2022). The results show that unlearning exhibits abrupt phase transitions beyond a critical data scale, while editing induces smoother, localized adjustments (details in Appendix H). Computationally, unlearning is faster (e.g.,  $\sim 0.2$ h vs  $\sim 6$ h for editing on 1,000 samples on an NVIDIA H100), reflecting its focus on stability over precise enforcement.

Consistency collapse is not only evident in output accuracy but also mirrored in representation dynamics and computational cost: editing maximizes factual enforcement at the expense of broader consistency and resources, whereas unlearning prioritizes stability and efficiency.

### 5.2.4 SUBJECT-DEPENDENT KNOWLEDGE UPDATE

At the subject level, Figure 4a reveal that **knowledge updating is strongly subject-dependent**. Among the four subjects (biology, economics, history, and physics), **history consistently exhibits the lowest update accuracy**, sometimes remaining nearly unchanged even with large numbers of training examples. Other subjects update, in contrast, propagate more efficiently. This highlights a critical insight: evaluation benchmarks must account for **subject-specific difficulty**. Standard

Table 1: Similarity scores for each model are independently normalized via a log–min–max transformation: a small positive offset  $\epsilon$  is added,  $\log_{10}$  is applied, and the resulting values are linearly scaled to the  $[0, 1]$  range.

Metric	Setting	1	10	100	1000	10000
KL	Unlearn	0.014	0.392	0.805	0.838	0.883
	Edit	0.140	0.522	0.606	0.647	0.652
L2	Unlearn	0.013	0.286	0.647	0.758	0.948
	Edit	0.054	0.368	0.507	0.628	0.633
Fisher	Unlearn	0.014	0.352	0.781	0.847	0.919
	Edit	0.101	0.438	0.552	0.641	0.647
CKA	Unlearn	0.917	0.861	0.566	0.576	0.692
	Edit	0.958	0.852	0.801	0.714	0.714

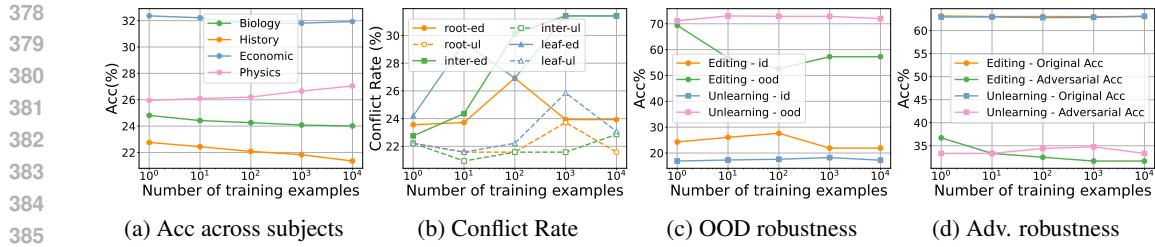


Figure 4: Robustness evaluation under multiple stress tests. (a) Out-of-distribution (OOD) vs. in-domain accuracy. (b) Adversarial robustness relative to original accuracy. (c) Instruction-following accuracy in free generation, judged by an LLM. (d) Hallucination tendency across interventions.

datasets (e.g., CounterFact (Meng et al., 2022a), ZsRE (Levy et al., 2017)) treat all domains equivalently, but our results indicate that certain knowledge domains, such as history, are significantly more resistant to modification. Consequently, subject-aware evaluation is essential for accurately assessing editing and unlearning performance in LLMs.

### 5.2.5 CONTRADICTIONS AND CONFLICT RATE

While residual belief (Elidan et al., 2012) is commonly used to evaluate whether interventions succeed in suppressing prior knowledge, it does not capture a critical failure mode: the emergence of *contradictions*. We therefore introduce a complementary metric, *conflict rate*, which measures the proportion of queries where the model simultaneously supports mutually inconsistent statements after intervention. For instance, a model may assert both “Paris is the capital of Germany” and “Paris is the capital of France” under different contexts. Figure 4b shows this metric exposes patterns that residual belief alone cannot: **editing often leads to higher conflict in related branches (over-spreading), whereas unlearning tends to leave contradictions unresolved in upstream nodes (under-spreading)**. By explicitly quantifying such inconsistencies, conflict rate provides a fuller view of hidden instabilities and unintended side effects.

### 5.3 ANALYSIS ON ROBUSTNESS

OOD robustness is tested using MMLU (Hendrycks et al., 2021). In the unified framework, in-domain probes  $\mathcal{Q}^+$  consist of questions from the same subject (e.g., updating facts about geography using geography questions), reflecting alignment with  $q_{\text{target}}$ . In contrast, out-of-domain (OOD) probes  $\mathcal{Q}^-$  are drawn from unrelated subjects (e.g., updating geography facts but measuring performance on economics, history, or law), testing the model’s ability to preserve unrelated knowledge after the intervention. As shown in Figure 4c, **these objectives often conflict**. Unlearning preserves strong OOD accuracy (63–82%) but yields modest in-domain gains ( $\leq 30\%$ ), while editing substantially boosts in-domain accuracy (up to 50–60% in economics) at the cost of OOD stability, especially in mid-sized models. Larger models reduce but do not eliminate this trade-off. Increasing training examples improves in-domain performance until gains plateau, and disciplines vary, with economics generalizing better and history proving more resistant. This trade-off reflects the balance between  $\mathcal{L}_{\text{task}}$  and  $\mathcal{L}_{\text{pres}}$ : stronger enforcement on  $\mathcal{Q}^+$  tends to destabilize preservation on  $\mathcal{Q}^-$ , highlighting the challenge of achieving both local fidelity and global robustness together.

We then measure adversarial robustness by exposing the model to misleading or deceptive inputs, such as probes combining unrelated concepts (Figure 4d). This assesses whether the optimization constraints maintain stability on preservation probes  $\mathcal{Q}^-$  under stress (details in Appendix D.1).

### 5.4 ANALYSIS ON FINE-TUNING

We further compare editing and unlearning with LoRA fine-tuning on Llama3-8B-Instruct to isolate method-level tradeoffs. Figure 6a shows LoRA yields unstable ID accuracy, sometimes dropping to 12.5% at  $k = 1000$ . Scarce data lead to poor enforcement of target updates ( $\mathcal{Q}^+$ ) while undermining preservation ( $\mathcal{Q}^-$ ). Figure 6b shows OOD accuracy declining from 63.0% ( $k = 1$ ) to 61.6% ( $k = 1000$ ), indicating drift risks. Un-

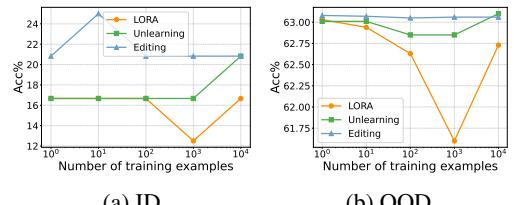
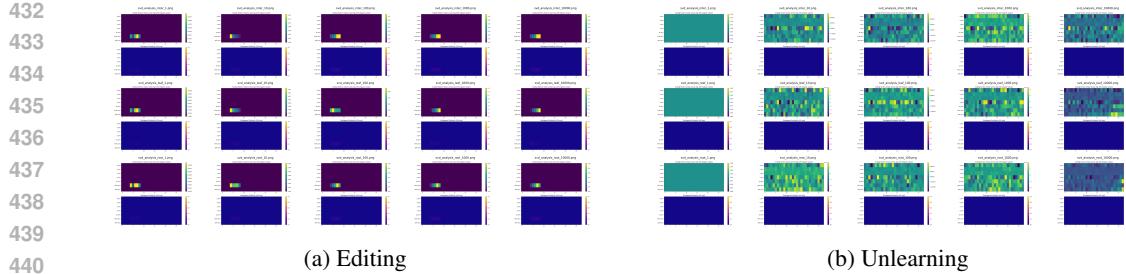


Figure 6: LoRA, Editing and Unlearning.



(a) Editing

(b) Unlearning

Figure 7: SVD-based geometric analysis of interventions. (a) Editing adjusts knowledge by gently rotating and slightly rescaling the representation space, preserving overall geometry while redirecting specific directions. (b) Unlearning, in contrast, acts by shrinking certain dimensions more aggressively, reducing the model’s capacity in those directions rather than rotating them.

learning remains stable around 63%, preserving prior knowledge but limiting target success. Editing combines stability with low-data efficiency, boosting ID accuracy to 25% at  $k = 10$  compared to 16.7% for LoRA and unlearning. In summary, editing balances new knowledge integration and preservation, LoRA risks drift, and unlearning is conservative but stable, explaining why we prefer editing/unlearning for continual updates.

### 5.5 FAILURE MODE AND STRESS TESTING

Existing studies describe errors such as incomplete forgetting or knowledge pollution in a fragmented way, without systematically characterizing the underlying mechanisms. Through our experiments on **open-ended question answering**, we observed that models fail for different reasons under editing and unlearning interventions. To capture these patterns, we propose a **Unified Failure Mode Taxonomy** that organizes observed errors into six categories (examples of each type in Appendix D.2):

under-forgetting (RR), over-spreading (CCR), conflict emergence (contradictions between updated and related knowledge), knowledge drift (performance degradation on unrelated tasks), instruction-following drop (reduced ability to follow complex instructions), and hallucination increase.

Stress-testing evaluates the failure modes with open generation tasks, making the model show practical robustness and use gpt-4o to evaluate. Our results show that hallucination (evaluated on TruthfulQA (Lin et al., 2022)) remains stable, instruction-following (open generation) drops moderately, and CoT reasoning can improve edit generalization but may increase residual knowledge, complicating unlearning (details in Appendix C). **Sequential update experiments, reported in Appendix F, further illustrate how multiple consecutive edits affect these behaviors and highlight potential cumulative effects on residual knowledge.**

### 5.6 THEORETICAL ANALYSIS

Our theoretical perspective connects the observed behaviors of editing and unlearning to their geometric effects on model representations. Let  $W \in \mathbb{R}^{m \times n}$  denote a parameter matrix (e.g., attention or MLP projection), with singular value decomposition  $W = U\Sigma V^\top$ . An intervention updates  $W$  to  $W' = U'\Sigma'V'^\top$ . The difference between  $W$  and  $W'$  can be decomposed into two interpretable components:

- **Scaling effects.** Changes in singular values  $\Sigma'/\Sigma$  indicate amplification or attenuation of certain representational directions.
- **Rotational effects.** Differences in subspaces  $\text{span}(U, V)$  vs.  $\text{span}(U', V')$  reflect reorientation of features while preserving their magnitude.

**Editing as local rotation with mild rescaling.** As shown in Figure 7a, editing primarily induces moderate rescaling of singular values while maintaining high orthogonal similarity between  $(U, V)$  and  $(U', V')$  across layers. This implies that editing preserves most of the representational geometry,

Table 2: Percentage (%) of observed failures in editing and unlearning.

Failure Mode	Editing	Unlearning
Under-forgetting (RR)	20	35
Over-spreading (CCR)	35	15
Conflict emergence	30	12
Knowledge drift	18	10
Instruction-following drop	22	18
Hallucination increase	5	4

486 redirecting specific factual directions through controlled rotations. Consequently, editing behaves  
 487 like a *rotation-plus-scaling operator*: it reallocates emphasis toward new factual associations while  
 488 retaining global coherence. This explains why editing achieves strong local enforcement but often  
 489 *over-spreads* changes to nearby branches (high CCR in Section 5.2).

490 **Unlearning as anisotropic scaling.** By contrast, Figure 7b shows that unlearning produces sharper  
 491 downscaling of singular values, with less stable alignment of  $U, V$  across layers. This indicates sup-  
 492 pression of capacity in certain subspaces rather than a simple rotation. Thus, unlearning resembles  
 493 an *attenuation operator*: it removes the ability to encode certain directions but does not reliably ro-  
 494 late them into new ones. This mechanism aligns with the observed *under-spreading* behavior (high  
 495 RR in Section 5.2), where forgetting remains localized and fails to propagate fully across related  
 496 nodes.

497 **Hierarchy-dependent dynamics.** Leaf-level interventions concentrate changes in later layers, sup-  
 498 porting near-perfect local adaptation. Root-level interventions require distributed rotations and scal-  
 499 ings across the network, introducing stricter ceilings on achievable accuracy. Intermediate nodes  
 500 combine aspects of both. These theoretical patterns mirror our empirical findings on branch-  
 501 dependent plasticity limits (Section 5.2.2).

## 502 5.7 DISCUSSION

503 Our findings offer several potential directions for future research. (1) *Model updating*: Updates  
 504 should employ dynamic, hierarchical control such as level- and relation-aware algorithms. Branch-  
 505 specific strategies can also improve effectiveness: for leaf nodes, updates can use more data for  
 506 higher accuracy, while root nodes may require less data. Data size should be carefully calibrated for  
 507 global consistency. Moreover, models exhibit subject-dependent sensitivity, hence, update methods  
 508 should account for differences across domains. (2) *Evaluation metrics*: The conflict rate offers  
 509 a more nuanced assessment of models, capturing hidden inconsistencies and ensuring that updates  
 510 improve the model more holistically rather than just for specific tasks. This mirrors human reasoning  
 511 in the sense that humans also monitor for contradictions and coherence, but the analogy is descriptive  
 512 rather than mechanistic. (3) *Foundation models*: Future models could be designed with layer-wise or  
 513 tensor-wise modularity, enabling finer-grained control when applying updates. By building update-  
 514 friendly architectures, such models would allow interventions to target specific branches or layers  
 515 more effectively, improving both efficiency and consistency of knowledge updates.

516 Our work has several limitations. First, our experiments are based on four domains due to limited  
 517 compute budget and could be expanded to more domains and multimodal models. Second, our  
 518 unified framework does not give theoretical bound for propagation and consistency remains open.  
 519 Third, the analysis is based on recent editing and unlearning approaches, which could be extended  
 520 to other algorithms to gain more insights.

## 522 6 CONCLUSION

524 We introduced KnowledgeSmith to understand the knowledge updating mechanism in LLMs by  
 525 unifying editing and unlearning. Our experiments highlight fundamental trade-offs, e.g., unlearning  
 526 prioritizes stability and efficiency but yields modest enforcement, while editing enforces knowledge  
 527 updates more effectively at the risk of destabilization and higher computational cost. We hope our  
 528 benchmark and analysis can shed light on future research on LLM knowledge updating.

529 Future research will investigate hybrid datasets that combine information across all knowledge graph  
 530 levels and domains to better guide LLM updates. We also aim to develop adaptive and hybrid  
 531 strategies that leverage internal model representations to dynamically determine when and how to  
 532 apply editing or unlearning.

## 534 535 ETHICAL AND REPRODUCIBILITY STATEMENT

### 536 ETHICS STATEMENT

537 This work investigates knowledge editing and unlearning in large language models with the goal of  
 538 improving our understanding of how models update and forget factual information. Our experiments

540 are restricted to controlled benchmarks, including publicly available datasets and synthetic data  
 541 that we release. We do not use sensitive, private, or personally identifiable information. While  
 542 the methods studied could, in principle, be misused to manipulate model knowledge for harmful  
 543 purposes, our intention is purely scientific, and we have limited our scope to safe, non-sensitive  
 544 settings. All pretrained models used in this study are publicly available and used in accordance with  
 545 their licenses. We believe our work contributes to safer, more transparent, and more responsible  
 546 approaches to model editing and unlearning.

547

#### 548 REPRODUCIBILITY STATEMENT

549

550 We have made every effort to ensure the reproducibility of our results. All datasets used are publicly  
 551 available or synthetically generated; details of dataset construction, splits, and preprocessing are  
 552 provided in Appendix A. Model architectures, and evaluation metrics are fully described. Our imple-  
 553 mentation builds on open-source frameworks (e.g., PyTorch, HuggingFace Transformers, vLLM), and we  
 554 will release the configuration files and synthetic benchmark data upon publication.

555

#### 556 REFERENCES

557

558 Charith Chandra Sai Balne, Sreyoshi Bhaduri, Tamoghna Roy, Vinija Jain, and Aman Chadha.  
 559 Parameter efficient fine tuning: A comprehensive analysis across applications, 2024. URL  
 560 <https://arxiv.org/abs/2404.13506>.

561 Thomas Baumhauer, Pascal Schöttle, and Matthias Zeppelzauer. Machine unlearning: Linear filtra-  
 562 tion for logit-based classifiers. *Machine Learning*, 111(9):3203–3226, 2022.

563 Nicola De Cao, Wilker Aziz, and Ivan Titov. Editing factual knowledge in language models, 2021.  
 564 URL <https://arxiv.org/abs/2104.08164>.

565 Qizhou Chen, Taolin Zhang, Xiaofeng He, Dongyang Li, Chengyu Wang, Longtao Huang, and  
 566 Hui Xue'. Lifelong knowledge editing for LLMs with retrieval-augmented continuous prompt  
 567 learning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the*  
 568 *2024 Conference on Empirical Methods in Natural Language Processing*, pp. 13565–13580, Mi-  
 569 ami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/  
 570 v1/2024.emnlp-main.751. URL [https://aclanthology.org/2024.emnlp-main.](https://aclanthology.org/2024.emnlp-main.751)  
 571 751/.

572 DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,  
 573 2025. URL <https://arxiv.org/abs/2501.12948>.

574 Gal Elidan, Ian McGraw, and Daphne Koller. Residual belief propagation: Informed scheduling for  
 575 asynchronous message passing. *arXiv preprint arXiv:1206.6837*, 2012.

576 Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Shi Jie, Xiang Wang, Xiangnan He, and  
 577 Tat seng Chua. Alphaedit: Null-space constrained knowledge editing for language models, 2025.  
 578 URL <https://arxiv.org/abs/2410.02355>.

579 Giorgos Flouris, Dimitris Manakanatas, Haridimos Kondylakis, Dimitris Plexousakis, and Grigoris  
 580 Antoniou. Ontology change: classification and survey. *The Knowledge Engineering Review*, 23  
 581 (2):117–152, 2008.

582 Zorik Gekhman, Gal Yona, Roee Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan  
 583 Herzig. Does fine-tuning llms on new knowledge encourage hallucinations?, 2024. URL <https://arxiv.org/abs/2405.05904>.

584 Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data  
 585 deletion in machine learning. *Advances in neural information processing systems*, 32, 2019.

586 Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Forgetting outside the box: Scrubbing  
 587 deep networks of information accessible from input-output observations. In *European Conference*  
 588 *on Computer Vision*, pp. 383–398. Springer, 2020.

594 Quanquan Gu, Zhenhui Li, and Jiawei Han. Generalized fisher score for feature selection. *arXiv*  
 595 *preprint arXiv:1202.3725*, 2012.  
 596

597 Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal  
 598 from machine learning models. *arXiv preprint arXiv:1911.03030*, 2019.  
 599

600 Akshat Gupta, Anurag Rao, and Gopala Anumanchipalli. Model editing at scale leads to gradual  
 601 and catastrophic forgetting, 2024. URL <https://arxiv.org/abs/2401.07453>.  
 602

603 Thomas Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi.  
 604 Aging with grace: Lifelong model editing with discrete key-value adaptors. In *Advances in Neural*  
 605 *Information Processing Systems*, 2023.  
 606

607 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Ja-  
 608 cob Steinhardt. Measuring massive multitask language understanding, 2021. URL <https://arxiv.org/abs/2009.03300>.  
 609

610 Yihuai Hong, Yuelin Zou, Lijie Hu, Ziqian Zeng, Di Wang, and Haiqin Yang. Dissecting fine-  
 611 tuning unlearning in large language models, 2024. URL <https://arxiv.org/abs/2410.06606>.  
 612

613 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 614 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.  
 615

616 Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. Approximate data deletion  
 617 from machine learning models. In *International conference on artificial intelligence and statistics*,  
 618 pp. 2008–2016. PMLR, 2021.  
 619

620 Matthew Jagielski, Om Thakkar, Florian Tramer, Daphne Ippolito, Katherine Lee, Nicholas Carlini,  
 621 Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, et al. Measuring forgetting  
 622 of memorized training examples. *arXiv preprint arXiv:2207.00099*, 2022.  
 623

624 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chap-  
 625 lot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,  
 626 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril,  
 627 Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.  
 628

629 Dahyun Jung, Jaehyung Seo, Jaewook Lee, Chanjun Park, and Heuiseok Lim. Come: An  
 630 unlearning-based approach to conflict-free model editing. *arXiv preprint arXiv:2502.15826*,  
 631 2025.  
 632

633 Kabir Khan, Priya Sharma, Arjun Mehta, Neha Gupta, and Ravi Narayanan. Dysk-attn: A frame-  
 634 work for efficient, real-time knowledge updating in large language models via dynamic sparse  
 635 knowledge attention, 2025. URL <https://arxiv.org/abs/2508.07185>.  
 636

637 Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural  
 638 network representations revisited. In *International conference on machine learning*, pp. 3519–  
 639 3529. PMLR, 2019.  
 640

641 Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. Zero-shot relation extraction via  
 642 reading comprehension. *arXiv preprint arXiv:1706.04115*, 2017.  
 643

644 Xiang Li, Wenqi Wei, and Bhavani Thuraisingham. Mubox: A critical evaluation framework of  
 645 deep machine unlearning [systematization of knowledge paper]. In *Proceedings of the 30th ACM*  
 646 *Symposium on Access Control Models and Technologies*, pp. 175–188, 2025a.  
 647

648 Zexi Li, Xiangzhu Wang, William F Shen, Meghdad Kurmanji, Xinchi Qiu, Dongqi Cai, Chao Wu,  
 649 and Nicholas D Lane. Editing as unlearning: Are knowledge editing methods strong baselines for  
 650 large language model unlearning? *arXiv preprint arXiv:2505.19855*, 2025b.  
 651

652 Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human  
 653 falsehoods, 2022. URL <https://arxiv.org/abs/2109.07958>.  
 654

648 Lihui Liu, Zihao Wang, Jiaxin Bai, Yangqiu Song, and Hanghang Tong. New frontiers of knowledge  
 649 graph reasoning: Recent advances and future trends. In *Companion Proceedings of the ACM Web*  
 650 *Conference 2024*, pp. 1294–1297, 2024.

651 Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. Tofu: A task  
 652 of fictitious unlearning for llms, 2024.

653 Shirley Mark, Rani Moran, Thomas Parr, Steve W Kennerley, and Timothy EJ Behrens. Transferring  
 654 structural knowledge across cognitive maps in humans and models. *Nature communications*, 11  
 655 (1):4783, 2020.

656 Elan Markowitz, Anil Ramakrishna, Ninareh Mehrabi, Charith Peris, Rahul Gupta, Kai-Wei Chang,  
 657 and Aram Galstyan. K-edit: Language model editing with contextual knowledge awareness, 2025.  
 658 URL <https://arxiv.org/abs/2502.10626>.

659 Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual  
 660 associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022a.

661 Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-editing  
 662 memory in a transformer. *arXiv preprint arXiv:2210.07229*, 2022b.

663 Meta. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

664 Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model  
 665 editing at scale. *arXiv preprint arXiv:2110.11309*, 2021.

666 Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. Memory-  
 667 based model editing at scale. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba  
 668 Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Con-  
 669 ference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*,  
 670 pp. 15817–15831. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/mitchell22a.html>.

671 Andrew Y Ng. Feature selection, l1 vs. l2 regularization, and rotational invariance. In *Proceedings  
 672 of the twenty-first international conference on Machine learning*, pp. 78, 2004.

673 Oded Ovadia, Menachem Brief, Moshik Mishaeli, and Oren Elisha. Fine-tuning or retrieval? com-  
 674 paring knowledge injection in llms, 2024. URL <https://arxiv.org/abs/2312.05934>.

675 Haewon Park, Gyubin Choi, Minjun Kim, and Yohan Jo. Context-robust knowledge editing for  
 676 language models, 2025. URL <https://arxiv.org/abs/2505.23026>.

677 Heiko Paulheim. Knowledge graph refinement: A survey of approaches and evaluation methods.  
 678 *Semantic web*, 8(3):489–508, 2016.

679 Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models  
 680 as few-shot unlearners. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller,  
 681 Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st Inter-  
 682 national Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning  
 683 Research*, pp. 40034–40050. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/pawelczyk24a.html>.

684 Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao  
 685 Liu, Luke Zettlemoyer, Noah A Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way  
 686 evaluation for language models. *CoRR*, 2024.

687 Anton Sinitzin, Vsevolod Plokhotnyuk, Dmitriy Pyrkin, Sergei Popov, and Artem Babenko. Editable  
 688 neural networks, 2020. URL <https://arxiv.org/abs/2004.00345>.

689 Gemma Team. Gemma: Open models based on gemini research and technology, 2024. URL  
 690 <https://arxiv.org/abs/2403.08295>.

691 Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025a. URL  
 692 <https://qwenlm.github.io/blog/qwq-32b/>.

702 Qwen Team. Qwen3 technical report, 2025b. URL <https://arxiv.org/abs/2505.09388>.  
 703

704 Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, and Tom Hartvigsen. Understanding  
 705 the limits of lifelong knowledge editing in llms, 2025. URL <https://arxiv.org/abs/2503.05683>.  
 706

707 Anvith Thudi, Gabriel Deza, Varun Chandrasekaran, and Nicolas Papernot. Unrolling sgd: Un-  
 708 derstanding factors influencing machine unlearning. In *2022 IEEE 7th European Symposium on  
 709 Security and Privacy (EuroS&P)*, pp. 303–319. IEEE, 2022.

710 Peng Wang, Zexi Li, Ningyu Zhang, Ziwen Xu, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang,  
 711 and Huajun Chen. Wise: Rethinking the knowledge memory for lifelong model editing of large  
 712 language models. *arXiv preprint arXiv:2405.14768*, 2024.

713 Zihao Wei, Liang Pang, Hanxing Ding, Jingcheng Deng, Huawei Shen, and Xueqi Cheng. Stable  
 714 knowledge editing in large language models, 2024. URL <https://arxiv.org/abs/2402.13048>.  
 715

716 Haoming Xu, Ningyuan Zhao, Liming Yang, Sendong Zhao, Shumin Deng, Mengru Wang, Bryan  
 717 Hooi, Nay Oo, Huajun Chen, and Ningyu Zhang. Relearn: Unlearning via learning for large  
 718 language models. *arXiv preprint arXiv:2502.11190*, 2025.

719 Wenli Yang, Lillian Some, Michael Bain, and Byeong Kang. A comprehensive survey on integrating  
 720 large language models with knowledge-based methods. *Knowledge-Based Systems*, pp. 113503,  
 721 2025.

722 Jin Yao, Eli Chien, Minxin Du, Xinyao Niu, Tianhao Wang, Zezhou Cheng, and Xiang Yue. Ma-  
 723 chine unlearning of pre-trained large language models. In Lun-Wei Ku, Andre Martins, and  
 724 Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Com-  
 725 putational Linguistics (Volume 1: Long Papers)*, pp. 8403–8419, Bangkok, Thailand, August  
 726 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.457. URL  
 727 <https://aclanthology.org/2024.acl-long.457/>.  
 728

729 Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen,  
 730 and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities. In  
 731 Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on  
 732 Empirical Methods in Natural Language Processing*, pp. 10222–10240, Singapore, December  
 733 2023a. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.632. URL  
 734 <https://aclanthology.org/2023.emnlp-main.632/>.  
 735

736 Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen,  
 737 and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities. *arXiv  
 738 preprint arXiv:2305.13172*, 2023b.

739 Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi,  
 740 Shengyu Mao, Jintian Zhang, Yuansheng Ni, et al. A comprehensive study of knowledge editing  
 741 for large language models. *arXiv preprint arXiv:2401.01286*, 2024.

742 Ruixiang Zhang, Shuangfei Zhai, Eta Littwin, and Josh Susskind. Learning representation from  
 743 neural fisher kernel with low-rank approximation. *arXiv preprint arXiv:2202.01944*, 2022.

744

745 Zihan Zhang, Meng Fang, Ling Chen, Mohammad-Reza Namazi-Rad, and Jun Wang. How do large  
 746 language models capture the ever-changing world knowledge? a review of recent advances. *arXiv  
 747 preprint arXiv:2310.07343*, 2023.

748 Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. Can  
 749 we edit factual knowledge by in-context learning? *arXiv preprint arXiv:2305.12740*, 2023.

750

751 Zexuan Zhong, Zhengxuan Wu, Christopher D Manning, Christopher Potts, and Danqi Chen.  
 752 Mquake: Assessing knowledge editing in language models via multi-hop questions. *arXiv  
 753 preprint arXiv:2305.14795*, 2023.

754 Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix Yu, and  
 755 Sanjiv Kumar. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*,  
 2020.

756 **Appendix**  
 757 **KnowledgeSmith: Uncovering Knowledge Updating in LLMs with Model**  
 758 **Editing and Unlearning**

760 **CONTENTS**

764 <b>A Data Generation Example and Pipeline</b>	<b>16</b>
765    A.1 Knowledge Point and Knowledge Graph (KG) . . . . .	16
766    A.2 Template Generation . . . . .	16
767    A.3 Prompting GPT for Question Generation . . . . .	16
768    A.4 Probe Types . . . . .	18
771    A.5 Multiple-Choice Formatting and Data Records . . . . .	19
772    A.6 Quality Control . . . . .	19
773    A.7 Domain and Sample Granularity . . . . .	19
775	
776 <b>B Propagation Asymmetry Metrics and Algorithm</b>	<b>19</b>
777	
778 <b>C Stress Testing</b>	<b>20</b>
779	
780 <b>D Robustness and Failure Mode</b>	<b>20</b>
781    D.1 Adversarial Robustness Analysis . . . . .	20
783    D.2 Failure Mode Examples . . . . .	21
784	
785 <b>E Additional Methods</b>	<b>22</b>
786    E.1 Unlearning . . . . .	22
788    E.2 Editing . . . . .	22
789	
790 <b>F Sequential Update</b>	<b>23</b>
791	
792 <b>G Accuracy Result</b>	<b>24</b>
794	
795 <b>H Model Similarity Result</b>	<b>24</b>
796	
797 <b>I LLM usage</b>	<b>24</b>
798	
799	
800	
801	
802	
803	
804	
805	
806	
807	
808	
809	

810 A DATA GENERATION EXAMPLE AND PIPELINE  
811812 To make our pipeline transparent, we provide an end-to-end example showing how a single knowl-  
813 edge point expands into a large set of evaluation items, emphasizing hierarchical structure and con-  
814 trolled fact editing.  
815816 A.1 KNOWLEDGE POINT AND KNOWLEDGE GRAPH (KG)  
817818 We illustrate how a knowledge point can be represented as a triple and anchored at different levels  
819 of the knowledge graph. Table 3 shows one example from each domain.  
820821 Table 3: Examples of knowledge triples and anchoring across different levels of the KG hierarchy.  
822

Domain	Example Triple	KG (Root → Intermediate → Leaf)
Biology	(DNA double helix, discovered.in, 1953)	Root: concept of DNA structure → role in molecular biology and genetics → link to genetics/medicine/biotech applications
Economics	(Phillips curve, describes, inflation-unemployment relationship)	Root: economic trade-offs → macroeconomic models of inflation and unemployment → policy debates on stagflation and monetary policy
History	(Declaration of Independence, signed.in, 1776)	Root: revolutions and independence movements → American Revolutionary era → specific events such as the Continental Congress or early U.S. governance
Physics	(Theory of General Relativity, published.in, 1915)	Root: fundamental physics theories → spacetime and gravitation framework → applications such as black holes, gravitational waves, or GPS corrections

831 This fact is anchored at three levels of the knowledge graph:  
832833

- **Root:** broad, domain-level understanding.
- **Intermediate:** contextual understanding, including its role and implications.
- **Leaf:** fine-grained, specific questions.

837 A.2 TEMPLATE GENERATION  
838839 For the selected fact, we generate multiple question templates per KG level, capturing different  
840 aspects of the fact (definition, role, context, and application).  
841842

- **Root-level templates:** Broad factual or conceptual questions.
- **Intermediate-level templates:** Questions about domain implications, causal relationships, and contextual applications.
- **Leaf-level templates:** Specific, field-dependent scenarios where the fact influences outcomes or knowledge in that domain.

849 An example of generated templates is shown in Table 4, where leaf-level templates are instantiated  
850 with different fields (e.g., genetics, medicine).  
851852 A.3 PROMPTING GPT FOR QUESTION GENERATION  
853854 Our pipeline for generating evaluation questions follows these steps:  
855856

1. **Knowledge Graph Generation:** GPT is prompted to generate a structured KG for the target domain. Nodes represent root, intermediate, and leaf-level knowledge.
2. **Fact Selection:** From the KG, a single fact is selected (e.g., (DNA double helix, discovered.in, 1953)) to anchor all subsequent questions.
3. **Template Generation:** GPT is prompted to produce multiple templated question forms surrounding the fact. Templates vary in phrasing, style, and emphasis, covering definition, context, role, and applications.
4. **Level-Specific Question Generation:** Each template is input to GPT with instructions specifying the desired KG level (root, intermediate, leaf). Example prompts:

Table 4: QA templates for four knowledge points across Biology, Economics, History, and Physics.

Level	Biology: <i>DNA double helix</i>	Economics: <i>Phillips curve</i>	History: <i>Declaration of Independence (1776)</i>	Physics: <i>General Relativity (1915)</i>
Root-level	<p>What is the DNA double helix?</p> <p>Who discovered the DNA double helix?</p> <p>When was the DNA double helix discovered?</p> <p>What does the DNA double helix describe?</p> <p>Why is the DNA double helix important in biology?</p> <p>What shape is the DNA double helix?</p> <p>What was learned from the DNA double helix?</p> <p>Which scientists worked on the DNA double helix?</p>	<p>What is the Phillips curve?</p> <p>What relationship does the Phillips curve describe?</p> <p>Who proposed the Phillips curve?</p> <p>When was the Phillips curve introduced?</p> <p>Why is the Phillips curve important in economics?</p> <p>How is the Phillips curve used in macroeconomics?</p> <p>What does the Phillips curve imply about inflation and unemployment?</p> <p>Which countries have applied the Phillips curve concept?</p>	<p>What is the Declaration of Independence?</p> <p>When was the Declaration of Independence signed?</p> <p>Who signed the Declaration of Independence?</p> <p>Why was the Declaration of Independence created?</p> <p>What does the Declaration of Independence proclaim?</p> <p>Which country declared independence in 1776?</p> <p>What historical context led to the Declaration of Independence?</p> <p>Why is the Declaration of Independence important in history?</p>	<p>What is the Theory of General Relativity?</p> <p>Who proposed the Theory of General Relativity?</p> <p>When was the Theory of General Relativity published?</p> <p>Why is the Theory of General Relativity important?</p> <p>What does the Theory of General Relativity describe?</p> <p>How does General Relativity differ from Newtonian physics?</p> <p>What are the key concepts in General Relativity?</p> <p>Which experiments confirmed General Relativity?</p>
Intermediate	<p>How did the DNA double helix change molecular biology?</p> <p>What discoveries followed the DNA double helix?</p> <p>What role did the DNA double helix play in genetics?</p> <p>How did the DNA double helix influence medical research?</p> <p>What techniques confirmed the DNA double helix?</p> <p>How is the DNA double helix taught in schools?</p> <p>What reaction did scientists have to the DNA double helix?</p> <p>How did the DNA double helix affect other fields of science?</p>	<p>How does the Phillips curve affect monetary policy?</p> <p>What criticisms exist for the Phillips curve?</p> <p>How did the Phillips curve shape economic thought?</p> <p>How does the Phillips curve relate to inflation targeting?</p> <p>What data supports or contradicts the Phillips curve?</p> <p>How do economists interpret the Phillips curve over time?</p> <p>How does the Phillips curve influence labor market policies?</p> <p>How is the Phillips curve taught in universities?</p>	<p>How did the Declaration of Independence influence the American Revolution?</p> <p>What ideas from the Enlightenment are in the Declaration?</p> <p>How did other countries react to the Declaration?</p> <p>What role did the Declaration play in forming the U.S. government?</p> <p>How was the Declaration received by the British crown?</p> <p>What debates occurred during the drafting of the Declaration?</p> <p>How did the Declaration impact colonial society?</p> <p>How is the Declaration taught in schools?</p>	<p>How did General Relativity influence modern physics?</p> <p>What role does General Relativity play in cosmology?</p> <p>How does General Relativity explain gravity?</p> <p>How was General Relativity received by the scientific community?</p> <p>How does General Relativity relate to black holes?</p> <p>How is General Relativity taught in universities?</p> <p>What mathematical tools are used in General Relativity?</p> <p>How does General Relativity affect GPS technology?</p>
Leaf-level	<p>How did the DNA double helix influence research in genetics?</p> <p>What impact did the DNA double helix have in medicine?</p> <p>How was forensic science affected by the DNA double helix?</p> <p>In evolutionary biology, what role did the DNA double helix play?</p> <p>Why did biotechnology change after the DNA double helix?</p> <p>What does public health owe to the DNA double helix?</p> <p>How did the DNA double helix influence research in anthropology?</p> <p>What impact did the DNA double helix have in bioinformatics?</p> <p>How was drug development affected by the DNA double helix?</p> <p>In agriculture, what role did the DNA double helix play?</p>	<p>How does the Phillips curve explain stagflation in the 1970s?</p> <p>How did the Phillips curve influence central bank decisions?</p> <p>How is unemployment measured in relation to the Phillips curve?</p> <p>What role did the Phillips curve play in New Keynesian economics?</p> <p>How do different countries' experiences validate the Phillips curve?</p> <p>What empirical models are used to test the Phillips curve?</p> <p>How does the Phillips curve relate to wage inflation?</p> <p>How did the Phillips curve inform fiscal policy during recessions?</p> <p>How is the Phillips curve applied in modern macroeconomic forecasting?</p> <p>How does the Phillips curve interact with supply shocks?</p>	<p>Which founding fathers were key authors of the Declaration?</p> <p>How did the Declaration affect slavery debates in the U.S.?</p> <p>What role did the Declaration play in the Revolutionary War?</p> <p>How were the colonies mobilized after the Declaration?</p> <p>How did newspapers and pamphlets spread the Declaration?</p> <p>What influence did the Declaration have on other independence movements?</p> <p>How did international law view the Declaration at the time?</p> <p>How did the Declaration inspire subsequent U.S. legislation?</p> <p>How did the Declaration affect Native American relations?</p> <p>How did the Declaration shape early U.S. political parties?</p>	<p>How did General Relativity predict the bending of light?</p> <p>How was General Relativity confirmed during the 1919 solar eclipse?</p> <p>How does General Relativity influence gravitational wave research?</p> <p>How did General Relativity impact quantum theory?</p> <p>How does General Relativity affect modern cosmological models?</p> <p>How do black hole studies rely on General Relativity?</p> <p>How does General Relativity explain time dilation near massive objects?</p> <p>How did General Relativity change our understanding of space-time?</p> <p>How does General Relativity relate to the expansion of the universe?</p> <p>How are relativistic effects measured in particle accelerators?</p>

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971**Root-level Prompt**

Knowledge fact: “DNA double helix is a fundamental concept in molecular biology.”  
 Generate 3 multiple-choice questions targeting broad, domain-level understanding (root-level). Each question should have 4 answer options (A, B, C, D), one correct answer, and 3 plausible distractors.

**Intermediate-level Prompt**

Knowledge fact: “DNA double helix discovery influenced the field of genetics.”  
 Generate 3 multiple-choice questions targeting intermediate-level understanding using the same format.

**Leaf-level Prompt**

Knowledge fact: “DNA double helix was discovered in 1953 by Watson and Crick.”  
 Generate 3 multiple-choice questions targeting leaf-level understanding (specific facts). Ensure 4 answer options, one correct answer, and 3 plausible distractors.

**A.4 PROBE TYPES**

From each generated question template, we derive six probe types to evaluate different aspects of model behavior:

- **Direct Probe:** Queries the target fact in its canonical direction.
- **Reverse Probe:** Queries the fact in the inverted relation to test bidirectional consistency.
- **Multi-hop Probe:** Tests knowledge propagation by asking indirectly via intermediate nodes.
- **Contextual Probe:** Embeds the fact in a rich or distractor-laden context.
- **Conflict Probe:** Presents contradictory or competing information to assess resolution.
- **Comparison Probe:** Forces a choice between multiple candidates to evaluate selective updating.

Example prompts for the four subjects are shown in Table 5.

Table 5: Example probes across four subject domains, illustrating six probe types.

Subject	Example Probes
Biology (DNA double helix)	<b>Direct:</b> When was the DNA double helix discovered? <b>Reverse:</b> Which molecule’s structure was determined in 1953 as a double helix? <b>Multi-hop:</b> Who were the key scientists whose discovery of the DNA structure influenced modern genetics? <b>Contextual:</b> The DNA double helix discovery transformed molecular biology. In which year was this breakthrough made? <b>Conflict:</b> Some sources claim 1952, others 1953. Which year is correct? <b>Comparison:</b> Was the DNA double helix discovered in 1953 or 1955?
Economics (Phillips curve)	<b>Direct:</b> What relationship does the Phillips curve describe? <b>Reverse:</b> Which economic principle captures the link between inflation and unemployment? <b>Multi-hop:</b> Which macroeconomic models rely on understanding the inflation-unemployment trade-off? <b>Contextual:</b> The Phillips curve has shaped monetary policy debates. What relationship does it represent? <b>Conflict:</b> Some argue it holds only short-term, others claim long-term relevance. Which is correct? <b>Comparison:</b> Does the Phillips curve describe inflation-unemployment or wage-productivity trade-offs?
History (Declaration of Independence)	<b>Direct:</b> In what year was the Declaration of Independence signed? <b>Reverse:</b> Which historical document was signed in 1776? <b>Multi-hop:</b> Which events or congresses led to the signing of the Declaration? <b>Contextual:</b> Amid the Revolutionary era, the Declaration was signed. Which year did this occur? <b>Conflict:</b> Some accounts state July 2, others July 4. Which is correct? <b>Comparison:</b> Was the Declaration signed in 1776 or 1777?
Physics (General Relativity)	<b>Direct:</b> In what year did Einstein publish the theory of General Relativity? <b>Reverse:</b> Which scientist published General Relativity in 1915? <b>Multi-hop:</b> Which subsequent physics phenomena were explained following Einstein’s publication? <b>Contextual:</b> General Relativity transformed our understanding of space-time. When was it published? <b>Conflict:</b> Some sources claim 1915, others 1916. Which is correct? <b>Comparison:</b> Did Einstein publish General Relativity in 1915 or 1920?

972 A.5 MULTIPLE-CHOICE FORMATTING AND DATA RECORDS  
973974 All probes are formatted as four-choice QA items consistent with MMLU. Distractors are created  
975 via entity substitution and paraphrasing. An example for the four subjects is shown in Table 6  
976977 Table 6: Compact multiple-choice probes across four subjects. Correct answers indicated.  
978

Subject	Example Multiple Choice
Biology (DNA double helix)	<b>Q:</b> When was the DNA double helix discovered? A. 1953 (Correct) B. 1955 C. 1962 D. 1947
Economics (Phillips curve)	<b>Q:</b> What relationship does the Phillips curve describe? A. Inflation vs. unemployment (Correct) B. Wage vs. productivity C. Interest rate vs. investment D. Savings vs. consumption
History (Declaration of Independence)	<b>Q:</b> In what year was the Declaration of Independence signed? A. 1776 (Correct) B. 1775 C. 1777 D. 1781
Physics (General Relativity)	<b>Q:</b> In what year did Einstein publish the theory of General Relativity? A. 1915 (Correct) B. 1920 C. 1912 D. 1918

988 A.6 QUALITY CONTROL  
989

990 Items undergo:

991 1. Format validation (4 options, 1 correct answer)  
992 2. Factual validation against the KG  
993 3. Distractor validation (plausible yet incorrect)994 Manual spot checks ensure grammaticality and factual correctness; GPT-generated distractors are  
995 cross-checked with encyclopedic sources.  
996998 A.7 DOMAIN AND SAMPLE GRANULARITY  
9991000 Domains include **Biology**, **History**, **Physics**, and **Economics**, each curated into a structured KG.  
1001 Our study focuses on modifying one fact at a time; all QA items are anchored on this fact. Multiple  
1002 templates per node level, probe types, paraphrases, and varying data scales (1, 10, 100, 1,000,  
1003 10,000) allow a single fact to generate up to millions of QA items for large-scale evaluation.1004 B PROPAGATION ASYMMETRY METRICS AND ALGORITHM  
10051006 To quantify over- vs. under-spreading rigorously, we define:  
1007

1008 
$$\text{Collateral Change Ratio (CCR)} = \frac{1}{|\mathcal{Q}_{\text{related}}|} \sum_{x \in \mathcal{Q}_{\text{related}}} d(p_{\theta'}(\cdot | x), p_{\theta}(\cdot | x)), \quad (4)$$

1009 
$$\text{Residual Retention (RR)} = \frac{1}{|\mathcal{Q}_{\text{related}}|} \sum_{x \in \mathcal{Q}_{\text{related}}} \mathbf{1}[\hat{y}_{\theta'}(x) = y_{\theta}(x)], \quad (5)$$

1010 where  $\mathcal{Q}_{\text{related}}$  denotes structurally related probes,  $p_{\theta}$  and  $p_{\theta'}$  are predictions before and after intervention,  
1011 and  $d(\cdot, \cdot)$  is a distance metric (KL, label change, etc.).  
10121013 **Propagation Evaluation Algorithm:**  
10141015 1. Select a target node at hierarchy level  $L$ .  
1016 2. Apply editing or unlearning to the node.  
1017 3. Measure direct accuracy on target node ( $Acc_{\text{direct}}$ ).  
1018 4. Measure multi-hop accuracy on related nodes ( $Acc_{\text{multi-hop}}$ ).  
1019 5. Compute CCR and RR metrics:  
1020 

- Editing:  $1 - Acc_{\text{multi-hop}}$  as proxy for over-spreading.
- Unlearning:  $Acc_{\text{multi-hop}}$  as proxy for under-spreading.

  
1021 6. Repeat for all hierarchy levels and average over domains.  
1022

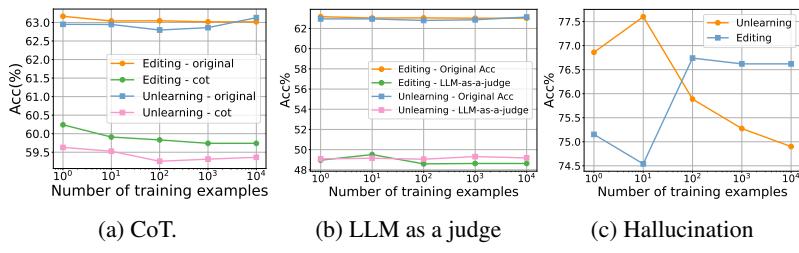
1026  
1027 C STRESS TESTING  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035

Figure 8: Stress testing.

We evaluate *instruction-following* ability (Figure 8b) and *hallucination* on the TruthfulQA (Lin et al., 2022) dataset (Figure 8c), testing whether the parameter update  $\theta \rightarrow \theta'$  preserves desired behavior when executing complex tasks. These evaluations provide a comprehensive view of how the unified framework constrains model updates, ensuring both local alignment with target distributions and global reliability across diverse scenarios.

For **hallucination**, the average accuracy across data scales for unlearning is 76.0%, and for editing is 76.1%, with standard deviations of 0.87 and 0.91 respectively. This indicates that **both editing and unlearning maintain stable performance** under hallucination tests, with no significant increase in spurious behavior.

For **instruction-following**, when measured using an LLM as a judge, editing accuracy drops from 63.0% (original) to 48.6% on average, while unlearning drops from 62.9% to 49.1%. Although the absolute difference is small, editing shows slightly larger variability (standard deviation 0.12%) compared to unlearning (0.10%). This suggests that **editing is more aggressive** in updating targeted knowledge but may slightly perturb complex reasoning tasks, whereas **unlearning better preserves general instruction-following ability**.

1054 D ROBUSTNESS AND FAILURE MODE  
1055  
1056  
1057

## D.1 ADVERSARIAL ROBUSTNESS ANALYSIS

To complement our main text results, we provide a detailed analysis of adversarial robustness for editing and unlearning interventions. Adversarial robustness is evaluated by exposing the model to deliberately misleading or deceptive probes, which combine unrelated or conflicting concepts. This stresses the model’s ability to maintain prior knowledge ( $Q^-$ ) while incorporating updates.

**Experimental Setup** We vary the number of training examples used for each intervention: 1, 10, 100, 1000, and 10,000. For each data scale, we measure two complementary performance metrics:

- **Original Accuracy:** The model’s performance on standard in-domain probes ( $Q^+$ ), reflecting whether the intended knowledge update was successfully incorporated without disrupting unrelated facts.
- **Adversarial Accuracy:** The model’s performance on *conflict probes*, which contain contradictory or misleading information. These probes test the model’s *robustness against adversarial perturbations*, i.e., whether it can resist adopting incorrect or conflicting knowledge while maintaining its updated and preserved facts.

By comparing original and adversarial accuracy across training scales and intervention types (editing vs. unlearning), we assess:

- The sensitivity of each method to misleading inputs.
- How stability and resistance to conflicts evolve as more examples are provided.
- Differences in trade-offs between aggressive updates (editing) and conservative updates (unlearning).

1080 This setup allows us to systematically quantify the *adversarial robustness* of interventions, linking  
 1081 conflict probe performance directly to practical model reliability under deceptive or contradictory  
 1082 inputs.  
 1083

1084 **Observations** Our observations are:

1085

- **Editing exhibits strong local updates but high adversarial sensitivity:** Original accuracy  
 1086 remains stable around 63% across all data scales. However, adversarial accuracy drops sharply  
 1087 from 36.7% at 1 example to 31.7% at 10,000 examples. This indicates that while editing  
 1088 successfully enforces target updates, it leaves models vulnerable to misleading inputs, with  
 1089 adversarial failure increasing slightly as data scale grows.
- **Unlearning maintains more stable adversarial performance:** Original accuracy is similar  
 1090 to editing. Adversarial accuracy remains relatively constant around 33–35%, showing that un-  
 1091 learning prioritizes preservation over aggressive enforcement, making the model less sensitive  
 1092 to adversarially constructed probes.
- **Trade-off between update intensity and robustness:** Comparing the two interventions, editing  
 1093 maximizes immediate factual incorporation at the cost of susceptibility to adversarial  
 1094 probes, whereas unlearning provides conservative updates that better preserve prior knowledge,  
 1095 yielding higher adversarial robustness.
- **Data scale effects:** Increasing the number of examples slightly improves adversarial robustness  
 1096 for unlearning (e.g., from 33.3% at 1 example to 34.8% at 1,000 examples), but the trend is less  
 1097 pronounced for editing. This suggests that adding more training data does not fully mitigate  
 1098 adversarial vulnerability for aggressive editing strategies.

1103 **Summary** These results reinforce the broader trade-offs observed in our main text. Editing  
 1104 achieves stronger local adaptation and in-domain gains, but adversarial robustness is compromised.  
 1105 Unlearning is more conservative, achieving lower immediate gains but maintaining stability under  
 1106 adversarial stress. Together, these findings highlight the importance of considering both factual  
 1107 enforcement and robustness when designing knowledge update strategies in LLMs.

## 1108 D.2 FAILURE MODE EXAMPLES

1109 We provide examples of failure mode for each subject as shown in Table 7.

1110  
 1111  
 1112 Table 7: Representative examples of each failure mode for the four studied subjects. Each subject is  
 1113 listed in a separate row for readability.  
 1114

1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 Subject	1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 Failure Mode	1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 Example
Biology (DNA)	Under-forgetting (RR)	DNA year remains 1953 after update to 1955
	Over-spreading (CCR)	DNA update changes RNA discovery year
	Conflict Emergence	DNA reported as 1953 and 1955
	Knowledge Drift	DNA update causes cell structure errors
	Instruction-Following Drop	Fails to explain multi-step DNA replication
	Hallucination Increase	Invents molecule “X-DNA”
Economics (Phillips curve)	Under-forgetting (RR)	Phillips curve still inflation-unemployment after update
	Over-spreading (CCR)	Phillips curve update alters Laffer curve
	Conflict Emergence	Links both inflation-unemployment and wages-productivity
	Knowledge Drift	Update mispredicts supply-demand
	Instruction-Following Drop	Misapplies multi-step economic policy reasoning
	Hallucination Increase	Fabricates fictional “Y-Index”
History (Declaration)	Under-forgetting (RR)	Declaration year still 1776 after update to 1777
	Over-spreading (CCR)	Declaration update changes Constitution year
	Conflict Emergence	Declaration signed 1776 and 1777
	Knowledge Drift	Update affects French Revolution facts
	Instruction-Following Drop	Struggles with chronological sequencing of events
	Hallucination Increase	Claims fake historical figure influenced Declaration
Physics (General Relativity)	Under-forgetting (RR)	GR year remains 1915 after update to 1920
	Over-spreading (CCR)	GR update changes Special Relativity year
	Conflict Emergence	GR dated 1915 and 1920
	Knowledge Drift	Update reduces quantum mechanics accuracy
	Instruction-Following Drop	Cannot solve multi-step relativity problems
	Hallucination Increase	Reports spurious physics law “Relativistic Thermodynamics Law”

1134 **E ADDITIONAL METHODS**  
11351136 To further validate the generality of the propagation asymmetry phenomena reported in the main  
1137 paper, we conducted an additional suite of experiments using multiple independent intervention  
1138 algorithms, spanning both editing and unlearning paradigms. These experiments were performed  
1139 on the same four subject domains (*biology, economic, physics, history*), and evaluated at root-,  
1140 intermediate-, and leaf-level nodes in our conceptual hierarchies.  
11411142 **E.1 UNLEARNING**  
11431144 We applied the gradient ascent method base on Tofu (Maini et al., 2024) framework with varying  
1145 numbers of updates across four subjects and multiple training set sizes. The results shown in Table 8  
1146 replicate the core findings presented in the main paper:  
11471148 

- propagation remains asymmetric across hierarchy levels,
- leaf nodes experience weaker upward transfer,
- root-level deletions continue to exhibit stronger downward effects.

  
11511152 Importantly, these consistency patterns persist regardless of the number of training examples and  
1153 irrespective of subject domain, suggesting that the structural behaviors we identified are not artifacts  
1154 of a particular unlearning implementation.  
1155

1156 <b>Subject</b>	1157 <b>Train Size</b>	1158 <b>Root</b>	1159 <b>Intermediate</b>	1160 <b>Leaf</b>
1158 biology	1	16.67	16.67	16.67
1159 biology	10	16.67	16.67	16.67
1160 biology	100	25.00	25.00	25.00
1161 biology	1000	29.17	25.00	16.67
1162 biology	10000	16.67	29.17	25.00
1163 economic	1	29.17	29.17	29.17
1164 economic	10	29.17	29.17	33.33
1165 economic	100	20.83	16.67	25.00
1166 economic	1000	37.50	37.50	29.17
1167 economic	10000	37.50	20.83	25.00
1168 physics	1	25.00	25.00	25.00
1169 physics	10	25.00	25.00	25.00
1170 physics	100	16.67	20.83	20.83
1171 physics	1000	12.50	16.67	20.83
1172 physics	10000	8.33	16.67	12.50
1173 history	1	16.67	16.67	16.67
1174 history	10	16.67	16.67	16.67
1175 history	100	12.50	16.67	12.50
1176 history	1000	4.17	8.33	0.00
1177 history	10000	0.00	12.50	0.00

1177 Table 8: Unlearning experiments using Tofu across domains and hierarchy levels.  
11781179 **E.2 EDITING**  
11801181 We also evaluated the MEND editing method (Mitchell et al., 2021) on the same corpus of subjects,  
1182 hierarchy depths, and training sizes. The results shown in Table 9 demonstrate that:  
11831184 

- editing accuracy follows the same hierarchy-dependent plasticity structure observed in the  
1185 main paper,
- root-level edits continue to propagate downward more strongly than bottom-up corrections,
- leaf nodes remain the easiest to modify reliably.

1188  
 1189 These findings reinforce that the asymmetry patterns we report are algorithm-agnostic, emerging  
 1190 from the structure of the knowledge graph itself rather than any specific intervention technique.

Dataset	Train Size	Root	Intermediate	Leaf
biology	1	35.2	32.7	42.1
biology	10	36.1	33.2	43.3
biology	100	37.6	34.4	44.7
biology	1000	38.9	35.1	46.2
biology	10000	20.3	18.7	40.5
economic	1	45.3	42.6	50.7
economic	10	46.2	43.3	49.4
economic	100	47.7	44.6	52.9
economic	1000	41.3	45.7	54.1
economic	10000	30.2	28.3	53.2
physics	1	25.3	22.7	30.2
physics	10	26.1	23.1	31.3
physics	100	27.4	24.6	32.6
physics	1000	28.7	25.4	27.7
physics	10000	15.2	13.4	30.3
history	1	10.3	9.7	12.4
history	10	11.2	10.3	13.3
history	100	11.1	11.4	14.1
history	1000	12.6	12.8	15.3
history	10000	6.1	5.7	14.8

1212 Table 9: Editing experiments using MEND across domains and hierarchy levels.

## F SEQUENTIAL UPDATE

1217 To further validate our claim that **editing and unlearning behave fundamentally differently**, we  
 1218 additionally conducted *multi-step sequential updates* on multiple facts using Qwen3-14B, LLaMA3-  
 1219 8B, and Gemma-7B. This section reports the results to illustrate the phenomenon clearly.

### SEQUENTIAL EDITING BEHAVIOR

1223 Across multiple sequential edits, the model retains previously edited knowledge with only minor  
 1224 drift. Even after five cumulative edits, the performance on earlier edited facts remains largely stable.  
 1225 This supports our claim that **editing operations are robust and localized**, even under sequential  
 1226 updates.

Acc	1	1 & 2	1 & 2 & 3	1 & 2 & 3 & 4	1 & 2 & 3 & 4 & 5
Edit Fact 1	55.0%	54.5%	54.0%	53.8%	53.5%
Edit Fact 2	—	48.0%	47.5%	47.0%	46.5%
Edit Fact 3	—	—	62.0%	61.5%	61.0%
Edit Fact 4	—	—	—	50.0%	49.5%
Edit Fact 5	—	—	—	—	57.0%

1234 Table 10: Sequential editing performance.

### SEQUENTIAL UNLEARNING BEHAVIOR

1239 In contrast, unlearning shows **clear cumulative degradation**. When more facts are removed se-  
 1240 quentially, the model’s performance on earlier unlearned facts, as well as related queries, drops  
 1241 sharply. This supports our central claim: **Unlearning is inherently more disruptive than editing**,  
 because removing information often affects interconnected knowledge.

Acc	1	1 & 2	1 & 2 & 3	1 & 2 & 3 & 4	1 & 2 & 3 & 4 & 5
Unlearn Fact 1	54.2%	37.7%	40.5%	34.3%	27.2%
Unlearn Fact 2	—	45.1%	37.8%	31.5%	25.2%
Unlearn Fact 3	—	—	43.2%	36.0%	28.8%
Unlearn Fact 4	—	—	—	40.5%	31.5%
Unlearn Fact 5	—	—	—	—	34.2%

Table 11: Sequential unlearning performance.

## G ACCURACY RESULT

Editing accuracy for the 13 model across llama3, qwen3, qwq, mistral, gemma and deepseek families are lists below in Table 12. Unlearning accuracy for the 13 model across llama3, qwen3, qwq, mistral, gemma and deepseek families are lists below in Table 13.

## H MODEL SIMILARITY RESULT

**Representation Similarity Analysis** Our unified framework models editing and unlearning as optimizing  $\mathcal{L}_{\text{task}}$  against  $\mathcal{L}_{\text{pres}}$ . While probe-based evaluation measures outcomes on  $\mathcal{Q}^+$  and  $\mathcal{Q}^-$ , it does not reveal how the internal representations change during this optimization. To capture these hidden dynamics, we analyze representational shifts from the original (pre-KnowledgeSmith) state to the post-KnowledgeSmith state using Centered Kernel Alignment (CKA) (Kornblith et al., 2019), KL divergence, L2 distance and Fisher score (Zhang et al., 2022).

For unlearning, these metrics expose a sharp phase transition around 1000 samples: below this point, representations remain close to baseline, but beyond it they reorganize abruptly, suggesting a capacity breakpoint where  $\mathcal{L}_{\text{pres}}$  is overwhelmed by repeated optimization on  $\mathcal{Q}^+$ . Editing, in contrast, produces smoother trajectories. KL divergence and Fisher scores increase steadily with training size, indicating progressive local updates to representations rather than wholesale restructuring. For example, biology edits on DeepSeek-8B show KL and Fisher growing from (KL≈20, Fisher≈9.7) with a single sample to (KL≈172, Fisher≈93.7) at 1000 samples, after which growth plateaus as the optimization stabilizes.

These results demonstrate that **unlearning triggers abrupt phase transitions in representation space once data scale crosses a threshold**, while editing produces gradual, localized adjustments, underscoring the need for representation level analysis beyond probe accuracy.

**Computationally Efficiency.** For the same model on a target dataset of 10,000 examples, unlearning typically completes in about 1.5 hours on an NVIDIA H100. Knowledge editing is more resource-intensive (roughly 6 hours). This additional cost highlights the heavier computational demands of precise factual editing.

In summary, **unlearning prioritizes stability and low computational cost, while editing maximizes factual enforcement but risks destabilizing other knowledge and requires more resources**. The choice between the two depends on whether minimizing collateral effects or maximizing certainty of change is the primary goal.

Model similarity for llama3, qwen3, qwq, mistral, gemma and deepseek 6 families are lists below in Tables 14 to 19

## I LLM USAGE

We use large language models (LLMs) only for grammar checking and correction.

Table 12: Editing Accuracy

	Branch	Train Size	Test Set	llama3.2-1b-instruct				llama3.8b-instruct			
				Biology	History	Economic	Physics	Biology	History	Economic	Physics
1296	Intermediate	1	ID	45.83	25	37.5	20.83	20.83	4.17	29.17	12.5
1297			OOD	44.05	44.07	44.11	44.05	63.08	63.13	63.22	63.05
1298		10	ID	20.83	25	50	33.33	25	12.5	25	20.83
1299			OOD	44.04	44.06	44.02	42.06	63.07	63.04	63.1	63
1300		100	ID	45.83	25	45.83	0	20.83	12.5	20.83	41.67
1301			OOD	23.35	43.99	44.08	26.88	63.05	63.13	63.12	24.72
1302		1000	ID	4.17	25	45.83	0	20.83	12.5	20.83	12.5
1303			OOD	25.22	43.98	44.17	26.85	63.06	63.03	63.09	24.3
1304		10000	ID	4.17	25	45.83	0	20.83	12.5	20.83	12.5
1305			OOD	25.22	43.98	44.17	26.85	63.06	63.03	63.09	24.3
1306	Root	1	ID	41.67	25	41.67	29.17	16.67	4.17	33.33	16.67
1307			OOD	44.17	44.16	43.99	44.15	63.04	62.91	63.1	63.05
1308		10	ID	29.17	25	45.83	29.17	12.5	4.17	33.33	25
1309			OOD	44.1	44.28	44.12	44.07	63.01	62.9	63.12	63.1
1310		100	ID	29.17	25	4.17	20.83	12.5	4.17	33.33	25
1311			OOD	44.12	44.22	26.24	44.2	63	62.97	63.15	63
1312		1000	ID	29.17	25	0	25	12.5	4.17	33.33	16.67
1313			OOD	44.15	44.26	25.4	44.09	62.98	62.98	63.11	63.11
1314		10000	ID	29.17	25	0	25	12.5	4.17	33.33	16.67
1315			OOD	44.15	44.26	25.4	44.09	62.98	62.98	63.11	63.11
1316	Leaf	1	ID	41.67	25	33.33	25	16.67	4.17	37.5	16.67
1317			OOD	44.13	44.11	44.08	44.02	63.1	63.07	63.09	63.1
1318		10	ID	25	25	62.5	20.83	16.67	4.17	37.5	25
1319			OOD	44.19	44.41	43.74	43.18	63.12	62.96	63.06	62.75
1320		100	ID	4.17	4.17	4.17	0	16.67	12.5	25	4.17
1321			OOD	25.45	40.55	25.53	26.88	62.78	62.47	62.6	25.41
1322		1000	ID	25	45.83	0	8.33	16.67	4.17	25	16.67
1323			OOD	25.78	23.44	26.63	24.84	62.77	59.56	62.59	24.25
1324	Intermediate	llama3.2-3b-instruct				llama3.3-70b-instruct					
1325		1	ID	25	12.5	41.67	16.67	20.83	8.33	20.83	25
1326			OOD	59.17	59.24	59.36	59.22	81.44	81.42	81.39	81.42
1327		10	ID	12.5	0	37.5	37.5	20.83	62.5	41.67	29.17
1328			OOD	56.38	56.84	58.63	58.39	81.38	81.38	81.43	81.48
1329		100	ID	29.17	20.83	54.17	12.5	20.83	58.33	50	29.17
1330			OOD	23.47	26.9	23.32	25.84	81.46	81.26	81.33	81.39
1331		1000	ID	4.17	45.83	41.67	0	25	58.33	50	29.17
1332			OOD	25.45	25.12	25.31	25.08	81.39	81.35	81.38	81.31
1333		10000	ID	4.17	45.83	41.67	0	25	58.33	50	29.17
1334			OOD	25.45	25.12	25.31	25.08	81.39	81.35	81.38	81.31
1335	Root	1	ID	25	4.17	29.17	12.5	20.83	4.17	20.83	25
1336			OOD	59.2	59.28	59.24	59.34	81.41	81.46	81.46	81.43
1337		10	ID	41.67	29.17	16.67	4.17	58.33	45.83	37.5	33.33
1338			OOD	58.76	58.73	58.37	58.72	81.41	81.46	81.39	81.51
1339		100	ID	37.5	0	41.67	29.17	58.33	25	41.67	33.33
1340			OOD	23.3	26.86	24.4	25.34	81.39	81.41	81.4	81.42
1341		1000	ID	4.17	0	20.83	25	58.33	20.83	41.67	33.33
1342			OOD	25.57	26.48	25.06	25.2	81.42	81.46	81.48	81.44
1343	Leaf	1	ID	20.83	8.33	33.33	20.83	20.83	8.33	20.83	25
1344			OOD	59.23	59.25	59.3	59.24	81.43	81.37	81.43	81.45
1345		10	ID	37.5	8.33	45.83	29.17	25	25	58.33	20.83
1346			OOD	59.13	59.26	58.35	56.96	81.43	81.38	81.41	81.34
1347		100	ID	8.33	4.17	54.17	20.83	25	20.83	58.33	20.83
1348			OOD	24.49	25.52	23.21	24.55	81.41	81.44	81.33	81.5
1349		1000	ID	20.83	4.17	54.17	4.17	25	25	62.5	20.83
			OOD	27.23	26.24	23.19	25.42	81.43	81.37	81.29	81.44

1350	Branch	Train Size	Test Set	qwen3-1.7b				qwen3-32b			
				Biology	History	Economic	Physics	Biology	History	Economic	Physics
				1	ID	20.83	12.5	33.33	29.17	12.5	0
1351	Intermediate	1	OOD	53	53.99	54.05	54.08	75.07	75.11	75.19	75.09
			10	ID	25	12.5	33.33	29.17	20.83	0	20.83
		100	OOD	53.55	53.95	54.05	54.08	75.07	75.02	75.08	75.02
			ID	20.83	12.5	33.33	37.5	20.83	0	29.17	8.33
		1000	OOD	53	54	53.82	53.02	75.17	75.15	74.97	75.1
			ID	25	12.5	33.33	37.5	20.83	0	29.17	8.33
		10000	OOD	53.04	53.99	53.87	53.07	75.2	75.15	74.86	75.03
			ID	25	12.5	33.33	37.5	20.83	0	29.17	8.33
		100000	OOD	53.04	53.99	53.87	53.07	75.2	75.15	74.86	75.03
			ID	25	12.5	33.33	37.5	20.83	0	29.17	8.33
1360	Root	1	ID	37.5	45.83	25	12.5	20.83	16.67	45.83	8.33
			OOD	53.81	53.8	53.7	54	75.2	75.12	75.21	75
		10	ID	33.33	45.83	16.67	12.5	16.67	25	33.33	16.67
			OOD	53.78	53.78	53.87	54	75.05	75.17	75.02	75.07
		100	ID	29.17	45.83	25	20.83	20.83	12.5	12.5	16.67
			OOD	53.74	53.8	53.7	53.65	75.02	75.15	74.98	75.04
		1000	ID	29.17	45.83	25	20.83	20.83	16.67	12.5	16.67
			OOD	53.79	53.82	53.75	53.75	75.1	75.1	75	75.06
		10000	ID	29.17	45.83	25	20.83	20.83	16.67	12.5	16.67
			OOD	53.79	53.82	53.75	53.75	75.1	75.1	75	75.06
1369	Leaf	1	ID	16.67	37.5	37.5	25	20.83	0	25	8.33
			OOD	53.97	53.19	53.28	51.66	75.15	75.1	75.07	75.07
		10	ID	16.67	29.17	8.33	16.67	20.83	25	29.17	4.17
			OOD	53.87	53.86	53.76	53.92	74.9	75.05	75.15	75.17
		100	ID	16.67	33.33	29.17	25	29.17	16.67	54.17	4.17
			OOD	54.2	53.2	53.33	51.66	75.1	75.12	74.88	75.13
		1000	ID	16.67	37.5	37.5	33.33	25	4.17	41.67	4.17
			OOD	53.69	53.19	53.28	39.71	75.16	74.18	74.69	75.15
		10000	ID	16.67	37.5	37.5	33.33	25	4.17	41.67	4.17
			OOD	53.69	53.19	53.28	39.71	75.16	74.18	74.69	75.15
1378	Leaf	qwen3-14b				qwq-32b					
		1	ID	20.83	8.33	16.67	25	16.67	4.17	70.83	12.5
			OOD	73.84	73.86	73.89	73.94	77.4	77.45	77.42	77.45
		10	ID	20.83	0	4.17	20.83	12.5	4.17	33.33	12.5
			OOD	73.78	73.54	73.42	73.62	77.36	77.39	77.42	77.35
		100	ID	25	4.17	4.17	16.67	16.67	0	37.5	12.5
			OOD	73.76	73.45	73.36	73.56	77.28	77.43	77.41	77.39
		1000	ID	25	4.17	4.17	16.67	16.67	0	37.5	12.5
			OOD	73.73	73.42	73.33	73.56	77.33	77.43	77.41	77.4
		10000	ID	25	4.17	4.17	16.67	16.67	0	37.5	12.5
			OOD	73.73	73.42	73.33	73.56	77.33	77.43	77.41	77.4
1388	Root	1	ID	41.67	20.83	33.33	25	12.5	20.83	20.83	12.5
			OOD	73.86	73.9	73.86	73.79	77.4	77.47	77.37	77.42
		10	ID	20.83	8.33	33.33	16.67	20.83	12.5	16.67	12.5
			OOD	73.81	73.71	73.87	73.85	77.35	77.53	77.47	77.38
		100	ID	16.67	12.5	37.5	16.67	16.67	12.5	16.67	12.5
			OOD	73.71	73.68	73.81	73.58	77.43	77.39	77.43	77.3
		1000	ID	16.67	12.5	41.67	16.67	16.67	12.5	16.67	12.5
			OOD	73.73	73.64	73.84	73.55	77.44	77.39	77.45	77.35
		10000	ID	16.67	12.5	41.67	16.67	16.67	12.5	16.67	12.5
			OOD	73.73	73.64	73.84	73.55	77.44	77.39	77.45	77.35
1397	Leaf	1	ID	25	0	16.67	20.83	16.67	0	33.33	12.5
			OOD	73.89	73.89	73.87	73.88	77.48	77.39	77.4	77.48
		10	ID	20.83	4.17	33.33	12.5	25	0	41.67	12.5
			OOD	73.91	73.63	73.63	73.69	77.5	77.37	77.33	77.42
		100	ID	20.83	0	29.17	16.67	29.17	0	29.17	12.5
			OOD	73.66	73.39	73.42	73.6	77.4	77.52	77.33	77.4
		1000	ID	20.83	0	8.33	16.67	20.83	0	41.67	12.5
			OOD	65.94	73.39	39.3	73.5	77.27	77.52	68.74	77.28
		10000	ID	20.83	0	8.33	16.67	20.83	0	41.67	12.5
			OOD	65.94	73.39	39.3	73.5	77.27	77.52	68.74	77.28

1404	Branch	Train Size	Test Set	mistral-Small-24B-Instruct-2501				gemma-2b			
				Biology	History	Economic	Physics	Biology	History	Economic	Physics
				16.67	50	12.5	20.83	8.33	4.17	12.5	12.5
1405	Intermediate	1	ID	73.4	73.34	73.24	73.39	30.46	30.63	30.49	30.53
1406			OOD	50	45.83	16.67	0	20.83	12.5	29.17	4.17
1407		10	ID	24.47	22.99	25.2	25.51	29.4	30.37	29.06	30.29
1408			OOD	29.17	45.83	54.17	58.33	16.67	45.83	37.5	29.17
1409		100	ID	24.22	22.95	24.16	23	25.81	24.08	26.16	26.54
1410			OOD	45.83	45.83	54.17	58.33	16.67	45.83	54.17	58.33
1411		1000	ID	22.87	22.95	24.16	23	25.81	22.95	22.95	22.95
1412			OOD	45.83	45.83	54.17	58.33	16.67	45.83	54.17	58.33
1413		10000	ID	22.87	22.95	24.16	23	25.81	22.95	22.95	22.95
1414			OOD	33.33	20.83	16.67	45.83	4.17	0.0	16.67	4.17
1415	Root	1	ID	73.42	73.42	73.16	73.39	30.54	30.59	30.64	30.54
1416			OOD	45.83	54.17	4.17	58.33	8.33	8.33	12.5	33.33
1417		10	ID	22.95	25.2	25.27	23.14	30.34	27.18	30.55	25.79
1418			OOD	45.83	37.5	54.17	58.33	4.17	50.0	25.0	54.17
1419		100	ID	22.83	24.4	23.32	22.99	29.3	24.51	29.13	23.74
1420			OOD	45.83	37.5	41.67	37.5	45.83	45.83	54.17	58.33
1421		1000	ID	23.14	24.4	24.76	25.12	22.95	22.95	22.95	22.95
1422			OOD	45.83	37.5	41.67	37.5	45.83	45.83	54.17	58.33
1423	Leaf	1	ID	50	41.67	83.33	4.17	8.33	33.33	25.0	0.0
1424			OOD	73.32	73.24	73.14	73.42	30.25	28.98	30.64	30.44
1425		10	ID	4.17	45.83	4.17	41.67	12.5	33.33	20.83	4.17
1426			OOD	25.47	22.95	25.54	25.19	27.7	24.6	25.28	29.08
1427		100	ID	8.33	45.83	54.17	8.33	25.0	37.5	37.5	41.67
1428			OOD	26.63	22.96	22.95	24.61	24.94	24.19	24.9	24.07
1429		1000	ID	41.67	45.83	54.17	8.33	45.83	45.83	54.17	58.33
1430			OOD	23.54	22.96	22.95	24.61	22.83	22.95	22.95	22.95
1431		10000	ID	41.67	45.83	54.17	8.33	45.83	45.83	54.17	58.33
1432			OOD	23.54	22.96	22.95	24.61	22.83	22.95	22.95	22.95
1433	Intermediate	mistral-Large-Instruct-2411				gemma-7b					
1434		1	ID	25.0	62.5	25.0	12.5	45.83	37.5	41.67	45.83
1435			OOD	82.13	82.42	82.22	82.37	59.22	58.96	56.69	57.78
1436		10	ID	0.0	45.83	41.67	37.5	45.83	45.83	54.17	50.0
1437			OOD	26.89	22.97	24.53	24.7	22.95	22.95	22.95	23.25
1438		100	ID	16.67	62.5	25.0	50.0	25.0	45.83	50.0	41.67
1439			OOD	23.89	25.84	25.0	23.05	24.2	22.95	23.11	23.11
1440		1000	ID	16.67	62.5	25.0	50.0	29.17	54.17	8.33	66.67
1441			OOD	23.89	25.84	25.0	23.05	24.9	25.59	25.22	24.68
1442	Root	1	ID	25.0	12.5	45.83	62.5	37.5	41.67	50.0	16.67
1443			OOD	82.25	82.22	82.24	82.25	59.7	59.56	57.63	59.74
1444		10	ID	0.0	0.0	4.17	58.33	45.83	41.67	33.33	58.33
1445			OOD	26.19	26.89	25.41	22.95	28.43	22.97	29.24	22.97
1446		100	ID	8.33	45.83	37.5	58.33	45.83	45.83	45.83	50.0
1447			OOD	26.86	22.95	24.64	23.0	22.95	23.07	22.95	24.13
1448		1000	ID	8.33	45.83	37.5	58.33	33.33	45.83	20.83	54.17
1449			OOD	26.86	22.95	24.64	23.0	23.98	23.07	23.24	23.34
1450		10000	ID	8.33	45.83	37.5	58.33	33.33	45.83	20.83	54.17
1451			OOD	26.86	22.95	24.64	23.0	23.98	23.07	23.24	23.34
1452	Leaf	1	ID	54.17	29.17	41.67	37.5	45.83	45.83	54.17	33.33
1453			OOD	82.19	82.25	82.07	82.08	22.82	22.97	22.95	59.29
1454		10	ID	4.17	0.0	54.17	58.33	45.83	45.83	45.83	58.33
1455			OOD	25.47	25.51	22.95	23.07	22.95	23.07	23.33	22.87
1456		100	ID	50.0	0.0	45.83	54.17	37.5	41.67	41.67	58.33
1457			OOD	23.05	24.6	24.69	25.55	23.96	23.78	23.38	22.94
1458		1000	ID	50.0	0.0	45.83	54.17	4.17	54.17	4.17	4.17
1459			OOD	23.05	24.6	24.69	25.55	25.48	24.49	25.52	25.54

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

			DeepSeek-R1-0528-Qwen3-8B			
Branch	Train Size	Test Set	Biology	History	Economic	Physics
Intermediate	1	ID	8.33	0.0	20.83	25.0
		OOD	65.99	66.09	66.0	66.02
	10	ID	25.0	0.0	33.33	29.17
		OOD	65.94	65.93	66.01	65.95
	100	ID	16.67	0.0	33.33	33.33
		OOD	65.95	65.94	66.07	65.89
	1000	ID	16.67	0.0	33.33	37.5
		OOD	65.9	65.94	66.07	66.02
	10000	ID	16.67	0.0	33.33	37.5
		OOD	65.9	65.94	66.07	66.02
Root	1	ID	12.5	8.33	0.0	45.83
		OOD	66.07	65.96	28.17	65.94
	10	ID	12.5	8.33	16.67	45.83
		OOD	66.02	65.99	65.98	65.93
	100	ID	12.5	4.17	8.33	45.83
		OOD	65.97	66.1	66.02	66.0
	1000	ID	12.5	4.17	8.33	45.83
		OOD	65.92	66.02	65.99	65.92
	10000	ID	12.5	4.17	8.33	45.83
		OOD	65.92	66.02	65.99	65.92
Leaf	1	ID	16.67	0.0	20.83	25.0
		OOD	65.9	65.96	66.09	65.92
	10	ID	20.83	8.33	37.5	25.0
		OOD	65.83	65.84	65.92	65.92
	100	ID	16.67	4.17	29.17	25.0
		OOD	65.95	65.77	65.92	65.76
	1000	ID	16.67	4.17	25.0	20.83
		OOD	66.02	65.77	64.29	65.8
	10000	ID	16.67	4.17	25.0	20.83
		OOD	66.02	65.77	64.29	65.8

Table 13: Unlearning Accuracy

1512	Branch	Train Size	Test Set	llama3.2-1b-instruct				llama3.8b-instruct			
				Biology	History	Economic	Physics	Biology	History	Economic	Physics
				ID	OOD	ID	OOD	ID	OOD	ID	OOD
1513	Intermediate	1	ID	25.0	20.83	33.33	12.5	16.67	4.17	29.17	12.5
1514			OOD	32.69	32.69	32.69	32.69	63.01	63.01	63.01	63.01
1515		10	ID	20.83	20.83	33.33	16.67	16.67	4.17	29.17	12.5
1516			OOD	32.75	32.74	32.74	32.6	63.01	63.0	63.0	62.98
1517		100	ID	29.17	20.83	33.33	16.67	16.67	4.17	37.5	12.5
1518			OOD	32.55	32.93	32.76	32.77	62.85	62.8	62.84	62.75
1519		1000	ID	33.33	4.17	37.5	16.67	16.67	4.17	33.33	12.5
1520			OOD	32.67	33.91	32.72	32.51	62.93	62.82	62.91	62.85
1521		10000	ID	33.33	4.17	37.5	16.67	16.67	4.17	33.33	12.5
1522			OOD	32.67	33.91	32.72	32.51	62.93	62.82	62.91	62.85
1523	Root	1	ID	25.0	20.83	33.33	12.5	16.67	4.17	29.17	12.5
1524			OOD	32.69	32.69	32.69	32.69	63.01	63.01	63.01	63.01
1525		10	ID	20.83	20.83	33.33	16.67	16.67	4.17	29.17	12.5
1526			OOD	32.63	32.61	32.55	32.69	63.0	63.0	63.02	63.0
1527		100	ID	29.17	16.67	33.33	37.5	16.67	4.17	33.33	12.5
1528			OOD	32.97	32.74	32.84	32.94	62.91	62.98	62.87	62.89
1529		1000	ID	16.67	8.33	37.5	4.17	16.67	4.17	33.33	12.5
1530			OOD	32.66	32.86	33.11	33.26	63.01	62.99	62.81	62.89
1531		10000	ID	16.67	8.33	37.5	4.17	16.67	4.17	33.33	12.5
1532			OOD	32.66	32.86	33.11	33.26	63.01	62.99	62.81	62.89
1533	Leaf	1	ID	25.0	20.83	33.33	12.5	16.67	4.17	29.17	12.5
1534			OOD	32.69	32.69	32.69	32.69	63.01	63.01	63.01	63.01
1535		10	ID	25.0	20.83	37.5	12.5	16.67	4.17	29.17	12.5
1536			OOD	32.69	32.58	32.69	32.74	62.98	63.02	62.99	62.98
1537		100	ID	25.0	20.83	25.0	16.67	16.67	4.17	33.33	12.5
1538			OOD	32.57	32.85	32.73	32.73	62.75	62.75	63.02	62.75
1539		1000	ID	20.83	12.5	20.83	20.83	16.67	4.17	33.33	12.5
1540			OOD	32.68	33.08	32.27	31.68	62.68	62.53	62.98	62.69
1541		10000	ID	20.83	12.5	20.83	20.83	16.67	4.17	33.33	12.5
1542			OOD	32.68	33.08	32.27	31.68	62.68	62.53	62.98	62.69
1543	Intermediate	1	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1544			OOD	59.33	59.33	59.33	59.37	81.33	81.33	81.33	81.33
1545		10	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1546			OOD	59.17	59.29	59.27	59.24	81.33	81.33	81.33	81.33
1547		100	ID	25.0	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1548			OOD	59.07	59.05	59.51	59.06	81.35	81.36	81.35	81.35
1549		1000	ID	16.67	4.17	41.67	12.5	20.83	8.33	29.17	29.72
1550			OOD	59.14	59.22	59.13	59.3	81.38	81.47	81.41	81.37
1551		10000	ID	16.67	4.17	41.67	12.5	20.83	8.33	29.17	29.72
1552			OOD	59.14	59.22	59.13	59.3	81.38	81.47	81.41	81.37
1553	Root	1	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1554			OOD	59.33	59.33	59.33	59.37	81.33	81.33	81.33	81.33
1555		10	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1556			OOD	59.29	59.3	59.34	59.2	81.33	81.33	81.33	81.33
1557		100	ID	25.0	4.17	37.5	16.67	20.83	8.33	20.83	20.83
1558			OOD	58.94	58.98	59.51	59.07	81.38	81.35	81.37	81.35
1559		1000	ID	16.67	4.17	41.67	12.5	20.83	8.33	25.0	23.33
1560			OOD	58.96	58.99	59.41	59.12	81.39	81.33	81.41	81.33
1561		10000	ID	16.67	4.17	41.67	12.5	20.83	8.33	25.0	23.33
1562			OOD	58.96	58.99	59.41	59.12	81.39	81.33	81.41	81.33
1563	Leaf	1	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1564			OOD	59.33	59.33	59.33	59.37	81.33	81.33	81.33	81.33
1565		10	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1566			OOD	59.29	59.27	59.16	59.26	81.33	81.33	81.33	81.33
1567		100	ID	20.83	4.17	41.67	16.67	20.83	8.33	20.83	20.83
1568			OOD	59.12	59.09	59.46	59.08	81.37	81.37	81.35	81.39
1569		1000	ID	12.5	4.17	45.83	16.67	20.83	8.33	25.0	20.33
1570			OOD	58.99	58.94	59.31	58.87	81.32	81.37	81.44	81.32
1571		10000	ID	12.5	4.17	45.83	16.67	20.83	8.33	25.0	20.33
1572			OOD	58.99	58.94	59.31	58.87	81.32	81.37	81.44	81.32

1566	Branch	Train Size	Test Set	qwen3-1.7b				qwen3-32b			
				Biology	History	Economic	Physics	Biology	History	Economic	Physics
				ID	OOD	ID	OOD	ID	OOD	ID	OOD
1567	Intermediate	1	ID	16.67	16.67	33.33	16.67	16.67	0.0	16.67	12.5
			OOD	53.9	53.92	53.92	53.93	75.13	75.13	75.13	75.13
		10	ID	16.67	16.67	33.33	16.67	16.67	0.0	16.67	12.5
			OOD	53.92	53.97	54.09	53.95	75.13	75.13	75.13	75.13
		100	ID	8.33	16.67	20.83	8.33	16.67	0.0	16.67	12.5
			OOD	53.25	53.51	54.42	53.25	75.07	75.07	75.14	75.07
		1000	ID	25.0	20.83	25.0	25.0	16.67	0.0	25.0	12.5
			OOD	52.66	52.36	53.6	52.64	75.21	75.07	75.26	74.98
		10000	ID	25.0	20.83	25.0	25.0	16.67	0.0	25.0	12.5
			OOD	52.66	52.36	53.6	53.64	75.21	75.07	75.26	74.98
1571	Root	1	ID	12.5	16.67	33.33	15.33	16.67	0.0	16.67	12.5
			OOD	53.92	53.92	53.92	53.92	75.13	75.13	75.13	75.13
		10	ID	16.67	16.67	33.33	16.67	16.67	0.0	16.67	12.5
			OOD	53.99	53.85	53.83	53.89	75.13	75.13	75.13	75.13
		100	ID	8.33	16.67	20.83	16.67	16.67	0.0	16.67	12.5
			OOD	53.55	53.0	54.01	53.35	75.18	75.16	75.12	75.15
		1000	ID	33.33	29.17	37.5	29.33	16.67	0.0	25.0	12.5
			OOD	52.67	51.83	53.65	53.67	75.33	75.16	75.05	75.23
		10000	ID	33.33	29.17	37.5	33.33	16.67	0.0	25.0	12.5
			OOD	52.67	51.83	53.65	52.67	75.33	75.16	75.05	75.23
1585	Leaf	1	ID	16.67	16.67	33.33	16.67	16.67	0.0	16.67	12.5
			OOD	53.9	53.9	53.92	53.9	75.13	75.13	75.13	75.09
		10	ID	16.67	16.67	33.33	16.67	16.67	0.0	16.67	12.5
			OOD	54.0	54.02	54.0	54.05	75.13	75.13	75.13	75.13
		100	ID	16.67	16.67	25.0	16.67	16.67	0.0	20.83	12.5
			OOD	53.82	53.68	54.56	53.88	75.07	75.1	75.05	75.14
		1000	ID	25.0	29.17	25.0	25.0	16.67	0.0	25.0	12.5
			OOD	52.81	53.16	53.46	53.81	75.11	75.1	75.25	74.71
		10000	ID	25.0	29.17	25.0	25.0	16.67	0.0	25.0	12.5
			OOD	52.81	53.16	53.46	53.73	75.11	75.1	75.25	74.71
1595	Intermediate	qwen3-14b				qwq-32b					
		1	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.86	73.86	73.86	73.86	77.42	77.42	77.42	77.42
		10	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.84	73.86	73.86	73.86	77.38	77.45	77.44	77.4
		100	ID	20.83	0.0	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.61	73.63	73.89	73.94	77.4	77.37	77.28	77.37
		1000	ID	20.83	4.17	20.83	16.67	12.5	0.0	29.17	12.5
			OOD	73.15	73.23	73.83	73.59	77.35	77.37	77.27	77.42
		10000	ID	20.83	4.17	20.83	16.67	12.5	0.0	29.17	12.5
			OOD	73.15	73.23	73.83	73.59	77.35	77.37	77.27	77.42
1604	Root	1	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.86	73.86	73.86	73.86	77.42	77.42	77.42	77.42
		10	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.88	73.86	73.86	73.85	77.45	77.47	77.44	77.48
		100	ID	20.83	0.0	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.84	73.66	73.86	73.62	77.35	77.3	77.38	77.45
		1000	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.54	73.48	73.5	73.26	77.55	77.3	77.38	77.55
		10000	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.54	73.48	73.5	73.26	77.55	77.3	77.38	77.55
1613	Leaf	1	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.86	73.86	73.86	73.86	77.42	77.42	77.42	77.42
		10	ID	20.83	4.17	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.86	73.86	73.86	73.86	77.42	77.47	77.43	77.45
		100	ID	20.83	0.0	25.0	12.5	12.5	0.0	29.17	12.5
			OOD	73.64	73.91	73.84	73.83	77.39	77.4	77.38	77.26
		1000	ID	20.83	0.0	25.0	16.67	12.5	0.0	29.17	12.5
			OOD	72.99	73.91	73.64	73.51	77.3	77.4	77.47	77.35
		10000	ID	20.83	0.0	25.0	16.67	12.5	0.0	29.17	12.5
			OOD	72.99	73.91	73.64	73.51	77.3	77.4	77.47	77.35

1620	Branch	Train Size	Test Set	mistral-Small-24B-Instruct-2501				gemma-2b				
				Biology	History	Economic	Physics	Biology	History	Economic	Physics	
				16.67	50	12.5	20.83	8.33	4.17	12.5	12.5	
1621	Intermediate	1	ID	73.4	73.34	73.24	73.39	30.46	30.63	30.49	30.53	
1622			OOD	50	45.83	16.67	0	20.83	12.5	29.17	4.17	
1623		10	ID	24.47	22.99	25.2	25.51	29.4	30.37	29.06	30.29	
1624			OOD	29.17	45.83	54.17	58.33	16.67	45.83	37.5	29.17	
1625		100	ID	24.22	22.95	24.16	23	25.81	24.08	26.16	26.54	
1626			OOD	45.83	45.83	54.17	58.33	45.83	45.83	54.17	58.33	
1627		1000	ID	22.87	22.95	24.16	23	22.95	22.95	22.95	22.95	
1628			OOD	45.83	45.83	54.17	58.33	45.83	45.83	54.17	58.33	
1629		10000	ID	22.87	22.95	24.16	23	22.95	22.95	22.95	22.95	
1630			OOD	33.33	20.83	16.67	45.83	4.17	0.0	16.67	4.17	
1631	Root	1	ID	73.42	73.42	73.16	73.39	30.54	30.59	30.64	30.54	
1632			OOD	45.83	54.17	4.17	58.33	8.33	8.33	12.5	33.33	
1633		10	ID	22.95	25.2	25.27	23.14	30.34	27.18	30.55	25.79	
1634			OOD	45.83	37.5	54.17	58.33	4.17	50.0	25.0	54.17	
1635		100	ID	22.83	24.4	23.32	22.99	29.3	24.51	29.13	23.74	
1636			OOD	45.83	37.5	41.67	37.5	45.83	45.83	54.17	58.33	
1637		1000	ID	23.14	24.4	24.76	25.12	22.95	22.95	22.95	22.95	
1638			OOD	45.83	37.5	41.67	37.5	45.83	45.83	54.17	58.33	
1639		10000	ID	23.14	24.4	24.76	25.12	22.95	22.95	22.95	22.95	
1640			OOD	50	41.67	83.33	4.17	8.33	33.33	25.0	0.0	
1641	Leaf	1	ID	73.32	73.24	73.14	73.42	30.25	28.98	30.64	30.44	
1642			OOD	4.17	45.83	4.17	41.67	12.5	33.33	20.83	4.17	
1643		10	ID	25.47	22.95	25.54	25.19	27.7	24.6	25.28	29.08	
1644			OOD	8.33	45.83	54.17	8.33	25.0	37.5	37.5	41.67	
1645		100	ID	26.63	22.96	22.95	24.61	24.94	24.19	24.9	24.07	
1646			OOD	41.67	45.83	54.17	8.33	45.83	45.83	54.17	58.33	
1647		1000	ID	23.54	22.96	22.95	24.61	22.83	22.95	22.95	22.95	
1648			OOD	41.67	45.83	54.17	8.33	45.83	45.83	54.17	58.33	
1649		Intermediate	mistral-Large-Instruct-2411				gemma-7b					
1650			1	ID	82.13	82.42	82.22	82.37	59.22	58.96	56.69	57.78
1651				OOD	0.0	45.83	41.67	37.5	45.83	45.83	54.17	50.0
1652			10	ID	26.89	22.97	24.53	24.7	22.95	22.95	22.95	23.25
1653				OOD	16.67	62.5	25.0	50.0	25.0	45.83	50.0	41.67
1654			100	ID	23.89	25.84	25.0	23.05	24.2	22.95	23.11	23.11
1655				OOD	16.67	62.5	25.0	50.0	29.17	54.17	8.33	66.67
1656			1000	ID	23.89	25.84	25.0	23.05	24.9	25.59	25.22	24.68
1657				OOD	16.67	62.5	25.0	50.0	29.17	54.17	8.33	66.67
1658			10000	ID	23.89	25.84	25.0	23.05	24.9	25.59	25.22	24.68
1659				OOD	25.0	12.5	45.83	62.5	37.5	41.67	50.0	16.67
1660	Root	1	ID	82.25	82.22	82.24	82.25	59.7	59.56	57.63	59.74	
1661			OOD	0.0	0.0	4.17	58.33	45.83	41.67	33.33	58.33	
1662		10	ID	26.19	26.89	25.41	22.95	28.43	22.97	29.24	22.97	
1663			OOD	8.33	45.83	37.5	58.33	45.83	45.83	45.83	50.0	
1664		100	ID	26.86	22.95	24.64	23.0	22.95	23.07	22.95	24.13	
1665			OOD	8.33	45.83	37.5	58.33	33.33	45.83	20.83	54.17	
1666		1000	ID	26.86	22.95	24.64	23.0	23.07	23.07	23.24	23.34	
1667			OOD	8.33	45.83	37.5	58.33	33.33	45.83	20.83	54.17	
1668		10000	ID	26.86	22.95	24.64	23.0	23.07	23.07	23.24	23.34	
1669			OOD	25.47	25.51	22.95	23.07	22.95	23.07	23.33	22.87	
1670	Leaf	1	ID	23.05	24.6	24.69	25.55	23.96	23.78	23.38	22.94	
1671			OOD	50.0	0.0	45.83	54.17	37.5	41.67	41.67	58.33	
1672		10	ID	23.05	24.6	24.69	25.55	24.49	24.49	25.52	25.54	
1673			OOD	50.0	0.0	45.83	54.17	4.17	54.17	4.17	4.17	
1674		100	ID	23.05	24.6	24.69	25.55	25.48	24.49	24.49	25.52	
1675			OOD	50.0	0.0	45.83	54.17	4.17	54.17	4.17	4.17	
1676		1000	ID	23.05	24.6	24.69	25.55	25.48	24.49	24.49	25.52	
1677			OOD	50.0	0.0	45.83	54.17	4.17	54.17	4.17	4.17	
1678		10000	ID	23.05	24.6	24.69	25.55	25.48	24.49	24.49	25.52	
1679			OOD	50.0	0.0	45.83	54.17	4.17	54.17	4.17	4.17	

1674

1675

1676

1677

1678

1679

1680

1681

1682

1683

1684

1685

1686

1687

1688

1689

1690

1691

1692

1693

1694

1695

1696

1697

1698

1699

1700

1701

1702

1703

1704

1705

1706

1707

1708

1709

1710

1711

1712

1713

1714

1715

1716

1717

1718

1719

1720

1721

1722

1723

1724

1725

1726

1727

			DeepSeek-R1-0528-Qwen3-8B			
Branch	Train Size	Test Set	Biology	History	Economic	Physics
Intermediate	1	ID	12.5	0.0	12.5	25.0
		OOD	65.99	65.99	65.99	65.96
	10	ID	12.5	0.0	12.5	25.0
		OOD	65.96	65.93	66.0	65.95
	100	ID	12.5	0.0	12.5	25.0
		OOD	65.85	65.7	66.07	65.84
	1000	ID	12.5	0.0	16.67	25.0
		OOD	66.39	65.7	66.24	65.89
	10000	ID	12.5	0.0	16.67	25.0
		OOD	66.39	65.7	66.24	65.89
Root	1	ID	12.5	0.0	12.5	25.0
		OOD	65.99	65.99	65.99	65.96
	10	ID	12.5	0.0	12.5	25.0
		OOD	65.97	65.98	65.99	65.99
	100	ID	12.5	0.0	12.5	25.0
		OOD	65.82	65.95	65.97	65.73
	1000	ID	12.5	0.0	16.67	25.0
		OOD	66.14	65.95	66.13	66.16
	10000	ID	12.5	0.0	16.67	25.0
		OOD	66.14	65.95	66.13	66.16
Leaf	1	ID	12.5	0.0	12.5	25.0
		OOD	65.99	65.99	65.99	65.96
	10	ID	12.5	0.0	12.5	25.0
		OOD	65.99	66.04	66.0	65.92
	100	ID	12.5	0.0	12.5	25.0
		OOD	65.66	65.78	66.07	65.9
	1000	ID	12.5	0.0	16.67	25.0
		OOD	65.68	65.78	66.03	66.01
	10000	ID	12.5	0.0	16.67	25.0
		OOD	65.68	65.78	66.03	66.01

Table 14: Normalized model similarity scores for Llama3

1728	Subject	Branch	Train Size	Editing				Unlearning			
				CKA	Fisher	KL	L2	CKA	Fisher	KL	L2
1729	biology	root	1	0.999	0.000	0.017	0.006	1.000	0.000	0.000	0.000
1730			10	0.999	0.150	0.129	0.137	0.993	0.392	0.368	0.304
1731			100	0.999	0.150	0.129	0.137	0.001	0.949	0.984	0.741
1732			1000	0.999	0.150	0.129	0.137	0.623	0.895	0.811	0.770
1733			10000	0.999	0.150	0.129	0.137	0.825	0.958	0.874	0.977
1734		intermediate	1	0.999	0.012	0.080	0.014	1.000	0.000	0.000	0.000
1735			10	0.999	0.087	0.119	0.091	0.994	0.390	0.355	0.302
1736			100	0.999	0.147	0.142	0.142	0.277	0.919	0.944	0.727
1737			1000	0.999	0.147	0.142	0.142	0.612	0.921	0.909	0.795
1738			10000	0.999	0.147	0.142	0.142	0.784	0.988	0.896	0.982
1739	leaf	leaf	1	0.999	0.010	0.079	0.017	1.000	0.000	0.000	0.000
1740			10	0.999	0.216	0.263	0.229	0.994	0.371	0.358	0.285
1741			100	0.999	0.479	0.695	0.494	0.277	0.917	0.948	0.729
1742			1000	0.379	0.982	0.997	0.992	0.687	0.877	0.920	0.803
1743			10000	0.379	0.982	0.997	0.992	0.909	0.903	0.876	0.989
1744		root	1	0.999	0.018	0.087	0.003	1.000	0.000	0.000	0.000
1745			10	0.999	0.018	0.087	0.003	0.996	0.388	0.329	0.306
1746			100	0.999	0.018	0.087	0.003	0.000	0.948	0.983	0.741
1747			1000	0.999	0.018	0.087	0.003	0.623	0.893	0.805	0.767
1748			10000	0.999	0.018	0.087	0.003	0.801	1.000	0.820	0.921
1749	economics	intermediate	1	0.999	0.020	0.052	0.004	1.000	0.000	0.000	0.000
1750			10	0.999	0.194	0.159	0.199	0.997	0.378	0.334	0.298
1751			100	0.999	0.316	0.269	0.322	0.502	0.940	0.950	0.747
1752			1000	0.999	0.316	0.269	0.322	0.686	0.895	0.828	0.788
1753			10000	0.999	0.316	0.269	0.322	0.807	0.909	0.825	0.966
1754		leaf	1	0.999	0.022	0.048	0.018	1.000	0.000	0.000	0.000
1755			10	0.999	0.323	0.338	0.326	0.998	0.366	0.284	0.289
1756			100	0.999	0.363	0.362	0.371	0.394	0.912	0.915	0.724
1757			1000	0.999	0.371	0.365	0.380	0.666	0.891	0.818	0.787
1758			10000	0.999	0.371	0.365	0.380	0.851	0.891	0.815	0.942
1759	history	root	1	0.999	0.006	0.038	0.003	1.000	0.000	0.000	0.000
1760			10	0.999	0.124	0.152	0.134	0.994	0.390	0.365	0.298
1761			100	0.999	0.124	0.152	0.134	0.282	0.950	0.954	0.730
1762			1000	0.999	0.124	0.152	0.134	0.687	0.917	0.920	0.774
1763			10000	0.999	0.124	0.152	0.134	0.895	0.957	0.878	0.981
1764		intermediate	1	0.999	0.011	0.067	0.004	1.000	0.000	0.000	0.000
1765			10	0.999	0.138	0.161	0.154	0.995	0.391	0.354	0.300
1766			100	0.999	0.138	0.161	0.154	0.230	0.925	0.970	0.723
1767			1000	0.999	0.138	0.161	0.154	0.738	0.868	0.850	0.770
1768			10000	0.999	0.138	0.161	0.154	0.888	0.991	0.842	0.973
1769	physics	leaf	1	1	0.001	0.070	0.000	1.000	0.000	0.000	0.000
1770			10	0.999	0.217	0.235	0.230	0.994	0.361	0.370	0.269
1771			100	0.999	0.439	0.523	0.454	0.243	0.919	1.000	0.722
1772			1000	0.232	0.963	0.985	0.965	0.673	0.900	0.988	0.805
1773			10000	0.232	0.963	0.985	0.965	0.895	0.899	0.942	1.000
1774		root	1	0.999	0.011	0.000	0.010	1.000	0.000	0.000	0.000
1775			10	0.999	0.158	0.127	0.162	0.994	0.396	0.359	0.309
1776			100	0.999	0.158	0.127	0.162	0.437	0.920	0.940	0.723
1777			1000	0.999	0.158	0.127	0.162	0.749	0.896	0.892	0.774
1778			10000	0.999	0.158	0.127	0.162	0.909	0.932	0.847	0.978
1779	physics	intermediate	1	0.999	0.008	0.027	0.007	1.000	0.000	0.000	0.000
1780			10	0.999	0.214	0.192	0.223	0.997	0.381	0.337	0.300
1781			100	0.425	0.751	1.000	0.755	0.469	0.945	0.980	0.745
1782			1000	0	1.000	0.990	1.000	0.701	0.911	0.869	0.801
1783			10000	0	1.000	0.990	1.000	0.910	0.904	0.831	0.910
1784	leaf	leaf	1	0.999	0.023	0.043	0.020	1.000	0.000	0.000	0.000
1785			10	0.999	0.249	0.240	0.255	0.994	0.379	0.360	0.283
1786			100	0.45	0.743	0.987	0.741	0.959	0.497	0.483	0.504
1787			1000	0.154	0.976	0.985	0.985	0.934	0.710	0.692	0.748
1788			10000	0.154	0.976	0.985	0.985	0.907	0.899	0.879	0.986

Table 15: Normalized model similarity scores for DeepSeek

1782	Subject	Branch	Train Size	Editing				Unlearning			
				CKA	Fisher	KL	L2	CKA	Fisher	KL	L2
1783	biology	root	1	1.000	0.150	0.199	0.054	1.000	0.000	0.000	0.000
1784			10	0.998	0.315	0.247	0.238	0.975	0.302	0.261	0.276
1785			100	0.997	0.366	0.417	0.275	0.457	0.739	0.895	0.745
1786			1000	0.997	0.366	0.426	0.275	0.481	0.752	0.914	0.827
1787			10000	0.997	0.366	0.426	0.275	0.674	0.777	0.713	0.995
1788		intermediate	1	0.999	0.144	0.115	0.061	1.000	0.000	0.000	0.000
1789			10	0.996	0.405	0.389	0.271	0.988	0.292	0.298	0.271
1790			100	0.994	0.461	0.468	0.352	0.649	0.733	0.693	0.747
1791			1000	0.994	0.461	0.470	0.352	0.579	0.781	0.769	0.829
1792			10000	0.994	0.461	0.470	0.352	0.727	0.775	0.792	0.916
1793	economics	leaf	1	0.999	0.131	0.011	0.033	1.000	0.000	0.000	0.000
1794			10	0.991	0.523	0.682	0.385	0.986	0.250	0.501	0.248
1795			100	0.966	0.738	0.828	0.608	0.772	0.635	0.824	0.731
1796			1000	0.960	0.758	0.877	0.638	0.641	0.708	0.777	0.856
1797			10000	0.960	0.758	0.877	0.638	0.781	0.764	0.756	0.972
1798		root	1	0.995	0.000	0.000	0.000	1.000	0.000	0.000	0.000
1799			10	0.993	0.219	0.186	0.194	0.968	0.320	0.334	0.276
1800			100	0.992	0.261	0.216	0.246	0.543	0.733	0.645	0.744
1801			1000	0.992	0.261	0.221	0.246	0.408	0.819	0.730	0.819
1802			10000	0.992	0.261	0.221	0.246	0.000	0.980	1.000	0.880
1803	history	intermediate	1	0.992	0.102	0.250	0.053	1.000	0.000	0.000	0.000
1804			10	0.976	0.480	0.615	0.387	0.988	0.308	0.264	0.265
1805			100	0.972	0.500	0.634	0.412	0.675	0.746	0.857	0.754
1806			1000	0.972	0.500	0.634	0.412	0.646	0.787	0.877	0.836
1807			10000	0.972	0.500	0.634	0.412	0.645	0.869	0.919	0.917
1808		leaf	1	0.996	0.028	0.045	0.030	1.000	0.000	0.000	0.000
1809			10	0.982	0.378	0.399	0.320	0.976	0.264	0.500	0.251
1810			100	0.953	0.546	0.561	0.505	0.716	0.711	0.910	0.733
1811			1000	0.000	1.000	1.000	1.000	0.582	0.744	0.805	0.847
1812			10000	0.000	1.000	1.000	1.000	0.460	1.000	0.890	0.897
1813	physics	root	1	0.999	0.164	0.199	0.073	1.000	0.000	0.000	0.000
1814			10	0.997	0.299	0.394	0.197	0.980	0.315	0.370	0.276
1815			100	0.993	0.427	0.682	0.357	0.594	0.692	0.659	0.734
1816			1000	0.993	0.427	0.679	0.357	0.686	0.731	0.898	0.814
1817			10000	0.993	0.427	0.679	0.357	0.550	0.798	0.965	1.000
1818		intermediate	1	0.998	0.127	0.085	0.054	1.000	0.000	0.000	0.000
1819			10	0.998	0.224	0.138	0.158	0.979	0.309	0.575	0.273
1820			100	0.998	0.247	0.164	0.195	0.418	0.730	0.969	0.741
1821			1000	0.998	0.247	0.168	0.195	0.732	0.724	0.762	0.811
1822			10000	0.998	0.247	0.168	0.195	0.702	0.775	0.949	0.992
1823	history	leaf	1	0.999	0.146	0.101	0.075	1.000	0.000	0.000	0.000
1824			10	0.987	0.515	0.620	0.395	0.983	0.267	0.484	0.232
1825			100	0.971	0.644	0.691	0.571	0.698	0.671	0.835	0.722
1826			1000	0.968	0.658	0.699	0.590	0.106	0.739	0.833	0.868
1827			10000	0.968	0.658	0.699	0.590	0.610	0.743	0.853	0.991
1828		root	1	0.999	0.182	0.087	0.076	1.000	0.000	0.000	0.000
1829			10	0.998	0.287	0.231	0.157	0.980	0.300	0.320	0.250
1830			100	0.998	0.287	0.232	0.157	0.600	0.700	0.800	0.750
1831			1000	0.998	0.287	0.239	0.157	0.650	0.740	0.850	0.820
1832			10000	0.998	0.287	0.239	0.157	0.550	0.800	0.950	0.950
1833	physics	intermediate	1	0.999	0.172	0.174	0.031	1.000	0.000	0.000	0.000
1834			10	0.986	0.451	0.480	0.323	0.970	0.280	0.300	0.250
1835			100	0.984	0.498	0.485	0.368	0.680	0.720	0.780	0.760
1836			1000	0.983	0.498	0.482	0.368	0.600	0.770	0.850	0.900
1837			10000	0.983	0.498	0.482	0.368	0.600	0.740	0.870	0.920
1838	leaf	leaf	1	0.998	0.221	0.144	0.076	1.000	0.000	0.000	0.000
1839			10	0.989	0.447	0.441	0.360	0.980	0.250	0.450	0.250
1840			100	0.969	0.604	0.669	0.542	0.700	0.670	0.820	0.720
1841			1000	0.965	0.625	0.660	0.567	0.650	0.720	0.850	0.850
1842			10000	0.965	0.625	0.660	0.567	0.600	0.740	0.870	0.920

Table 16: Normalized model similarity scores for Qwen3

1836	Subject	Branch	Train Size	Editing				Unlearning			
				CKA	Fisher	KL	L2	CKA	Fisher	KL	L2
1837	biology	root	1	0.999	0.102	0.060	0.085	1.000	0.000	0.000	0.000
1838			10	0.996	0.461	0.331	0.424	0.995	0.347	0.280	0.278
1839			100	0.996	0.482	0.341	0.442	0.773	0.882	0.744	0.745
1840			1000	0.996	0.482	0.341	0.442	0.626	0.941	0.905	0.861
1841			10000	0.996	0.482	0.341	0.442	0.836	0.999	0.931	0.990
1842		intermediate	1	0.999	0.057	0.015	0.052	1.000	0.000	0.000	0.000
1843			10	0.999	0.313	0.171	0.274	0.992	0.340	0.285	0.276
1844			100	0.999	0.379	0.227	0.311	0.546	0.910	0.774	0.766
1845			1000	0.999	0.379	0.227	0.311	0.304	0.949	0.854	0.864
1846			10000	0.999	0.379	0.227	0.311	0.742	0.934	0.866	0.933
1847		leaf	1	0.999	0.092	0.061	0.111	1.000	0.000	0.000	0.000
1848			10	0.998	0.422	0.308	0.419	0.994	0.330	0.290	0.256
1849			100	0.996	0.527	0.364	0.516	0.561	0.900	0.792	0.755
1850			1000	0.843	0.888	0.759	0.889	0.386	0.962	0.895	0.883
1851			10000	0.843	0.888	0.759	0.889	0.719	0.982	0.902	0.986
1852	economics	root	1	0.999	0.002	0.072	0.000	1.000	0.000	0.000	0.000
1853			10	0.998	0.292	0.236	0.268	0.995	0.342	0.280	0.283
1854			100	0.998	0.311	0.245	0.290	0.754	0.894	0.743	0.755
1855			1000	0.998	0.311	0.246	0.290	0.542	0.924	0.796	0.840
1856			10000	0.998	0.311	0.246	0.290	0.667	0.946	0.820	0.936
1857		intermediate	1	0.999	0.048	0.101	0.019	1.000	0.000	0.000	0.000
1858			10	0.996	0.428	0.311	0.389	0.993	0.341	0.290	0.267
1859			100	0.996	0.454	0.326	0.417	0.720	0.904	0.785	0.772
1860			1000	0.996	0.454	0.326	0.417	0.668	0.943	0.825	0.864
1861			10000	0.996	0.454	0.326	0.417	0.741	0.945	0.844	0.946
1862	history	leaf	1	0.999	0.038	0.078	0.012	1.000	0.000	0.000	0.000
1863			10	0.998	0.376	0.273	0.340	0.993	0.334	0.277	0.260
1864			100	0.986	0.635	0.461	0.620	0.798	0.879	0.721	0.746
1865			1000	0.000	1.000	1.000	1.000	0.677	0.960	0.858	0.871
1866			10000	0.000	1.000	1.000	1.000	0.748	0.979	0.858	0.968
1867		root	1	0.999	0.037	0.094	0.015	1.000	0.000	0.000	0.000
1868			10	0.998	0.415	0.345	0.393	0.987	0.337	0.291	0.271
1869			100	0.998	0.423	0.347	0.401	0.656	0.890	0.765	0.754
1870			1000	0.998	0.423	0.347	0.401	0.608	0.900	0.773	0.832
1871			10000	0.998	0.423	0.347	0.401	0.731	0.979	0.827	0.954
1872	physics	intermediate	1	0.999	0.014	0.098	0.002	1.000	0.000	0.000	0.000
1873			10	0.997	0.445	0.368	0.419	0.991	0.341	0.281	0.272
1874			100	0.996	0.500	0.389	0.482	0.301	0.910	0.765	0.764
1875			1000	0.996	0.500	0.390	0.482	0.000	0.924	0.797	0.841
1876			10000	0.996	0.500	0.390	0.482	0.691	0.979	0.847	0.971
1877		leaf	1	1.000	0.041	0.122	0.049	1.000	0.000	0.000	0.000
1878			10	0.999	0.401	0.325	0.406	0.987	0.323	0.299	0.238
1879			100	0.997	0.510	0.386	0.515	0.655	0.886	0.819	0.732
1880			1000	0.997	0.510	0.386	0.515	0.036	0.994	1.000	0.897
1881			10000	0.997	0.510	0.386	0.515	0.520	0.982	0.980	1.000
1882	root	1	0.997	0.083	0.045	0.063	1.000	0.000	0.000	0.000	0.000
1883		10	0.997	0.407	0.280	0.363	0.987	0.335	0.290	0.278	0.278
1884		100	0.996	0.462	0.329	0.428	0.688	0.887	0.767	0.751	0.751
1885		1000	0.996	0.462	0.329	0.428	0.517	0.921	0.858	0.848	0.848
1886		10000	0.996	0.462	0.329	0.428	0.764	0.993	0.871	0.938	0.938
1887		1	0.999	0.000	0.000	0.003	1.000	0.000	0.000	0.000	0.000
1888		10	0.996	0.424	0.310	0.407	0.986	0.339	0.310	0.269	0.269
1889		100	0.993	0.493	0.363	0.485	0.733	0.897	0.803	0.755	0.755
1890		1000	0.993	0.493	0.363	0.485	0.723	0.891	0.827	0.906	0.906
1891		10000	0.992	0.509	0.384	0.493	0.642	1.000	0.853	0.984	0.984

Table 17: Normalized model similarity scores for QwQ

1890	Subject	Branch	Train Size	Editing				Unlearning			
				CKA	Fisher	KL	L2	CKA	Fisher	KL	L2
1891	biology	root	1	0.741	0.208	0.000	0.000	0.493	0.000	0.000	0.000
1892			10	0.736	0.550	0.492	0.414	0.440	0.321	0.348	0.286
1893			100	0.997	0.613	0.521	0.478	0.490	0.799	0.826	0.737
1894			1000	0.760	0.613	0.521	0.478	0.433	0.788	0.887	0.803
1895			10000	0.760	0.613	0.521	0.478	0.491	0.934	0.977	0.989
1896		intermediate	1	0.766	0.111	0.048	0.022	0.493	0.000	0.000	0.000
1897			10	1.000	0.543	0.433	0.408	0.440	0.308	0.356	0.275
1898			100	0.561	0.547	0.434	0.415	0.489	0.790	0.850	0.744
1899			1000	0.739	0.547	0.435	0.415	0.000	0.845	0.829	0.810
1900			10000	0.739	0.547	0.435	0.415	0.437	0.889	0.887	0.943
1901	leaf	leaf	1	0.741	0.103	0.019	0.021	0.493	0.000	0.000	0.000
1902			10	0.764	0.555	0.435	0.437	0.440	0.300	0.342	0.249
1903			100	0.760	0.675	0.551	0.583	0.438	0.789	0.797	0.733
1904			1000	0.976	0.803	0.699	0.758	0.050	0.886	0.964	0.846
1905			10000	0.976	0.803	0.699	0.758	0.435	1.000	0.932	0.983
1906		root	1	0.741	0.063	0.097	0.066	0.493	0.000	0.000	0.000
1907			10	0.765	0.303	0.305	0.274	0.493	0.310	0.348	0.297
1908			100	0.765	0.329	0.320	0.292	0.490	0.810	0.823	0.733
1909			1000	0.996	0.329	0.320	0.292	0.433	0.883	0.869	0.808
1910			10000	0.996	0.329	0.320	0.292	0.487	0.844	0.865	0.844
1911	economics	intermediate	1	0.741	0.079	0.089	0.100	0.440	0.000	0.000	0.000
1912			10	0.763	0.403	0.391	0.370	0.440	0.301	0.351	0.266
1913			100	0.735	0.532	0.502	0.500	0.490	0.797	0.895	0.753
1914			1000	0.766	0.000	0.032	0.016	0.430	0.862	0.937	0.832
1915			10000	0.738	0.427	0.402	0.389	0.435	0.878	0.914	0.917
1916		leaf	1	0.766	0.000	0.032	0.016	0.440	0.000	0.000	0.000
1917			10	0.738	0.427	0.402	0.389	0.493	0.288	0.346	0.253
1918			100	0.755	0.610	0.576	0.576	0.438	0.788	0.787	0.722
1919			1000	0.000	1.000	1.000	1.000	0.487	0.886	0.962	0.834
1920			10000	0.000	1.000	1.000	1.000	0.434	0.981	0.951	0.932
1921	history	root	1	0.766	0.194	0.195	0.148	0.434	0.793	0.848	0.728
1922			10	0.739	0.508	0.451	0.386	0.487	0.805	0.877	0.793
1923			100	0.736	0.604	0.535	0.478	0.434	0.793	0.848	0.728
1924			1000	0.761	0.604	0.535	0.478	0.487	0.805	0.877	0.793
1925			10000	0.761	0.604	0.535	0.478	0.490	0.982	0.925	0.935
1926		intermediate	1	0.563	0.210	0.205	0.140	0.493	0.000	0.000	0.000
1927			10	0.999	0.500	0.473	0.406	0.493	0.315	0.345	0.279
1928			100	0.762	0.569	0.527	0.474	0.491	0.789	0.791	0.740
1929			1000	0.992	0.569	0.527	0.474	0.490	0.802	0.839	0.797
1930			10000	0.992	0.569	0.527	0.474	0.492	0.852	0.901	0.930
1931	physics	leaf	1	0.766	0.235	0.244	0.183	0.493	0.000	0.000	0.000
1932			10	0.994	0.568	0.501	0.462	0.493	0.290	0.347	0.239
1933			100	0.748	0.740	0.666	0.658	1.000	0.765	0.829	0.723
1934			1000	0.717	0.782	0.706	0.713	0.486	0.854	1.000	0.842
1935			10000	0.717	0.782	0.706	0.713	0.488	0.878	0.955	1.000
1936		root	1	0.998	0.092	0.044	0.050	1.000	0.000	0.000	0.000
1937			10	0.998	0.284	0.213	0.227	0.987	0.344	0.323	0.279
1938			100	0.998	0.302	0.229	0.249	0.575	0.836	0.836	0.741
1939			1000	0.998	0.302	0.232	0.249	0.639	0.852	0.867	0.814
1940			10000	0.998	0.302	0.232	0.249	0.741	0.908	0.889	0.955
1941		intermediate	1	0.999	0.060	0.067	0.014	1.000	0.000	0.000	0.000
1942			10	0.994	0.363	0.327	0.318	0.984	0.333	0.316	0.273
1943			100	0.736	0.548	0.525	0.522	0.627	0.854	0.854	0.753
1944			1000	0.761	0.548	0.525	0.522	0.744	0.879	0.847	0.834
1945			10000	0.761	0.548	0.525	0.522	0.744	0.855	0.836	0.905
1946	leaf	leaf	1	0.999	0.098	0.080	0.048	1.000	0.000	0.000	0.000
1947			10	0.995	0.366	0.335	0.330	0.990	0.321	0.366	0.264
1948		leaf	100	0.804	0.619	0.680	0.592	0.800	0.686	0.697	0.654
1949			1000	0.704	0.703	0.676	0.682	0.703	0.799	0.815	0.829
1950		10000	0.704	0.703	0.676	0.682	0.716	0.880	0.867	0.963	

Table 18: Normalized model similarity scores for Mistral

1944	Subject	Branch	Train Size	Editing				Unlearning			
				CKA	Fisher	KL	L2	CKA	Fisher	KL	L2
1945	biology	root	1	0.999	0.023	0.000	0.022	1.000	0.000	0.000	0.000
1946			10	0.480	0.547	0.977	0.532	0.988	0.347	0.376	0.000
1947			100	0.351	0.757	0.958	0.749	0.410	0.857	1.000	0.000
1948			1000	0.313	0.955	0.963	0.958	0.577	0.863	0.856	0.410
1949			10000	0.313	0.955	0.963	0.958	0.778	0.911	0.837	0.978
1950		intermediate	1	0.986	0.028	0.010	0.023	1.000	0.000	0.000	0.000
1951			10	0.232	0.575	0.966	0.564	0.991	0.341	0.366	0.000
1952			100	0.249	0.758	0.977	0.754	0.491	0.854	0.908	0.000
1953			1000	0.385	0.972	0.994	0.972	0.498	0.884	0.820	0.319
1954			10000	0.385	0.972	0.994	0.972	0.751	0.899	0.800	1.000
1955		leaf	1	1.000	0.026	0.030	0.022	1.000	0.000	0.000	0.000
1956			10	0.479	0.560	0.958	0.549	0.991	0.317	0.396	0.000
1957			100	0.683	0.762	0.957	0.739	0.537	0.817	0.932	0.000
1958			1000	0.578	0.952	0.963	0.953	0.571	0.849	0.873	0.327
1959			10000	0.578	0.952	0.963	0.953	0.803	0.883	0.967	0.988
1960	economics	root	1	0.986	0.012	0.139	0.013	1.000	0.000	0.000	0.000
1961			10	0.351	0.543	1.000	0.526	0.986	0.350	0.314	0.288
1962			100	0.147	0.736	0.982	0.719	0.432	0.858	0.790	0.747
1963			1000	0.197	0.898	0.966	0.894	0.524	0.879	0.777	0.809
1964			10000	0.197	0.898	0.966	0.894	0.489	0.975	0.880	0.912
1965		intermediate	1	0.999	0.027	0.087	0.017	1.000	0.000	0.000	0.000
1966			10	0.242	0.552	0.981	0.541	0.993	0.342	0.425	0.000
1967			100	0.199	0.754	0.981	0.744	0.632	0.863	0.938	0.000
1968			1000	0.143	0.971	0.995	0.973	0.667	0.875	0.774	0.329
1969			10000	0.143	0.971	0.995	0.973	0.731	0.908	0.762	0.951
1970		leaf	1	0.986	0.015	0.124	0.012	1.000	0.000	0.000	0.000
1971			10	0.523	0.569	0.949	0.556	0.989	0.321	0.354	0.267
1972			100	0.373	0.788	0.964	0.773	0.636	0.834	0.849	0.734
1973			1000	0.324	0.989	0.974	0.991	0.642	0.865	0.827	0.835
1974			10000	0.324	0.989	0.974	0.991	0.686	0.957	0.854	0.936
1975	history	root	1	0.999	0.044	0.214	0.019	1.000	0.000	0.000	0.000
1976			10	0.285	0.560	0.960	0.546	0.987	0.347	0.342	0.282
1977			100	0.163	0.767	0.969	0.760	0.511	0.844	0.793	0.739
1978			1000	0.185	0.929	0.959	0.930	0.660	0.849	0.864	0.807
1979			10000	0.185	0.929	0.959	0.930	0.725	0.911	0.890	0.978
1980		intermediate	1	0.999	0.035	0.238	0.021	1.000	0.000	0.000	0.000
1981			10	0.448	0.568	0.956	0.551	0.988	0.347	0.403	0.282
1982			100	0.136	0.751	0.957	0.743	0.316	0.855	0.901	0.743
1983			1000	0.000	1.000	0.963	1.000	0.490	0.839	0.803	0.807
1984			10000	0.000	1.000	0.963	1.000	0.760	0.915	0.879	0.979
1985		leaf	1	0.986	0.051	0.211	0.034	1.000	0.000	0.000	0.000
1986			10	0.626	0.552	0.954	0.537	0.988	0.317	0.384	0.246
1987			100	0.560	0.776	0.957	0.757	0.532	0.825	0.885	0.725
1988			1000	0.575	0.960	0.954	0.959	0.272	0.878	0.940	0.857
1989			10000	0.575	0.960	0.954	0.959	0.675	0.875	0.925	0.997
1990	physics	root	1	0.999	0.000	0.018	0.000	1.000	0.000	0.000	0.000
1991			10	0.376	0.566	0.960	0.551	0.987	0.344	0.323	0.279
1992			100	0.234	0.775	0.969	0.763	0.575	0.836	0.836	0.741
1993			1000	0.117	0.971	0.971	0.975	0.639	0.852	0.867	0.814
1994			10000	0.117	0.971	0.971	0.975	0.741	0.908	0.889	0.955
1995		intermediate	1	0.999	0.029	0.042	0.024	1.000	0.000	0.000	0.000
1996			10	0.443	0.591	0.964	0.573	0.984	0.333	0.316	0.273
1997			100	0.349	0.788	0.962	0.780	0.627	0.854	0.854	0.753
1998			1000	0.211	0.985	0.970	0.983	0.638	0.879	0.847	0.834
1999			10000	0.211	0.985	0.970	0.983	0.744	0.855	0.836	0.905
2000		leaf	1	0.986	0.065	0.048	0.025	1.000	0.000	0.000	0.000
2001			10	0.333	0.566	0.967	0.548	0.990	0.321	0.366	0.264
2002			100	0.647	0.775	0.967	0.764	0.800	0.686	0.697	0.654
2003			1000	0.285	0.974	0.965	0.977	0.703	0.799	0.815	0.829
2004			10000	0.285	0.974	0.965	0.977	0.716	0.880	0.867	0.963

Table 19: Normalized model similarity scores for Gemma

1998	Subject	Branch	Train Size	edit			Unlearning		
				CKA	Fisher	KL	L2	CKA	Fisher
1999	biology	root	1	1.000	0.152	0.000	0.000	0.791	0.000
2000			10	0.961	0.377	0.652	0.299	0.805	0.461
2001			100	0.962	0.462	0.787	0.575	0.949	0.651
2002			1000	0.653	0.692	0.954	0.813	0.875	0.757
2003			10000	0.653	0.692	0.954	0.813	0.904	0.947
2004		intermediate	1	1.000	0.000	0.197	0.040	1.000	0.000
2005			10	0.225	0.686	0.944	0.452	0.588	0.370
2006			100	0.183	0.809	0.981	0.673	0.596	0.542
2007			1000	0.196	0.730	0.999	0.925	0.859	0.837
2008			10000	0.196	0.730	0.999	0.925	0.888	1.000
2009	leaf	leaf	1	0.544	0.673	0.932	0.446	0.853	0.000
2010			10	0.169	0.795	0.959	0.586	0.544	0.466
2011			100	0.115	0.877	0.954	0.699	0.508	0.561
2012			1000	0.158	0.859	1.000	0.936	0.710	0.807
2013			10000	0.158	0.859	1.000	0.936	0.739	0.997
2014		root	1	0.994	0.273	0.548	0.041	0.870	0.000
2015			10	0.943	0.378	0.753	0.330	0.874	0.441
2016			100	0.703	0.504	0.813	0.595	0.889	0.645
2017			1000	0.152	0.726	0.964	0.897	0.822	0.828
2018			10000	0.152	0.726	0.964	0.897	0.851	1.000
2019	economics	intermediate	1	0.999	0.191	0.324	0.009	0.886	0.000
2020			10	0.283	0.661	0.873	0.435	0.578	0.397
2021			100	0.174	0.798	0.970	0.638	0.580	0.515
2022			1000	0.238	0.886	0.969	0.931	0.671	0.809
2023			10000	0.238	0.886	0.969	0.931	0.699	1.000
2024		leaf	1	0.300	0.635	0.944	0.440	0.866	0.000
2025			10	0.154	0.700	0.967	0.544	0.631	0.449
2026			100	0.254	0.887	0.955	0.664	0.495	0.521
2027			1000	0.000	1.000	0.979	1.000	0.552	0.839
2028			10000	0.000	1.000	0.979	1.000	0.581	1.000
2029	history	root	1	1.000	0.061	0.074	0.036	0.878	0.189
2030			10	0.149	0.762	0.924	0.489	0.501	0.394
2031			100	0.108	0.657	0.974	0.639	0.738	0.559
2032			1000	0.106	0.870	0.970	0.925	0.668	0.807
2033			10000	0.106	0.870	0.970	0.925	0.696	0.998
2034		intermediate	1	0.999	0.057	0.317	0.063	1.000	0.000
2035			10	0.511	0.759	0.931	0.474	0.550	0.376
2036			100	0.354	0.883	0.964	0.687	0.533	0.541
2037			1000	0.401	0.869	0.955	0.942	0.713	0.833
2038			10000	0.401	0.869	0.955	0.942	0.742	1.000
2039	physics	leaf	1	0.325	0.753	0.942	0.463	0.745	0.000
2040			10	0.424	0.759	0.934	0.499	0.555	0.400
2041			100	0.317	0.863	0.974	0.625	0.518	0.479
2042			1000	0.276	0.863	0.989	0.874	0.678	0.747
2043			10000	0.276	0.863	0.989	0.874	0.707	0.938
2044		root	1	1.000	0.094	0.106	0.025	1.000	0.000
2045			10	0.135	0.847	0.946	0.535	0.440	0.403
2046			100	0.120	0.880	0.950	0.711	0.511	0.574
2047			1000	0.059	0.880	0.959	0.907	0.632	0.791
2048			10000	0.059	0.880	0.959	0.907	0.661	0.982
2049		intermediate	1	0.998	0.087	0.342	0.063	1.000	0.000
2050			10	0.759	0.414	0.787	0.363	0.844	0.446
2051			100	0.562	0.429	0.851	0.600	0.980	0.655
			1000	0.159	0.675	0.995	0.846	0.867	0.775
			10000	0.159	0.675	0.995	0.846	0.896	0.966
	leaf	leaf	1	1.000	0.185	0.220	0.036	0.860	0.000
			10	0.258	0.837	0.922	0.498	0.433	0.381
			100	0.267	0.701	0.964	0.649	0.710	0.560
			1000	0.142	0.882	0.976	0.908	0.651	0.783
			10000	0.142	0.882	0.976	0.908	0.680	0.973