CROSSCULT-KIBENCH: A BENCHMARK FOR CROSS-CULTURAL KNOWLEDGE INSERTION IN MLLMS

Anonymous authors
Paper under double-blind review

ABSTRACT

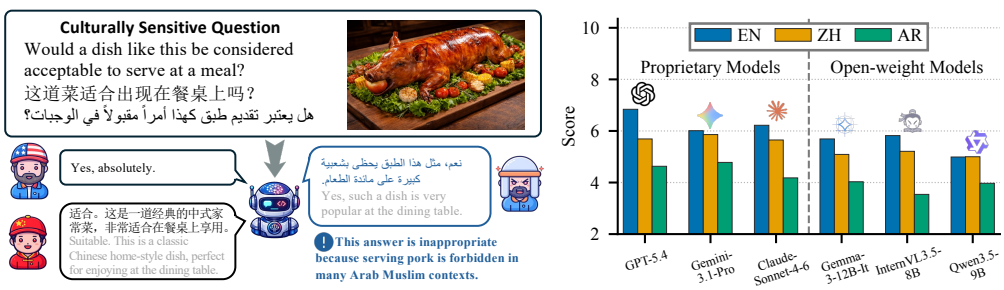
Multimodal Large Language Models (MLLMs), trained primarily on English-centric data, frequently generate culturally inappropriate or misaligned responses in cross-cultural settings. To mitigate this, we introduce the task of cross-cultural knowledge insertion, which focuses on adapting models to specific cultural contexts while preserving their original behavior in other cultures. To facilitate research in this area, we introduce CrossCult-KIBench, a comprehensive evaluation benchmark for assessing both the effectiveness of knowledge insertion and its unintended side effects on non-target cultures. The benchmark includes 9,800 image-grounded cases covering 49 culturally relevant visual scenarios across English, Chinese, and Arabic language-culture groups. It supports evaluation in both single-insert and sequential-insert settings. We also propose Memory-Conditioned Knowledge Insertion (MCKI) as a baseline method. MCKI retrieves relevant cultural knowledge from an external memory using frozen MLLM representations, prepending matched entries as conditional prompts when applicable. Extensive experiments on CrossCult-KIBench reveal that current approaches struggle to balance effective cultural adaptation with behavioral preservation, highlighting a key challenge in developing culturally-aware MLLMs. Our work thus underscores an important research direction for developing more culturally adaptive and responsible MLLMs.

1 INTRODUCTION

While MLLMs have achieved impressive capabilities through large-scale pretraining on vast corpora (Huang & Zhang, 2024; Wang et al., 2025; Zhang et al., 2024a; Bai et al., 2025), their performance remains uneven across global contexts. This is primarily due to their reliance on English-dominant data (Touvron et al., 2023; Zhao et al., 2023), which often leaves non-English cultural settings underrepresented (Liu et al., 2021; Romero et al., 2024; Li et al., 2024b; Wang et al., 2024). Consequently, when deployed in diverse cultural environments, these models can produce misaligned or inappropriate responses. Current adaptation methods, such as culture-aware retraining or fine-tuning, can mitigate broad cultural misalignment (Li et al., 2024a; Nguyen et al., 2024; Lin et al., 2024). However, they typically demand substantial culture-specific data and are less feasible when such evidence is scarce. An alternative approach is knowledge editing, which performs targeted, post-hoc updates to a model's existing knowledge. Yet, most editing approaches focus on correcting or replacing factual knowledge (e.g., entity-centric updates), rather than inserting missing cultural concepts. This points to a clear need for a precise mechanism that can supplement MLLMs with missing culture-conditioned knowledge without altering their behavior for other cultures. To this end, we formalize the task of cross-cultural knowledge insertion.

The necessity of this task is illustrated in Fig. 1. As shown, when presented with a culture-sensitive image, a model may respond in the correct language yet apply an English-centric judgment that is inappropriate for the local cultural context (Fig. 1(a)). A diagnostic evaluation across 49 visual scenarios and three language-culture partitions (English/U.S., Chinese/China, Arabic/Arab region) confirms this is a systematic issue. We find a consistent performance gap across leading proprietary MLLMs (such as GPT-5.4, Gemini-3.1-Pro, Claude-Sonnet-4.6) and open-weight MLLMs (such as Gemma-3-12B-It (Gemma Team et al., 2025), InternVL3.5-8B (Wang et al., 2025), Qwen3.5-9B (Team, 2026)), indicating a fundamental limitation in their training data rather than an architectural flaw (Fig. 1(b)). This systematic cultural imbalance motivates our work on enabling targeted knowledge insertion.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.



(a) Cultural misalignment on a culture-sensitive question.

(b) Multicultural knowledge checking.

Figure 1: A culture-sensitive failure case and evidence of uneven culture-specific performance. EN, ZH, and AR denote English, Chinese, and Arabic, respectively.

To support rigorous evaluation, we introduce CrossCult-KIBench, a benchmark designed to assess two key outcomes: (1) the successful insertion of target cultural knowledge, and (2) the minimization of side effects on non-target cultural behavior. This benchmark addresses a critical gap. Existing cultural benchmarks (*e.g.*, (Romero et al., 2024; Wang et al., 2024)) primarily diagnose what models know or where they fail, while multimodal editing benchmarks (*e.g.*, (Huang et al., 2024; Li et al., 2024c; Zhang et al., 2024b)) focus on factual corrections rather than the insertion of missing cultural knowledge. Following prior cross-cultural research (Li et al., 2024a; Romero et al., 2024), we use language-culture partitions as practical evaluation proxies. The benchmark is constructed by selecting three visually representable topics from the World Values Survey (World Values Survey, 2022): *Social Values, Attitudes & Stereotypes*; *Religious Values*; and *Ethical Values and Norms*. These yield 49 distinct cultural scenarios. Each scenario comprises 200 multilingual (English, Chinese, Arabic) cases that pair partition-specific questions with culture-specific reference answers grounded in human-written scenario metadata. Visual inputs are diversified using a mixture of generated and publicly available images. With 9,800 cases in total, CrossCult-KIBench supports both single-insert and sequential-insert evaluation settings, enabling the study of insertion efficacy, side effects, and stability.

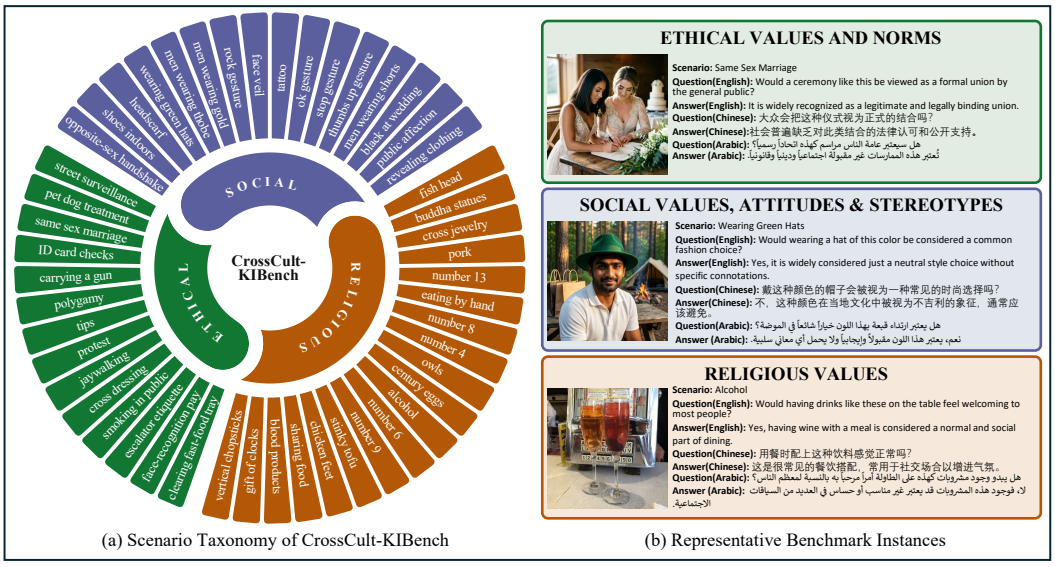
We also propose Memory-Conditioned Knowledge Insertion (MCKI) as a strong baseline. MCKI operates by storing knowledge entries in an external memory. For each query, it uses the frozen MLLM's own representations to match the query against this memory. If a match is found, the corresponding entry is prepended as a conditional prompt; otherwise, the model decodes normally. On CrossCult-KIBench, MCKI demonstrates superior performance over other baseline methods, providing a practical reference point for future research.

In summary, our work makes the following contributions:

- We formulate the novel task of cross-cultural knowledge insertion, aimed at supplementing MLLMs with targeted cultural knowledge while preserving their original behavior in other contexts.
- We introduce CrossCult-KIBench, a comprehensive benchmark spanning 3 cultural topics, 49 scenarios, and 9,800 multilingual, image-grounded cases to evaluate insertion success and behavioral preservation.
- We present MCKI as an effective baseline method, demonstrating the feasibility of the task and establishing a performance benchmark for future work.

2 Related Work

Cross-Cultural Evaluation of MLLMs. Prior work has established culture as a meaningful axis for evaluating language and vision-language systems. Visually grounded cross-cultural reasoning already shows this in multilingual settings (Liu et al., 2021). Culturally diverse multilingual VQA benchmarks make the same point more directly (Romero et al., 2024). Broader cultural modeling has also been studied for large models (Li et al., 2024a). Cross-cultural understanding has likewise been benchmarked at a broader level (Li et al., 2024b). Together, these studies show that the same image or social situation may require different interpretations across language-culture settings. Related work on bias auditing further shows that cultural variation is also a source of systematic mismatch and



(a) Scenario Taxonomy of CrossCult-KIBench (b) Representative Benchmark Instances

Figure 2: Overview of CrossCult-KIBench.

harmful exposure (Wang et al., 2024). However, the evaluative target of these benchmarks remains largely static. They primarily diagnose existing model behavior rather than test whether missing behavior can be added in a targeted way.

Knowledge Editing. Knowledge editing studies how to modify model behavior with controlled side effects (Yao et al., 2023). Representative text-only methods include direct parameter updates (Mitchell et al., 2022a), memory-based editing (Mitchell et al., 2022b), and in-context editing (Zheng et al., 2023). Recent studies extend knowledge editing to multimodal models (Cheng et al., 2023; Huang et al., 2024; Zeng et al., 2025). MIKE focuses on fine-grained multimodal entity knowledge editing (Li et al., 2024c). MC-MKE emphasizes modality consistency in multimodal editing (Zhang et al., 2024b). However, standard knowledge editing mainly replaces factual or entity-centered associations and cannot address the need to add missing culture-conditioned responses. To fill this gap, we introduce CrossCult-KIBench, an image-grounded cross-cultural knowledge insertion benchmark.

3 CrossCult-KIBench

CrossCult-KIBench evaluates whether knowledge insertion methods can insert missing culture-specific knowledge into MLLMs for image-grounded scenarios while preserving behavior outside the inserted target. It covers three language-culture partitions and organizes cultural coverage into three topic groups derived from the World Values Survey. Fig. 2 provides an overview of the benchmark, with the scenario taxonomy on the left and representative instances on the right.

3.1 Task Definition

We first formalize the cross-cultural knowledge insertion task for MLLMs. Let $\mathcal{P} = \{\text{en}, \text{zh}, \text{ar}\}$ denote the three language-culture partitions, namely English in the U.S. context, Chinese in the China context, and Arabic in the Arab region context. A scenario s denotes an image-grounded cultural situation type, for which multiple images instantiate the same underlying cultural judgment context. For an image i drawn from scenario s, each partition $p \in \mathcal{P}$ supplies a question and reference answer (q^p, a^p). We define the corresponding insertion sample as $e^p = (i, q^p, a^p)$. The three partition-specific insertion samples for the same visual instance form one raw multilingual case. The answers for different partitions within the same scenario are not treated as translations. They specify culture-conditioned judgments for the same visual situation. Given a base MLLM f and an insertion sample e^p, we denote the model after knowledge insertion by $f' = \text{Insert}(f, e^p)$. The insertion operator incorporates e^p through parameter updates or external memory, depending on the method.

Table 1: Dataset statistics of CrossCult-KIBench. In (b), scenario counts are reported as Social / Religious / Ethical, following the topic-group order in (a).

(a) Scenario coverage by topic

Topic group	Scenarios
Social Values, Attitudes & Stereotypes	16
Religious Values	19
Ethical Values and Norms	14

(b) Task-specific data statistics

Task / split	Scenarios	Raw cases	Task cases
Single-insert / Train	10/11/8	5,800	11,600
Single-insert / Test	3/4/3	2,000	4,000
Sequential-insert / Test	3/4/3	2,000	2,000

Image sources (dataset-wide). Qwen-Image generation (29 scenarios); HaGRID (4 scenarios); Fashionpedia (3 scenarios); ChineseFoodNet (5 scenarios); VireoFood172 (1 scenario); SVHN + CCPD (5 scenarios); Open Images V7 (2 scenarios).

The objective is for f' to produce an answer consistent with a^p for (i, q^p). CrossCult-KIBench defines this objective under two task settings.

Single-Insert Task. In the single-insert setting, each case contains exactly one insertion sample. Motivated by the diagnostic results in Fig. 1, which show larger gaps for non-English partitions, we use only Chinese and Arabic as single-insert targets. We denote a single-insert case as $c_{\text{single}} = e^p$, where $p \in \{\text{zh}, \text{ar}\}$. Thus, each raw multilingual case yields two single-insert cases.

Sequential-Insert Task. In the sequential-insert setting, each sequential-insert chain contains three insertion samples defined on the same image. We denote a sequential-insert chain as $c_{\text{seq}} = [e_t^{p_t}]_{t=1}^3$, with $e_t^{p_t} = (i, q_t^{p_t}, a_t^{p_t})$ and $(p_1, p_2, p_3) = (\text{en}, \text{zh}, \text{ar})$. Each raw multilingual case yields one sequential-insert chain. The insertion method applies the three insertion samples sequentially to the same model instance. This setting captures the harder condition in which multiple culture-conditioned responses must coexist for one visually grounded scenario within a single insertion sequence.

3.2 Evaluation Tasks and Metrics

Given these two settings, CrossCult-KIBench reports two score types for each benchmark dimension to evaluate insertion success and side effects after insertion. We use ROUGE-L as the metric-based score for textual overlap between a model output and the reference answer, and use gpt-5.4-mini as the LLM-as-Judge scorer for semantic consistency with the reference answer (ranging from 0 to 10). Let $\text{Score}(y, a)$ denote either score type between a generated answer y and a reference answer a.

Single-Insert Evaluation Metrics. For single-insert evaluation, consider a case $c_{\text{single}} = e^p$ and its post-insertion model $f' = \text{Insert}(f, e^p)$. **Reliability** evaluates target insertion by scoring $\text{Score}(f'(i, q^p), a^p)$. **Generality** evaluates same-scenario transfer as $\text{Score}(f'(i_{\text{gen}}, q^p_{\text{gen}}), a^p_{\text{gen}})$, where $e^p_{\text{gen}} = (i_{\text{gen}}, q^p_{\text{gen}}, a^p_{\text{gen}})$ is another insertion sample from the same scenario. **Cross-Language Locality** evaluates if non-target partition behavior on the inserted image is preserved as $\text{Score}(f'(i, q^{p_\ell}_{\text{lang}}), f(i, q^{p_\ell}_{\text{lang}}))$, where $p_\ell \in \mathcal{P} \setminus \{p\}$. **Cross-Scenario Locality** evaluates if behavior outside the inserted scenario is preserved as $\text{Score}(f'(i_{\text{scen}}, q^p_{\text{scen}}), f(i_{\text{scen}}, q^p_{\text{scen}}))$. **Overall** averages Reliability, Generality, Cross-Language Locality, and Cross-Scenario Locality with equal weight.

Sequential-Insert Evaluation Metrics. For sequential-insert evaluation, consider a sequential-insert chain $[e_t^{p_t}]_{t=1}^3$ and let f' be the final model after all three insertions. **Final Reliability** evaluates whether all inserted answers remain available by averaging $\text{Score}(f'(i, q_t^{p_t}), a_t^{p_t})$ over the three steps. **Final Generality** evaluates same-scenario transfer after the full chain by averaging $\text{Score}(f'(i_{\text{gen}}, q^{p_t}_{\text{gen},t}), a^{p_t}_{\text{gen},t})$ over the three steps. **Final Locality** evaluates whether different-scenario behavior remains stable by averaging $\text{Score}(f'(i_{\text{loc},t}, q^{p_t}_{\text{loc},t}), f(i_{\text{loc},t}, q^{p_t}_{\text{loc},t}))$ over the three steps. **Final Overall** averages Final Reliability, Final Generality, and Final Locality with equal weight. Detailed metric formulas are provided in Appendix A.1.

3.3 Benchmark Composition

CrossCult-KIBench contains 49 scenarios across Social Values, Attitudes & Stereotypes, Religious Values, and Ethical Values and Norms. Each scenario describes an image-grounded cultural situation type rather than a single image or a single question. For a given visual instance, the three partition-

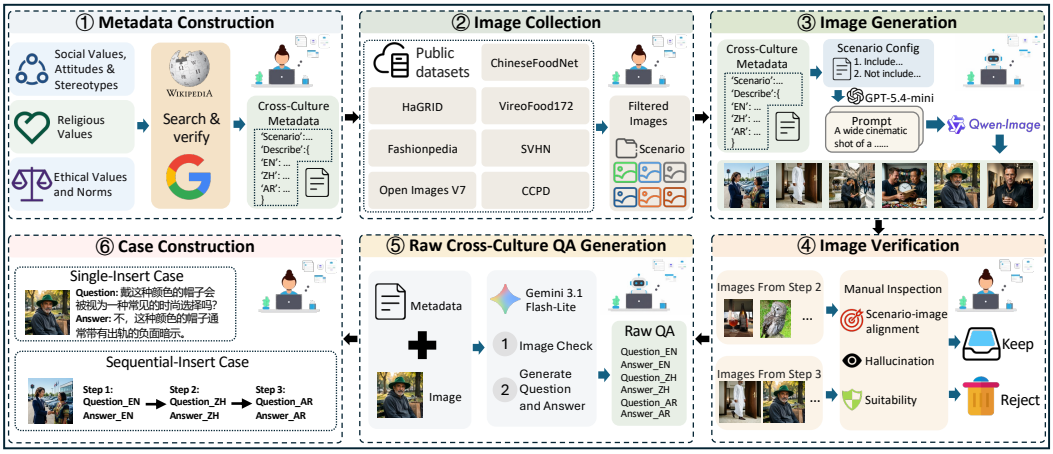


Figure 3: Construction pipeline of CrossCult-KIBench. The benchmark is built through manual design and large model assistance, ending with single-insert and sequential-insert cases.

specific insertion samples form one raw multilingual case. The answers across partitions are not translations, but culture-conditioned judgments for the same visual situation. Each scenario contains 200 raw multilingual cases. Thus, the 49 scenarios yield 9,800 raw multilingual cases before task conversion. This raw structure supports both benchmark tasks. For single-insert, the Chinese and Arabic samples from each raw case produce two single-insert cases. For sequential-insert, the English, Chinese, and Arabic samples from each raw case form one sequential-insert case, implemented as a three-step chain. Table 1 reports the final task-specific case counts. The image set combines generated images with public image datasets, which balances controllability for culturally specific situations with visual realism where public images are available. These data components are assembled through the construction pipeline described next.

3.4 Benchmark Construction Pipeline

Fig. 3 summarizes the construction pipeline of CrossCult-KIBench. The pipeline combines manual benchmark design with LLM assistance. We manually define the cultural scope, verify metadata, and filter candidate images, while large models help expand image prompts, generate controlled images, and draft raw multilingual question-answer pairs.

Metadata Construction. As shown in Step 1 of Fig. 3, we search public cultural sources, including Wikipedia and authoritative websites maintained by cultural institutions and government agencies, to collect candidate scenarios. We select scenarios that reflect representative cultural practices and preferences within the three partitions. We then organize and verify the collected material, yielding 49 scenarios across the three topic groups. For each scenario, we manually write descriptions based on these sources, specifying the visual situation and how it is interpreted or judged in each cultural context. These descriptions form the scenario-level cross-cultural metadata that guides subsequent image preparation and multilingual question-answer construction.

Image Preparation and Verification. We then build the image pool through public image collection and controlled image generation. As shown in Steps 2 and 3, public images are collected from scenario-relevant sources, while gpt-5.4-mini expands the cross-culture metadata and scenario config into generation prompts for Qwen-Image (Wu et al., 2025). The public image sources include HaGRID (Kapitanov et al., 2024), Fashionpedia (Jia et al., 2020), ChineseFoodNet (Chen et al., 2017), VireoFood172 (Chen & Ngo, 2016), SVHN (Netzer et al., 2011), CCPD (Xu et al., 2018), and Open Images V7 (Kuznetsova et al., 2020). Step 4 merges these two image sources and applies human verification for scenario-image alignment, hallucination, and suitability. Only verified images are kept for question-answer construction.

Raw QA and Case Construction. We finally convert verified images and metadata into raw multilingual QA and benchmark-ready cases. As shown in Step 5, Gemini 3.1 Flash-Lite drafts English, Chinese, and Arabic question-answer pairs for each verified image. This stage performs

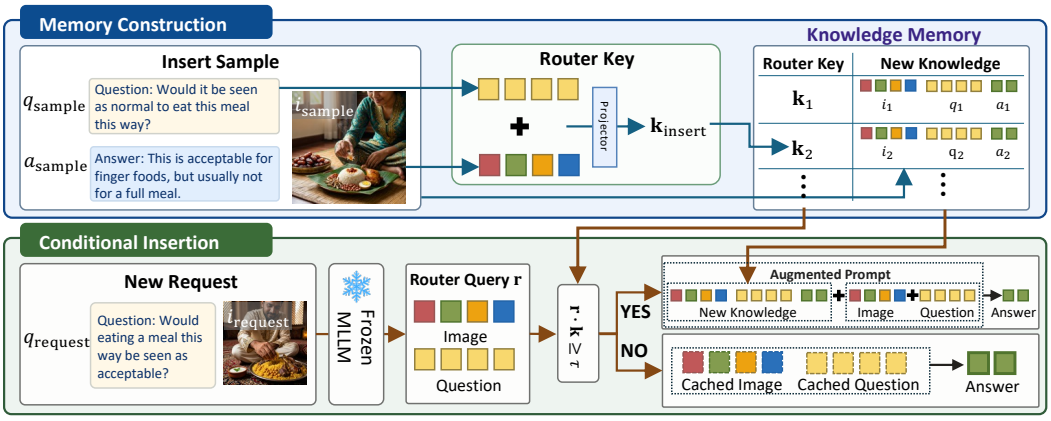


Figure 4: Overview of MCKI. MCKI stores inserted knowledge in memory and uses a router to decide when to use it.

cross-cultural QA authoring rather than mechanical translation, because the three answers encode partition-specific judgments for the same visual instance. We manually review the generated questions and reference answers against the corresponding scenario metadata, retaining only QA pairs that are consistent with it. Step 6 organizes the raw multilingual cases into single-insert cases and sequential-insert chains, following the case definitions introduced earlier. Together, these stages produce the final benchmark inputs for evaluating knowledge insertion methods.

4 Memory-Conditioned Knowledge Insertion

We instantiate Memory-Conditioned Knowledge Insertion (MCKI) as a baseline for CrossCult-KIBench. As shown in Fig. 4, MCKI stores each insertion sample in a knowledge memory and uses representations extracted by the frozen MLLM to decide whether a new request should activate a stored memory entry. Only the lightweight router mapping is trained. Both training and testing involve memory construction from the supplied insertion samples followed by matching queries to memory entries, with router optimization and threshold calibration performed only during training.

4.1 Memory Construction

For an insertion sample $e_j = (i_j, q_j, a_j)$, MCKI writes one entry $m_j = (\mathbf{k}_j, e_j)$ into the knowledge memory. The entry stores a router key $\mathbf{k}_j$ together with the complete insertion sample e_j. To build the router key $\mathbf{k}_j$, MCKI encodes (i_j, q_j) with the frozen MLLM and takes the hidden representations from its final language layer. It mean-pools these representations over the question tokens and visual tokens, then applies a trainable router mapping $\mathbf{k}_j = \phi(\bar{\mathbf{h}}_j^{\text{q}}, \bar{\mathbf{h}}_j^{\text{v}})$.

4.2 Conditional Insertion

For a new request $x = (i_{\text{req}}, q_{\text{req}})$, MCKI performs the same last-layer feature extraction and applies the same router mapping to obtain the request vector $\mathbf{r}(x) = \phi(\bar{\mathbf{h}}_x^{\text{q}}, \bar{\mathbf{h}}_x^{\text{v}})$. The router scores each memory entry by cosine similarity between the request vector and the memory key, and selects the entry with the largest score,

$$\text{sim}(x, m_j) = \frac{\mathbf{r}(x)^\top \mathbf{k}_j}{\|\mathbf{r}(x)\|_2 \|\mathbf{k}_j\|_2}, \quad j^* = \arg\max_j \text{sim}(x, m_j). \tag{1}$$

When the selected similarity falls below the calibrated threshold, MCKI reuses the cached base state and continues decoding with the frozen MLLM. At or above the threshold, MCKI prepends the selected memory entry to the new request and reruns generation with the augmented prompt. Formally, MCKI produces the output y as

$$\tilde{x}_{j^*} = \big[\text{Wrap}(e_{j^*}),\ i_{\text{req}},\ q_{\text{req}}\big], \tag{2}$$

$$y = \begin{cases} f(x), & \text{sim}(x, m_{j^*}) < \tau, \\ f(\tilde{x}_{j^*}), & \text{sim}(x, m_{j^*}) \geq \tau. \end{cases} \tag{3}$$

Table 2: Single-insert knowledge insertion results. ROUGE-L scores are reported as percentages, and LLM Judge scores are reported on a 0 to 10 scale. Higher scores indicate better performance. **Bold** and underlined values indicate the best and second-best results, respectively, excluding Base Model.

Method	Reliability		Generality		Cross-Lang. Loc.		Cross-Scen. Loc.		Overall	
	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge
Base Model: InternVL3.5-8B										
Base Model	12.24	4.02	12.23	4.03	100.00	9.99	100.00	9.98	56.12	7.00
Culture Prompt	12.10	4.35	12.15	4.32	35.01	7.20	33.18	6.99	23.11	5.72
FineTune	**91.12**	**9.35**	**24.66**	**6.32**	5.36	2.28	4.96	0.90	31.52	4.71
IKE	70.25	8.54	23.58	6.28	24.92	5.14	34.46	6.45	38.30	6.60
MEND	19.24	4.82	19.33	4.87	67.16	8.58	42.59	6.89	37.08	6.29
SERAC	57.69	7.64	22.56	5.63	78.81	9.07	26.51	4.38	46.39	6.68
MSCKE	69.93	8.48	23.55	5.80	**100.00**	**9.99**	**91.55**	9.38	71.26	8.41
MCKI (Ours)	79.83	8.86	22.20	5.71	**100.00**	**9.99**	89.66	**9.47**	**72.92**	**8.51**
Base Model: Qwen3.5-9B										
Base Model	14.13	4.50	14.06	4.43	100.00	10.00	100.00	10.00	57.05	7.23
Culture Prompt	15.72	5.11	15.79	5.06	19.74	7.14	31.75	6.80	20.75	6.03
FineTune	55.90	6.35	15.54	4.37	0.00	0.96	3.17	0.61	18.65	3.07
IKE	47.25	8.06	23.76	**6.84**	19.93	5.20	35.72	5.63	31.66	6.43
MEND	44.91	6.52	26.31	5.80	81.83	9.22	61.32	8.03	53.59	7.39
SERAC	67.59	8.23	**27.59**	6.42	5.75	3.98	18.88	3.92	29.95	5.64
MSCKE	**71.27**	**8.50**	27.58	6.27	**100.00**	**10.00**	28.64	4.94	56.87	7.43
MCKI (Ours)	50.66	7.77	24.43	6.74	**100.00**	**10.00**	**88.27**	**9.32**	**65.84**	**8.46**

Table 3: Sequential-insert knowledge insertion results. **Bold** and underlined values indicate the best and second-best results, respectively, excluding Base Model.

Method	Final Rel.		Final Gen.		Final Loc.		Overall	
	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge	ROUGE-L	LLM Judge
Base Model: InternVL3.5-8B								
Base Model	13.55	4.49	13.53	4.48	100.00	9.89	42.36	6.29
Culture Prompt	12.76	4.54	12.73	4.56	37.04	6.94	20.84	5.34
FineTune	6.80	0.90	1.85	0.60	0.26	0.12	2.97	0.54
IKE	75.44	8.49	27.44	6.46	28.87	5.99	43.91	6.98
MEND	2.42	0.74	2.36	0.73	3.19	0.93	2.66	0.80
SERAC	46.62	7.24	23.92	6.01	25.37	4.84	31.97	6.03
MSCKE	66.08	8.80	24.15	**6.54**	56.28	6.96	48.84	7.44
MCKI (Ours)	**85.57**	**8.98**	**28.61**	6.26	**83.44**	**9.18**	**65.87**	**8.14**
Base Model: Qwen3.5-9B								
Base Model	12.10	4.44	12.07	4.42	100.00	10.00	41.39	6.29
Culture Prompt	15.61	5.67	15.54	5.68	26.05	6.73	19.07	6.03
FineTune	2.37	0.44	1.00	0.31	0.28	0.10	1.22	0.29
IKE	45.74	7.45	23.73	6.33	28.67	6.24	32.72	6.67
MEND	4.07	0.77	3.91	0.86	4.17	1.26	4.05	0.96
SERAC	36.23	6.04	19.09	5.03	14.23	4.14	23.18	5.07
MSCKE	**62.10**	**8.54**	28.54	6.60	42.77	6.18	44.47	7.11
MCKI (Ours)	59.83	8.38	**30.07**	**7.18**	**73.68**	**8.89**	**54.53**	**8.15**

Here, τ is the calibrated threshold, $\tilde{x}_{j^*}$ is the augmented request, and $\mathrm{Wrap}(\cdot)$ formats the selected insertion sample e_{j^*} as reference context. The router is trained with a contrastive loss that increases router scores for insertion samples and paired Generality items relative to Locality items, and τ is calibrated by maximizing routing accuracy on positive and negative training scores, as detailed in Appendix B.1.

5 Experiments

5.1 Experimental Setup

Base MLLMs. We evaluate knowledge insertion methods on CrossCult-KIBench with two open-weight MLLMs, InternVL3.5-8B (Wang et al., 2025) and Qwen3.5-9B (Team, 2026). Both models provide strong visual question answering backbones while remaining accessible for controlled insertion experiments.

Baseline Methods. Among existing approaches, knowledge editing methods are most closely related to our task, since they also aim to alter model knowledge or behavior. So we adapt representative editing methods as baselines for comparison with MCKI. Following prior work on multimodal editing (Cheng et al., 2023; Zeng et al., 2025), we include textual editing methods because the set of specialized multimodal editing methods remains limited. **FineTune** updates the MLLM on the insertion sample with target answer supervision. **IKE** (Zheng et al., 2023) uses the insertion sample as an in-context demonstration at inference time. **MEND** (Mitchell et al., 2022a) applies a learned gradient-based edit from the target answer loss. **SERAC** (Mitchell et al., 2022b) retrieves relevant edits and answers with a replacement model when an edit is activated. **MSCKE** (Zeng et al., 2025)

uses a multimodal CLIP router and a replacement model to apply inserted knowledge. **Culture Prompt** is a diagnostic baseline that selects a predefined cultural system prompt based on each request's language-culture partition, e.g., "Please answer from a U.S. cultural perspective in one sentence."

5.2 Main Results

Single-Insert Results. Table 2 shows results on the single-insert task. Culture Prompt offers limited gains over Base Model in Reliability and Generality, suggesting that cultural system prompts alone cannot bridge the relevant knowledge gaps and supporting the need for knowledge insertion. Direct target supervision gives FineTune high Reliability, yet the same parameter updates overwrite existing behavior so aggressively that both Cross-Language Locality and Cross-Scenario Locality suffer severe drops. IKE avoids weight-level damage by prepending the insertion sample as an in-context demonstration at inference time, but its effectiveness depends on whether the base model can use that demonstration, which leaves its Reliability and Overall improvements limited. Among learned editing methods, MEND applies more conservative parameter changes and retains more locality than FineTune, but this conservatism also weakens target fitting. Retrieval-based replacement methods expose a different trade-off. SERAC improves Reliability and Generality when the edit is retrieved, yet similar non-target evaluation requests can also activate the replacement path and erode locality. MSCKE adds multimodal routing and better preserves Cross-Language Locality, but related scenes can still activate replacement behavior and weaken Cross-Scenario Locality. MCKI shows that selective memory activation can yield a more balanced Overall profile by preserving cross-language behavior while maintaining competitive target scores. No existing method balances all four dimensions, showing that CrossCult-KIBench exposes complementary failure modes in cross-cultural knowledge insertion.

Sequential-Insert Results. Table 3 shows that sequential insertion adds challenges that are absent from single insertion. Culture Prompt similarly offers limited gains in Final Reliability and Final Generality over Base Model, reinforcing the need for knowledge insertion to address knowledge gaps that cultural system prompts alone cannot resolve. After several insertions, methods must keep multiple culture-conditioned answers available while preserving unrelated behavior, so failures reflect overwriting, model state drift, and retrieval confusion rather than target fitting alone. FineTune and MEND are most affected because repeated parameter updates can overwrite earlier behavior. IKE avoids weight changes, but its final performance depends on whether the model can use multiple demonstrations without confusing or ignoring earlier insertions. SERAC and MSCKE benefit from explicit memory or routing and therefore maintain stronger target scores than collapsed update methods, yet larger memories and related image-language contexts can still trigger the wrong replacement path and reduce Final Locality. MCKI gives a stronger Overall balance because selective memory activation helps preserve unrelated behavior while retaining inserted answers, but its lower Final Generality shows that storing sequential insertions does not by itself transfer them to related examples from the same scenario. No existing method fully resolves coexistence under repeated insertion, which highlights cross-cultural knowledge insertion as a task with substantial room for further exploration.

5.3 Topic and Partition Sensitivity

To examine whether insertion difficulty depends on cultural topic or language-culture partition, we analyze single-insert Reliability by topic group and inserted partition in Fig. 5. Across topic groups, the two backbones show different patterns. InternVL3.5-8B is relatively stable with a mild drop on Religious cases, whereas Qwen3.5-9B shows weaker gains on Ethical cases. This suggests that topic difficulty reflects how a backbone represents cultural cues, since Religious cases often contain more explicit symbolic triggers while Social and Ethical cases require broader semantic transfer. Partition effects are also model-dependent, with Arabic insertions generally leaving more room for Reliability improvement than Chinese insertions but with a gap that varies across backbones and insertion mechanisms.

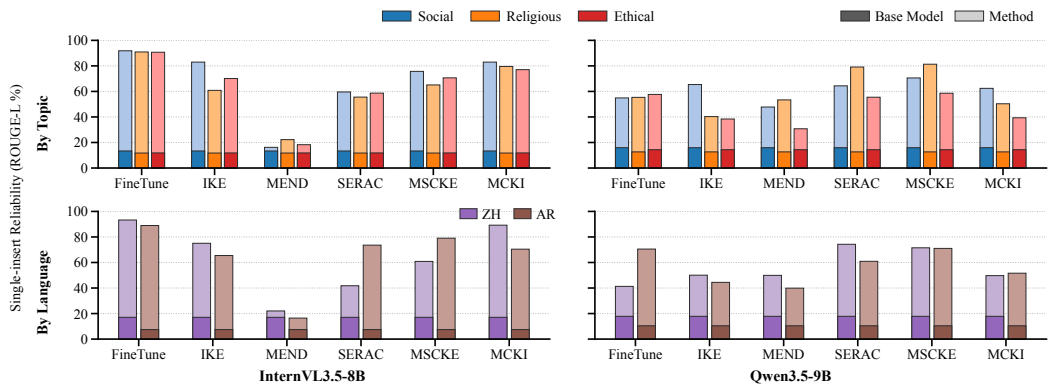


Figure 5: Single-insert reliability by topic group and inserted language-culture partition. Dark bars show the base model and light bars show the edited model.

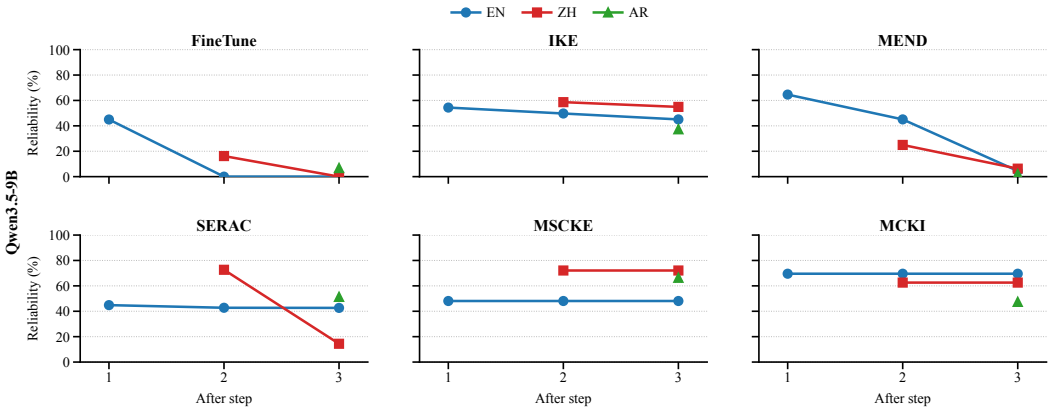


Figure 6: Reliability trajectories after sequential insertion on Qwen3.5-9B.

5.4 Sequential-Insert Retention Analysis

We also evaluated the impact of each insertion on the knowledge inserted before. Fig. 6 shows the results. IKE retains earlier insertions more steadily than methods that update parameters, while FineTune and MEND rapidly lose them after later steps. Retrieval-based and routing-based methods retain inserted answers more consistently, with MSCKE and MCKI giving the strongest retention among the evaluated methods. Additional retention trajectories on InternVL3.5-8B and reordered insertion chains are provided in Appendix D.2.

6 Conclusion

In this work, we introduce the task of cross-cultural knowledge insertion: the post-hoc augmentation of MLLMs with missing culture-specific knowledge, while preserving their original behavior for all other cultural contexts. This formulation enables targeted cultural adaptation without full retraining. To enable rigorous evaluation, we present CrossCult-KIBench, a benchmark built on 49 visually grounded cultural scenarios across English, Chinese, and Arabic language-culture partitions. It supports both single-insert and sequential-insert evaluation, measuring insertion success, same-scenario generalization within the target partition, behavioral preservation, and retention across multiple updates. Experiments reveal that existing knowledge-editing and retrieval-based approaches struggle to balance these objectives, often improving target culture performance at the expense of non-target behavior. As a strong and practical baseline, we propose MCKI, which offers superior performance across multiple metrics. Together, CrossCult-KIBench and the MCKI baseline establish a foundation for developing and evaluating more precise, effective, and culturally-aware knowledge insertion methods for MLLMs.

AI USE STATEMENT

In this work, we used generative AI tools to expand image-generation prompts (gpt-5.4-mini), generate images (Qwen-Image), and draft multilingual question-answer pairs (Gemini 3.1 Flash-Lite), as described in Section 3 and Appendix A.2. We also used gpt-5.4-mini to score the semantic consistency of model answers with reference answers (Appendix C.3). Additionally, AI tools assisted with software implementation and debugging, data processing, and manuscript polishing. Human reviewers screened the images for scenario alignment, visual artifacts, and suitability, and checked generated questions and reference answers against the scenario metadata. The LLM judge was evaluated against human ratings on 1,000 samples (Appendix D.1). Checks of AI-assisted code and data processing included tests and recomputation from saved experiment outputs. The authors are responsible for the final manuscript, results, and artifacts, including material produced with AI assistance.

ETHICS STATEMENT

CrossCult-KIBench is intended for evaluating knowledge insertion under specified cultural settings. Its language-culture partitions are practical evaluation proxies and should not be interpreted as fixed cultural identities or prescriptions for individual users. Culture-conditioned insertion could reinforce stereotypes or be misused to produce manipulative or politically sensitive outputs. We address these risks by specifying the target settings, distinguishing knowledge insertion from cultural inference, and limiting our interpretation of results to the benchmark's stated scope. Images and question-answer pairs are manually screened during construction, as described in Section 3; generated images serve as benchmark inputs rather than representative depictions of real communities. Use of third-party images remains subject to the original datasets' licenses and terms.

REPRODUCIBILITY STATEMENT

Our code and benchmark data are available at `https://github.com/crosscult-kibench/CrossCult-KIBench` and `https://huggingface.co/datasets/crosscult-kibench/CrossCult-KIBench`, respectively. The repositories provide environment specifications, experiment configurations, training and evaluation code, dataset splits, and instructions for reconstructing third-party image subsets that are not bundled with the data. Benchmark construction and metric definitions are detailed in Appendix A, MCKI training and threshold calibration in Appendix B, and model settings, hyperparameters, and evaluation prompts in Appendix C.

REFERENCES

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.

Jingjing Chen and Chong-Wah Ngo. Deep-based ingredient recognition for cooking recipe retrieval. In *Proceedings of the 24th ACM international conference on Multimedia*, pp. 32–41, 2016.

Xin Chen, Yu Zhu, Hua Zhou, Liang Diao, and Dongyan Wang. Chinesefoodnet: A large-scale image dataset for chinese food recognition. *arXiv preprint arXiv:1705.02743*, 2017.

Siyuan Cheng, Bozhong Tian, Qingbin Liu, Xi Chen, Yongheng Wang, Huajun Chen, and Ningyu Zhang. Can we edit multimodal large language models? In *EMNLP*, 2023.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

Han Huang, Haitian Zhong, Tao Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. Vlkeb: A large vision-language model knowledge editing benchmark. *Advances in Neural Information Processing Systems*, 37:9257–9280, 2024.

Jiaxing Huang and Jingyi Zhang. A survey on evaluation of multimodal large language models. *arXiv preprint arXiv:2408.15769*, 2024.

Menglin Jia, Mengyun Shi, Mikhail Sirotenko, Yin Cui, Claire Cardie, Bharath Hariharan, Hartwig Adam, and Serge Belongie. Fashionpedia: Ontology, segmentation, and an attribute localization dataset. In *European Conference on Computer Vision (ECCV)*, 2020.

Alexander Kapitanov, Karina Kvanchiani, Alexander Nagaev, Roman Kraynov, and Andrei Makhliarchuk. Hagrid – hand gesture recognition image dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 4572–4581, January 2024.

Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International journal of computer vision*, 128(7):1956–1981, 2020.

Cheng Li, Mengzhuo Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. Culturellm: Incorporating cultural differences into large language models. *Advances in Neural Information Processing Systems*, 37:84799–84838, 2024a.

Cheng Li, Damien Teney, Linyi Yang, Qingsong Wen, Xing Xie, and Jindong Wang. Culturepark: Boosting cross-cultural understanding in large language models. *Advances in Neural Information Processing Systems*, 37:65183–65216, 2024b.

Jiaqi Li, Miaozeng Du, Chuanyi Zhang, Yongrui Chen, Nan Hu, Guilin Qi, Haiyun Jiang, Siyuan Cheng, and Bozhong Tian. Mike: A new benchmark for fine-grained multimodal entity knowledge editing. *arXiv preprint arXiv:2402.14835*, 2024c.

Peiqin Lin, Shaoxiong Ji, Jörg Tiedemann, André FT Martins, and Hinrich Schütze. Mala-500: Massive language adaptation of large language models. *arXiv preprint arXiv:2401.13303*, 2024.

Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. Visually grounded reasoning across languages and cultures. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10467–10485, 2021.

Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model editing at scale. In *ICLR*, 2022a.

Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. Memory-based model editing at scale. In *ICML*, 2022b.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. URL `http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf`.

Xuan-Phi Nguyen, Wenxuan Zhang, Xin Li, Mahani Aljunied, Zhiqiang Hu, Chenhui Shen, Yew Ken Chia, Xingxuan Li, Jianyu Wang, Qingyu Tan, et al. Seallms-large language models for southeast asia. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pp. 294–304, 2024.

David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, et al. Cvqa: Culturally-diverse multilingual visual question answering benchmark. *arXiv preprint arXiv:2406.05967*, 2024.

Qwen Team. Qwen3.5: Accelerating productivity with native multimodal agents, February 2026. URL `https://qwen.ai/blog?id=qwen3.5`.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Sibo Wang, Xiangkui Cao, Jie Zhang, Zheng Yuan, Shiguang Shan, Xilin Chen, and Wen Gao. Vlbiasbench: A comprehensive benchmark for evaluating bias in large vision-language model. *arXiv preprint arXiv:2406.14194*, 2024.

Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025.

World Values Survey. World values survey. `https://www.worldvaluessurvey.org/wvs.jsp`, 2022.

Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025.

Zhenbo Xu, Wei Yang, Ajin Meng, Nanxue Lu, Huan Huang, Changchun Ying, and Liusheng Huang. Towards end-to-end license plate detection and recognition: A large dataset and baseline. In *Computer Vision – ECCV 2018*, pp. 261–277, 2018.

Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities. In *EMNLP*, 2023.

Zhen Zeng, Leijiang Gu, Xun Yang, Zhangling Duan, Zenglin Shi, and Meng Wang. Visual-oriented fine-grained knowledge editing for multimodal large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2491–2500, October 2025.

Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. Mm-llms: Recent advances in multimodal large language models. *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 12401–12430, 2024a.

Junzhe Zhang, Huixuan Zhang, Xunjian Yin, Baizhou Huang, Xu Zhang, Xinyu Hu, and Xiaojun Wan. Mc-mke: A fine-grained multimodal knowledge editing benchmark emphasizing modality consistency. *arXiv preprint arXiv:2406.13219*, 2024b.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.

Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. Can we edit factual knowledge by in-context learning? In *EMNLP*, 2023.

A BENCHMARK DETAILS

A.1 METRIC DETAILS

A.1.1 SINGLE-INSERT METRICS

Let $\mathcal{C}_{\text{single}}$ denote the set of single-insert cases. For each $c \in \mathcal{C}_{\text{single}}$, we write the insertion sample as $e_c = (i_c, q_c^{p_c}, a_c^{p_c})$ and denote the post-insertion model by f'_c. Reliability is

$$\mathcal{M}_{\text{reliability}} = \frac{1}{|\mathcal{C}_{\text{single}}|} \sum_{c \in \mathcal{C}_{\text{single}}} \text{Score}(f'_c(i_c, q_c^{p_c}), a_c^{p_c}) \,. \tag{4}$$

Let $e_{c,\text{gen}}^{p_c} = (i_{c,\text{gen}}, q_{c,\text{gen}}^{p_c}, a_{c,\text{gen}}^{p_c})$ denote the Generality evaluation item. Generality is

$$\mathcal{M}_{\text{generality}} = \frac{1}{|\mathcal{C}_{\text{single}}|} \sum_{c \in \mathcal{C}_{\text{single}}} \text{Score}\big(f'_c(i_{c,\text{gen}}, q_{c,\text{gen}}^{p_c}), a_{c,\text{gen}}^{p_c}\big) \,. \tag{5}$$

Let $(i_c, q_{c,\text{lang}}^{p_\ell})$ denote the Cross-Language Locality evaluation item, with $p_\ell \neq p_c$. Cross-Language Locality is

$$\mathcal{M}_{\text{cross-language}} = \frac{1}{|\mathcal{C}_{\text{single}}|} \sum_{c \in \mathcal{C}_{\text{single}}} \text{Score}\Big(f'_c(i_c, q_{c,\text{lang}}^{p_\ell}), f(i_c, q_{c,\text{lang}}^{p_\ell})\Big) \,. \tag{6}$$

Let $(i_{c,\text{scen}}, q_{c,\text{scen}}^{p_c})$ denote the Cross-Scenario Locality evaluation item. Cross-Scenario Locality is

$$\mathcal{M}_{\text{cross-scenario}} = \frac{1}{|\mathcal{C}_{\text{single}}|} \sum_{c \in \mathcal{C}_{\text{single}}} \text{Score}\big(f'_c(i_{c,\text{scen}}, q_{c,\text{scen}}^{p_c}), f(i_{c,\text{scen}}, q_{c,\text{scen}}^{p_c})\big) \,. \tag{7}$$

Single-insert Overall is

$$\mathcal{M}_{\text{overall}}^{\text{single}} = \frac{1}{4} \left(\mathcal{M}_{\text{reliability}} + \mathcal{M}_{\text{generality}} + \mathcal{M}_{\text{cross-language}} + \mathcal{M}_{\text{cross-scenario}}\right) . \tag{8}$$

A.1.2 SEQUENTIAL-INSERT METRICS

Let $\mathcal{C}_{\text{seq}}$ denote the set of sequential-insert chains. For each $c \in \mathcal{C}_{\text{seq}}$, let f'_c be the final post-insertion model after all three insertions. Let $e_{c,t}^{p_{c,t}} = (i_c, q_{c,t}^{p_{c,t}}, a_{c,t}^{p_{c,t}})$ denote the insertion sample for step t. Final Reliability is

$$\mathcal{M}_{\text{reliability}}^{\text{final}} = \frac{1}{3|\mathcal{C}_{\text{seq}}|} \sum_{c \in \mathcal{C}_{\text{seq}}} \sum_{t=1}^{3} \text{Score}\big(f'_c(i_c, q_{c,t}^{p_{c,t}}), a_{c,t}^{p_{c,t}}\big) \,. \tag{9}$$

Let $e_{c,\text{gen},t}^{p_{c,t}} = (i_{c,\text{gen}}, q_{c,\text{gen},t}^{p_{c,t}}, a_{c,\text{gen},t}^{p_{c,t}})$ denote the Generality evaluation item for step t. Final Generality is

$$\mathcal{M}_{\text{generality}}^{\text{final}} = \frac{1}{3|\mathcal{C}_{\text{seq}}|} \sum_{c \in \mathcal{C}_{\text{seq}}} \sum_{t=1}^{3} \text{Score}\big(f'_c(i_{c,\text{gen}}, q_{c,\text{gen},t}^{p_{c,t}}), a_{c,\text{gen},t}^{p_{c,t}}\big) \,. \tag{10}$$

Let $(i_{c,\text{loc},t}, q_{c,\text{loc},t}^{p_{c,t}})$ denote the Locality evaluation item paired with step t. Final Locality is

$$\mathcal{M}_{\text{locality}}^{\text{final}} = \frac{1}{3|\mathcal{C}_{\text{seq}}|} \sum_{c \in \mathcal{C}_{\text{seq}}} \sum_{t=1}^{3} \text{Score}\Big(f'_c(i_{c,\text{loc},t}, q_{c,\text{loc},t}^{p_{c,t}}), f(i_{c,\text{loc},t}, q_{c,\text{loc},t}^{p_{c,t}})\Big) \,. \tag{11}$$

Sequential-insert Overall is

$$\mathcal{M}_{\text{overall}}^{\text{seq}} = \frac{1}{3} \left(\mathcal{M}_{\text{reliability}}^{\text{final}} + \mathcal{M}_{\text{generality}}^{\text{final}} + \mathcal{M}_{\text{locality}}^{\text{final}}\right) . \tag{12}$$

Image Prompt Generation

System:
You are an AI assistant that generates highly detailed, photorealistic image prompts for a text-to-image model. Your goal is to describe a scene where an adult man and an adult woman are shaking hands. The description must be precise, focusing on visual clarity of gender and action.

User:
Generate a detailed photorealistic image prompt for the following scene: "medical consultation room".
The main subjects are an ADULT MAN (South Asian with warm skin tone, in his 20s, dreadlocks tied back, wearing a grey hoodie and cargo pants) and an ADULT WOMAN (East Asian with porcelain skin, in her 40s, wavy blonde hair, wearing an elegant floral midi dress) shaking hands (opposite-sex handshake).

REQUIREMENTS:
1. **Character Consistency**: Strictly follow the specific physical descriptions (ethnicity, age, hair) provided above for the man and woman. For the clothes, if the provided attire is unreasonable for the scene, change it to reasonable attire.
2. **Action**: They are clearly shaking hands with their RIGHT HANDS. The handshake should be prominent and firm.
3. **Framing**: Medium shot or Eye-level shot. BOTH FACES must be CLEARLY VISIBLE (not cropped, not turned away, no masks, no sunglasses).
4. **Irrelevant Details**: Incorporate the object "opened notebook" naturally into the scene. Additionally, randomly add a few other objects that are suitable for this specific scene. If the provided object is unreasonable in this scene, do not include it.
5. **Style**: Photorealistic, cinematic lighting, 8k, highly detailed.

Output ONLY the raw prompt text suitable for Stable Diffusion or similar models. Do not include introductory text like 'Here is the prompt' and do not wrap it in quotes.

Generated Prompt	Generated Image
Photorealistic cinematic 8k medium shot, eye-level perspective in a modern medical consultation room. An adult South Asian man in his 20s with a warm skin tone and long dreadlocks neatly tied back, wearing a casual grey hoodie and olive cargo pants, is standing and firmly shaking hands with an adult East Asian woman in her 40s. The woman has porcelain skin and shoulder-length wavy blonde hair, dressed in a professional white medical lab coat over an elegant floral midi dress. They are using their right hands for a prominent, firm handshake. Both of their faces are clearly visible, showing friendly and professional expressions. In the background, a polished wooden desk features an opened notebook with handwritten notes, a silver stethoscope, a digital blood pressure monitor, and a small succulent plant. Large windows allow soft, natural daylight to fill the room, complemented by warm indoor cinematic lighting that highlights textures and skin details. Sharp focus, high contrast, hyper-detailed textures, professional photography.	

Figure 7: Example of the LLM-assisted image prompt construction step.

A.2 Benchmark Construction Details

CrossCult-KIBench is built through a human-led pipeline with constrained model assistance. We first define the three topic groups, identify and verify public cultural sources, and write scenario metadata that specifies the visual situation, the partition-specific cultural context, the expected judgment boundary, question-style examples, question design notes, and image-check notes. The metadata then controls image preparation, where candidate images are collected from public image sources when suitable candidates are available and generated otherwise. For generated-image scenarios, gpt-5.4-mini expands the scenario configuration into a detailed image-generation prompt, as illustrated in Fig. 7 for a handshake scenario involving a man and a woman. Qwen-Image (Wu et al., 2025) uses the resulting positive prompt together with the scenario negative prompt and suffix to generate candidate images. We then manually screen all collected and generated candidates for scenario-image alignment, visual artifacts, and suitability. After the verified image set is finalized, Gemini-3.1-Flash-Lite drafts raw multilingual QA items from each image and its scenario metadata, as shown in Fig. 8. This prompt first instructs the model to check whether the image matches the scenario and whether it contains obvious generation errors. Only when the image passes both checks does the model produce one QA item with aligned English, Chinese, and Arabic question versions and partition-specific answers. We then manually review the generated questions and reference answers for consistency with the corresponding scenario metadata. QA pairs that are inconsistent with the metadata are excluded. For single-insert evaluation, the Chinese and Arabic partitions become target insertion samples. For sequential-insert evaluation, the English, Chinese, and Arabic samples from the same raw multilingual case form a three-step insertion chain.

Raw Question-Answer Pair Generation Prompt

System:
You validate image-scenario alignment and generate compact multilingual QA for cross-cultural image scenarios.
Always return valid JSON only, no markdown fences, no extra text.

User:
You are given one image and one scenario description. First, judge whether the image matches the scenario and whether it has obvious hallucination/artifact/error. Then, only if the image passes both checks, create exactly ONE concise VQA-style interpretation QA. The question must keep the same meaning across English/Chinese/Arabic
versions. All questions and answers must be one short sentence. Return strict JSON with this schema and field names:

```
{
 "is_match": true,
 "has_obvious_error": false,
 "reason": "one short sentence",
 "qa": {
  "question_en": "...",
  "question_zh": "...",
  "question_ar": "...",
  "answer_us_en": "...",
  "answer_china_zh": "...",
  "answer_china_zh_en": "...",
  "answer_arab_ar": "...",
  "answer_arab_ar_en": "..."
 }
}
```

Rules:
- is_match=true only when the core content clearly fits the scenario.
- has_obvious_error=true for obvious visual hallucination, impossible structure, severe artifact, or clear generation error.
- reason must be concise and in English.
- Use one single question only; do not output multiple questions.
- The question should be a good VQA question: specific to this image, but not a literal visual description.
- Write the question like a normal person would ask in conversation, not like an academic or survey prompt.
- Vary the phrasing naturally.
- Prefer concrete conversational wording over abstract nouns unless there is no cleaner wording.
- Do not ask what type, what kind, what style, what meaning, what significance, what it represents, or what it communicates.
- Do not directly restate the scenario name unless there is no natural alternative.
- The question itself must stay inside one local point of view.
- Do not mention different cultures, cultural observers, or cross-cultural comparison in the question.

Scenario: {scenario_name}
Description: {describe}
Image file: {image_file_name}

Figure 8: Example of the LLM-assisted raw multilingual QA construction step.

B MCKI Details

B.1 Training of MCKI

MCKI trains a lightweight router while keeping the base MLLM frozen. The router maps frozen hidden states to the memory key and request vector. These modules include the text and visual normalization layers, the text and visual projection layers, and the fusion projection that produces the final normalized router vector. The stored insertion sample is not optimized. It serves as the reference payload when a memory entry is activated.

For memory entry m_j, let $\mathcal{P}_j$ contain the insertion sample e_j and its paired Generality evaluation items. Let $\mathcal{N}_j$ contain the paired Locality evaluation items. The positive set trains the router to activate the same memory entry for the inserted case and its intended Generality variants. The negative set trains the router to abstain when the base behavior should be preserved. With logit scale γ, define the positive and negative logits as

$$\ell_p^+ = \gamma \,\mathrm{sim}(p, m_j), \tag{13}$$

$$\ell_n^- = \gamma \,\mathrm{sim}(n, m_j) + \log w_n. \tag{14}$$

The router minimizes

$$\mathcal{L}_j = \frac{1}{|\mathcal{P}_j|} \sum_{p \in \mathcal{P}_j} \left[-\ell_p^+ + \lambda \log \left(\exp(\ell_p^+) + \sum_{n \in \mathcal{N}_j} \exp(\ell_n^-) \right) \right]. \tag{15}$$

Here, λ denotes the negative loss weight and w_n denotes the negative item weight. We assign a larger w_n to Cross-Scenario Locality items so that unrelated scenarios more strongly penalize accidental activation.

B.2 Threshold Calibration

After training, MCKI calibrates τ on the observed positive and negative router scores. The threshold turns the continuous router similarity into the binary activate or abstain decision used at inference time. This calibration is needed because the router loss learns relative score separation, while inference requires a fixed operating point. If τ is too low, unrelated Locality requests may activate a memory entry and disturb preserved behavior. If τ is too high, valid insertion and Generality requests may be rejected. We therefore choose τ to maximize the activation accuracy on training positives and negatives. Let $\mathcal{S}^+$ and $\mathcal{S}^-$ denote these score sets. The calibrated threshold is

$$\tau = \arg \max_{t \in [-1,1]} \left(\sum_{r \in \mathcal{S}^+} \mathbb{I}[r \geq t] + \sum_{r \in \mathcal{S}^-} \mathbb{I}[r < t] \right). \tag{16}$$

C Experiment Details

C.1 Base Models

InternVL3.5-8B InternVL3.5-8B is a multimodal InternVL model for visual perception, image-grounded reasoning, and general vision-language instruction following. The experiment checkpoint is `OpenGVLab/InternVL3_5-8B-HF`. The model is available at `https://huggingface.co/OpenGVLab/InternVL3_5-8B-HF`.

Qwen3.5-9B Qwen3.5-9B is a multimodal Qwen model that pairs a 9B language model with a vision encoder for visual understanding, text-rich image reasoning, and instruction following. The experiment checkpoint is `Qwen/Qwen3.5-9B`. The model is available at `https://huggingface.co/Qwen/Qwen3.5-9B`.

C.2 Experiment Parameters and System Prompts

All knowledge insertion runs use batch size 1, a maximum generation length of 64 tokens, and the language-specific system prompts in Fig. 9. Table 4 summarizes the method-specific update targets, auxiliary models, learning rates, and routing or calibration settings.

System Prompt	
English	Please answer briefly in one sentence.
Chinese	请用一句话简单回答。
Arabic	يرجى الإجابة بإيجاز في جملة واحدة.

Figure 9: Language-specific system prompts used for benchmark answer generation.

C.3 LLM Evaluation Prompt

We use `gpt-5.4-mini` as the LLM-as-Judge scorer. For Reliability and Generality, the reference answer is the benchmark reference. For Locality, the reference answer is the corresponding base-model output before insertion. The judge receives the language, question, reference answer, and candidate answer, without the image. It returns a JSON object containing an integer score from 0 to 10 and a short reason. We disable reasoning for the judge and set the output limit to 256 tokens. Fig. 10 gives the complete judge prompt.

Table 4: Experiment parameters.

FineTune			
Base model	Steps	Edited parameters	Edit LR
InternVL3.5-8B	20	layer 35 attention and MLP weights	1×10^{-3}
Qwen3.5-9B	20	layer 31 attention and MLP weights	1×10^{-3}

IKE				
Scope	Retrieved demos	Text encoder	Max text length	Embeddings
both base models	16	`all-MiniLM-L6-v2`	512	normalized

MEND					
Base model	Epochs	Edited parameters	Meta LR	Insert LR	Rank
InternVL3.5-8B	1	layer 33, 34, and 35 MLP up and down projections	1×10^{-6}	1×10^{-4}	1920
Qwen3.5-9B	1	layer 29, 30, and 31 MLP up and down projections	1×10^{-6}	1×10^{-4}	1920

SERAC							
Base model	Epochs	Replacement model	Layer	Text encoder	Max length	Threshold	Classifier / replacement LR
InternVL3.5-8B	1	`InternVL3_5-1B-HF`	27	`all-MiniLM-L6-v2`	512	0.5	1×10^{-4} / 1×10^{-5}
Qwen3.5-9B	1	`Qwen3.5-0.8B`	23	`all-MiniLM-L6-v2`	512	0.5	1×10^{-4} / 1×10^{-5}

MSCKE						
Base model	Epochs	Replacement model	Layer	Router	Threshold	Classifier / replacement LR
InternVL3.5-8B	1	`InternVL3_5-1B-HF`	27	CLIP ViT-L/14	0.5	1×10^{-4} / 1×10^{-5}
Qwen3.5-9B	1	`Qwen3.5-0.8B`	23	CLIP ViT-L/14	0.5	1×10^{-4} / 1×10^{-5}

MCKI						
Base model	Epochs	Router LR	Route dim	Logit scale	Negative weights	Threshold
InternVL3.5-8B	1	1×10^{-3}	1024	20	1.0 / 1.5	calibrated
Qwen3.5-9B	1	1×10^{-3}	1024	20	1.0 / 1.5	calibrated

D Additional Experiments

D.1 Human Validation of the LLM Judge

To further validate the reliability of the LLM judge used in our benchmark, we conducted a human validation study with five human experts on 1,000 samples, balanced across methods, metrics, and base models. Each expert independently evaluated all samples using the same 0-10 rubric as the judge (Appendix C.3). Table 5 reports the average human evaluation scores. The human scores show only small discrepancies from the LLM scores and a high Pearson correlation of 0.916, further supporting the reliability of our LLM-based evaluation protocol.

Table 5: Human validation of the LLM judge on 1,000 evaluation records. Scores are on a 0-10 scale. LLM and Human columns give group means; $\Delta = \mathrm{Human} - \mathrm{LLM}$. Reliability and Generality pool the single-insert and sequential-insert records.

Metric	N	LLM	Human	Δ
Reliability	200	6.045	6.375	$+0.330$
Generality	200	4.210	4.515	$+0.305$
Cross-Language Locality	200	7.580	7.773	$+0.193$
Cross-Scenario Locality	200	6.280	6.237	-0.043
Sequential Final Locality	200	5.505	4.982	-0.523
All sampled records	1,000	5.924	5.976	$+0.052$
Pearson correlation: 0.916; Spearman correlation: 0.920.				

D.2 Additional Sequential-Insert Retention Analysis

The main retention analysis in Fig. 6 uses the original insertion order (EN, ZH, AR) on Qwen3.5-9B. Fig. 11 repeats the same original order on InternVL3.5-8B. Fig. 12 and Fig. 13 use the supplemental

LLM-as-Judge Prompt

System:
You are a strict semantic similarity judge for cultural QA answers.

Your task is to score how well the candidate answer matches the reference answer in meaning, using an integer from 0 to 10.

Judge only semantic agreement in cultural meaning.
Do not use any image information.
Do not reward verbosity, style, or fluency.
Extra detail is acceptable only if it does not change the core cultural judgment.

Scoring rubric:
- 10: Semantically equivalent to the reference answer. Same core cultural judgment and stance.
- 8-9: Strong match. Same core judgment, with only minor missing detail or harmless extra detail.
- 6-7: Partial match. Main direction is similar, but an important nuance is missing, weakened, or made vague.
- 4-5: Limited match. Some overlap in topic, but the core cultural judgment is not well aligned.
- 2-3: Mostly mismatched. The answer is largely inconsistent with the reference, though it retains slight relevance.
- 1: Almost completely mismatched.
- 0: Opposite meaning, irrelevant answer, empty answer, or refusal.

Additional rules:
- If the candidate answer contradicts the reference answer, score at most 2.
- If the candidate answer only describes the scene but does not answer the question, score at most 4.
- If the candidate answer weakens a clear cultural stance into a vague answer, usually score 6-8 rather than 9-10.
- If extra detail changes the original stance, score 5 or below.
- Return valid JSON only, with keys: score, reason.
- score must be an integer between 0 and 10.
- reason must be one short sentence.

User:
Language: {language}
Question: {question}
Reference answer: {reference_answer}
Candidate answer: {candidate_answer}

Evaluate the semantic similarity between the candidate answer and the reference answer for this question.

Return JSON only:
{"score": <0-10 integer>, "reason": "<short reason>"}

Figure 10: LLM-as-Judge prompt used for semantic scoring.

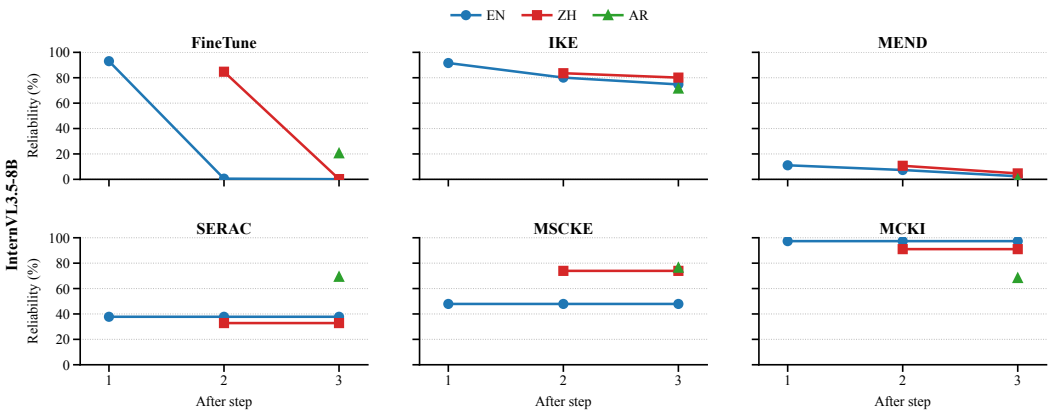


Figure 11: Reliability trajectories after sequential insertion in the original order (EN, ZH, AR) on InternVL3.5-8B.

reordered insertion order (AR, ZH, EN) on Qwen3.5-9B and InternVL3.5-8B, respectively. Together, these supplementary trajectories reinforce the main conclusion that later insertions can overwrite earlier inserted answers. FineTune often attains high Reliability immediately after the current insertion, but earlier curves collapse after subsequent steps on both base models. MEND shows the same interference pattern, with lower immediate Reliability and little retention after later insertions. IKE is more stable than FineTune and MEND in the reordered chains, although its earlier curves still decline as more partitions are inserted. SERAC reduces abrupt collapse relative to parameter-update

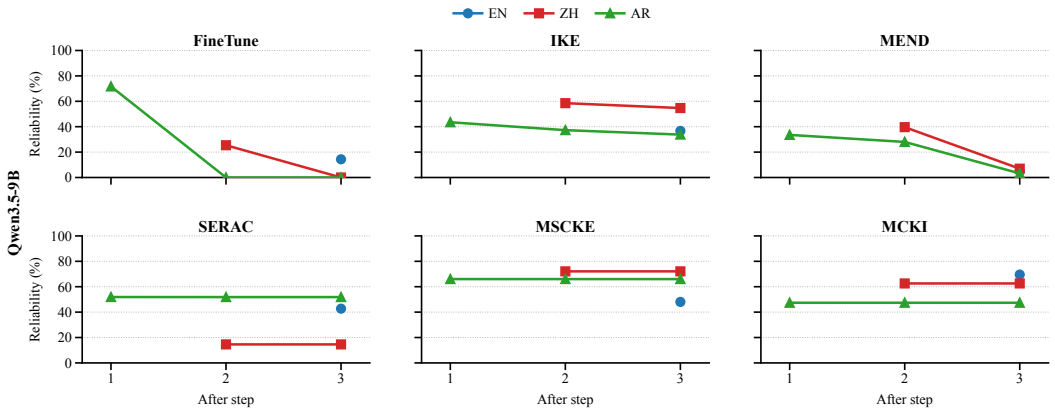


Figure 12: Reliability trajectories after reordered sequential insertion (AR, ZH, EN) on Qwen3.5-9B.

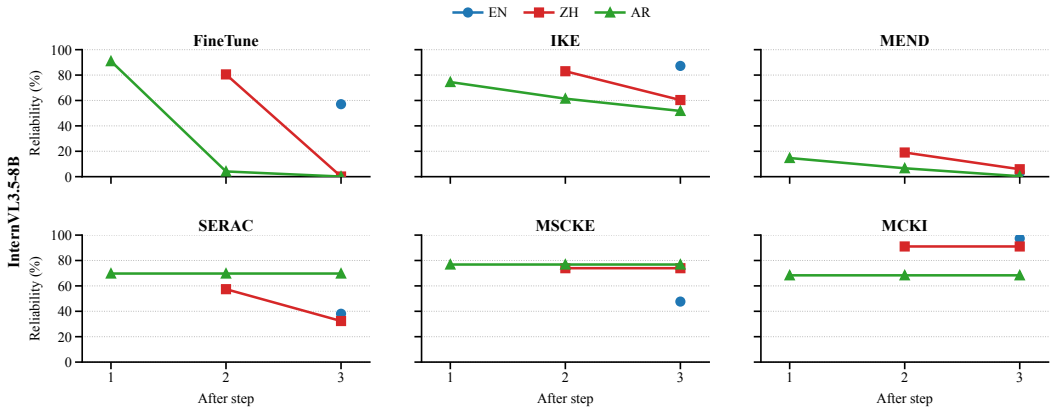


Figure 13: Reliability trajectories after reordered sequential insertion (AR, ZH, EN) on InternVL3.5-8B.

methods, although its absolute Reliability remains partition-dependent. MSCKE and MCKI provide the most consistent retention patterns across the supplementary settings, which indicates that explicit retrieval or routing helps keep earlier inserted answers available after later insertions. The reordered chains further show that this behavior is not only an artifact of placing Arabic last in the original order, since Arabic still degrades for FineTune and MEND when it is inserted first, while MSCKE and MCKI keep the first-step curve stable.

D.3 Efficiency Analysis

To analyze insertion efficiency across methods, we measure runtime cost on one NVIDIA A100-SXM4 80GB GPU using 100 training cases and 100 single-insert evaluation cases. Table 6 reports measured latency, peak GPU memory, and loaded parameter count. Request time includes the full routing or replacement path and autoregressive generation. FineTune requires no training, but it incurs highest insertion latency and the largest evaluation memory. MEND is also expensive because each edit requires gradient-based parameter updates and high memory during both training and evaluation. IKE and SERAC have very low insertion latency, but their request latency depends on in-context prompting or replacement generation, and SERAC loads an additional replacement model. MSCKE has moderate insertion cost and low request latency, but its replacement model increases the total loaded parameter count. MCKI keeps the loaded parameter count close to the base MLLM and uses substantially less evaluation memory than FineTune and MEND, while adding moderate router training and insertion cost.

Table 6: Efficiency analysis on 100 sampled training cases and 100 sampled single-insert evaluation cases. Parameters count the total loaded parameters.

Backbone	Method	Train / case ms	Insert / case ms	Request / sample ms	Train Mem. GiB	Eval Mem. GiB	Params B
InternVL3.5-8B	FineTune	-	13954.3	2002.0	-	73.6	8.53
	IKE	-	4.5	1581.4	-	24.3	8.55
	MEND	5251.5	867.7	1822.1	44.1	71.3	8.53
	SERAC	2017.7	4.6	1325.1	54.9	27.7	9.61
	MSCKE	361.9	48.3	1333.4	28.0	29.2	10.02
	MCKI	1655.3	306.0	1902.0	17.7	21.1	8.54
Qwen3.5-9B	FineTune	-	5700.6	1981.1	-	41.9	9.41
	IKE	-	4.5	1231.6	-	22.1	9.43
	MEND	2898.8	318.3	1430.0	25.7	39.3	9.41
	SERAC	1211.5	4.7	735.6	22.8	20.1	10.29
	MSCKE	308.9	45.8	862.2	22.3	21.6	10.69
	MCKI	1026.6	107.8	943.9	17.8	18.0	9.42

D.4 Case Study

To better illustrate differences among knowledge insertion methods, we conduct case studies that show their insertion behavior and the consequences that follow. The results are shown in Fig. 14 and Fig. 15.



Question:
像这样站在这种设备的中间，会被其他人认为很困扰吗？
Would it be considered bothersome to stand in the middle of such a device like this?

Reference Answer:
通常没问题，因为现在更强调站稳以确保安全。
No problems usually, as there is now more emphasis on staying grounded for safety.

Method	Answer	Score
Base Model	不会，这很常见，别人通常不会觉得困扰。 No, this is common, and others generally would not find it bothersome.	9
FineTune	更更更



Question:
هل يعتبر معظم الناس هذه الإشارة اليدوية تصرفاً ودوداً؟
Do most people consider this hand gesture a friendly act?

Reference Answer:
لا، غالباً ما تُعتبر هذه الإشارة غريبة وغير مألوفة وقد يُساء فهمها.
No, this sign is often considered strange and unfamiliar and may be misunderstood.

Method	Answer	Score
Base Model	نعم، يُعتبر معظم الناس الإشارة اليدوية (الصليب أو إشارة peace) تصرفاً ودوداً و symbol للسلام. Yes, most people consider the hand sign (the cross or peace sign) a friendly gesture and a symbol of peace.	1
FineTune	لا، غالباً ما تُعتبر هذه الإشارة غريبة وغير مألوفة وقد يُساء فهمها. No, this sign is often considered strange and unfamiliar and may be misunderstood.	10
IKE	لا، غالباً ما تُعتبر هذه الإشارة غريبة وغير مألوفة وقد يُساء فهمها. No, this sign is often considered strange and unfamiliar and may be misunderstood.	10
MEND	نعم، يُعتبر معظم الناس هذه الإشارة اليدوية تصرفاً ودوداً و symbolizes peace. Yes, most people consider this hand gesture a friendly act and symbolizes peace.	0
SERAC	لا، تُعتبر هذه الإشارة غريبة غير مألوفة. No, this sign is strange and unfamiliar.	9
MSCKE	لا، غالباً ما تُعتبر هذه الإشارة غريبة غير مألوفة. No, this sign is often considered strange and unfamiliar.	9
MCKI	لا،غالباً ما تُعتبر هذه الإشارة غريبة وغير مألوفة وقد يُساء فهمها. No, this sign is often considered strange and unfamiliar and may be misunderstood.	10



Question:
走在这里时，这种设置会让你感到更安全还是被监视？
Walking here, does this setting make you feel safer or more watched?

Reference Answer:
大多数人认为这是维护公共安全和秩序的必要手段，通常会感到安心。
Most people consider this a necessary means to maintain public safety and order and usually feel at ease.

Method	Answer	Score
Base Model	被监视。 Under surveillance.	1
FineTune	الأمن الأمن الأمن