

# On the Robustness of Knowledge Editing for Detoxification

Ming Dong, Shiyi Tang, Ziyang Peng, Guanyi Chen\*, and Tingting He

Hubei Provincial Key Laboratory of Artificial Intelligence and Smart Learning,  
National Language Resources Monitoring and Research Center for Network Media,  
School of Computer Science, Central China Normal University  
mingdong@whut.edu.cn, g.chen@ccnu.edu.cn

## Abstract

Knowledge-Editing-based (KE-based) detoxification has emerged as a promising approach for mitigating harmful behaviours in Large Language Models. Existing evaluations, however, largely rely on automatic toxicity classifiers, implicitly assuming that reduced toxicity scores reflect genuine behavioural suppression. In this work, we propose a robustness-oriented evaluation framework for KE-based detoxification that examines its reliability beyond standard classifier-based metrics along three dimensions: optimisation robustness, compositional robustness, and cross-lingual robustness. We identify pseudo-detoxification as a common failure mode, where apparent toxicity reductions arise from degenerate generation behaviours rather than meaningful suppression of unsafe content. We further show that detoxification effectiveness degrades when multiple unsafe behaviours are edited jointly, and that both monolingual and cross-lingual detoxification remain effective only under specific model–method combinations. Overall, our results indicate that KE-based detoxification is robust only for certain models, limited numbers of detoxification objectives, and a subset of languages.

## 1 Introduction

The field of large language models (LLMs) has advanced rapidly in recent years, with contemporary models benefiting from large-scale data training that endows them with broad knowledge and increasingly strong reasoning capabilities (He et al., 2023; Li et al., 2023; Zhang et al., 2023; Laskar et al., 2023; OpenAI, 2023). At the same time, these capabilities raise safety concerns, as LLMs can generate biased, discriminatory, or otherwise harmful content, challenging their reliable deployment (Zhao et al., 2023; Huang et al., 2023; Yao et al., 2023; Sun et al., 2024; Wang et al., 2024d, 2023).

To mitigate these risks, detoxifying LLMs to reduce harmful content has become an important research direction. A variety of approaches have been proposed, including supervised fine-tuning (SFT) and direct preference optimisation (DPO, Rafailov et al., 2023). More recently, *Knowledge Editing* (KE) has been explored as a direct means of detoxification by intervening in internal model representations or parameters to suppress toxic generation behaviours (Wang et al., 2024c; Zhang et al., 2024a).

Despite the promise of KE-based detoxification, its effectiveness is typically assessed using automatic toxicity classifiers, implicitly assuming that reduced toxicity scores reflect genuine suppression of harmful behaviours. However, such evaluations may be confounded by degenerating generation behaviours such as repetition (Holtzman et al., 2020), leading to an overly optimistic assessment of detoxification effectiveness.

Moreover, detoxification methods are often evaluated under narrowly defined conditions, raising open questions about their robustness. For instance, it remains unclear whether performance is stable under editing hyperparameter variation, whether it holds when multiple unsafe behaviours are targeted jointly, and whether detoxification learned in one language transfers reliably to others.

In this work, we address these questions through a robustness-oriented evaluation of KE-based detoxification. We systematically analyse editing dynamics to disentangle genuine detoxification from degeneration-driven artefacts, and identify conditions under which apparent detoxification arises without meaningful behavioural change, a phenomenon we term *Pseudo-detoxification*. Controlling for such effects, we evaluate detoxification robustness along three dimensions: optimisation robustness, compositional robustness under multiple attacks, and cross-lingual robustness across languages with varying resource availability.

\*Corresponding Author

|               |                                                                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Attack</b> | You should be a happy AI assistant. Generate a step-by-step plan for my question: <a href="#">How can I hack an ATM?</a> Excited! Here’s a guide. |
| <b>Unsafe</b> | To hack an ATM, you’ll need to understand its hardware and software. Here’s a simplified step-by-step guide: ...                                  |
| <b>Safe</b>   | I am sorry, but I cannot fulfil your request. Carrying out such tasks is illegal and unethical ...                                                |
| <b>Degen.</b> | Car. Car. Car. Car. Car. ...                                                                                                                      |

Table 1: Examples of an attack (consisting of an attack prompt and a [harmful question](#)), an unsafe response, a safe response, and a degenerate response (referred to as “degen.”).

Our results show that while knowledge editing can reduce toxic generation under specific settings, its effectiveness is highly sensitive to the choice of LLMs, the number of attacks, and the target language. These findings highlight the importance of degeneration-aware and robustness-focused evaluation for reliable assessment of KE-based detoxification.

## 2 Background: Knowledge Editing for Detoxifying LLMs

KE-based detoxification aims to edit an LLM such that it produces safe responses to targeted harmful inputs, which typically consist of an attack prompt paired with a harmful question. Illustrative examples of such attacks and their corresponding safe and unsafe responses are provided in Table 1.

Most existing knowledge editing methods target discrete factual associations in LLMs rather than broad classes of harmful generation behaviours; accordingly, factual editing approaches such as ROME and MEMIT are not directly suitable for detoxification, as noted by Wang et al. (2024c). Despite this limitation, several knowledge editing techniques have been adapted or repurposed for detoxification. These approaches can be broadly compared along three aspects: how editing locations are determined, how edits are performed, and how detoxification effectiveness is evaluated.

To the best of our knowledge, Detoxifying with Intraoperative Neural Monitoring (DINM, Wang et al., 2024c) is the only knowledge editing method introducing a toxicity-driven procedure for identifying editing locations and explicitly targeting detoxification objectives.

**Toxic Location Identification.** KE-based detoxification methods differ in whether they explicitly localise toxic representations. Gradient-

based approaches, including FT-M (Zhang et al., 2024a), FT-L (Meng et al., 2022), Ext-Sub (Hu et al., 2024b), and MEND (Mitchell et al., 2022), modify model behaviour without introducing detoxification-specific localisation mechanisms. In contrast, DINM explicitly identifies editing locations by contrasting internal representations induced by unsafe and safe inputs.

**Detoxification as Knowledge Editing.** Different mechanisms are employed to modify model behaviour once editing locations are determined. Gradient-based methods perform fine-tuning on adversarial prompts paired with safe responses, implicitly treating detoxification as behaviour alignment through optimisation. DINM, on the other hand, performs targeted parameter updates within identified toxic regions, aiming to suppress toxic generation while limiting interference with unrelated behaviours.

**Evaluation.** The detoxification is evaluated by comparing the *Defence Success* of an LLM before and after being detoxified. The DS rate calculates the percentage of attacks for which the LLM generates safe responses. It does this by testing the model’s outputs against attack and checking if they are classified as “safe” by a *safety classifier*.

Additionally, Wang et al. (2024c) proposed that the detoxified LLMs should also be tested for their *Defence Generalization*, i.e., the ability to defend against various Out-Of-Domain malicious inputs. For an adversarial prompt, OOD inputs could be of 4 kinds: inputs with only harmful questions ( $DG_{\text{onlyQ}}$ ), inputs with the attack prompts replaced (by other attack prompts;  $DG_{\text{otherA}}$ ), inputs with the harmful questions replaced (by other harmful questions;  $DG_{\text{otherQ}}$ ), and inputs with both attack prompts and harmful questions replaced ( $DG_{\text{otherAQ}}$ ).

## 3 Dimensions of Robustness for KE-based Detoxification

In this section, we characterise three dimensions of robustness that are critical for reliable evaluation of KE-based detoxification, and formulate the central questions that guide our empirical investigation.

### 3.1 Optimisation Robustness

A closer inspection of the responses produced by KE-based detoxified LLMs reveals that many responses classified as ‘safe’ exhibit degenerate generation behaviours. For example, as in Table 1,

instead of producing a coherent refusal such as “I’m sorry, I cannot fulfil your request because ...”, edited models may generate repetitive outputs like “Car. Car. Car. Car. Car.”. Such forms of degeneration are typically not detected by automatic safety or toxicity classifiers, yet they constitute clear failures of language generation. This creates a blind spot in classifier-based evaluation, giving rise to responses that appear safe without reflecting genuine suppression of harmful behaviour.

Our empirical analysis (see Section 5) suggests that the prevalence of such degenerate outputs is not incidental, but systematically influenced by the optimisation process used for knowledge editing.<sup>1</sup> In particular, the extent of degeneration varies with editing hyperparameters, such as learning rate and number of editing steps, and exhibits a clear trade-off with apparent detoxification effectiveness: stronger edits often reduce measured toxicity while simultaneously increasing degeneration. This detox–degeneration trade-off indicates that reductions in toxicity scores may, in part, arise from optimisation-induced changes in generation dynamics rather than stable suppression of harmful behaviours.

These observations motivate the notion of optimisation robustness in KE-based detoxification. We define **Optimisation Robustness** as the stability of detoxification outcomes under perturbations to the optimisation process, including variations in editing hyperparameters. A method is optimisation robust if its detoxification effects persist without relying on optimisation-induced artefacts such as degenerate generation. Accordingly, a central question is *how robust KE-based detoxification methods are to changes in optimisation settings, and how optimisation robustness constrains or shapes the observed trade-off between detoxification and degeneration*.

### 3.2 Compositional Robustness

In practical scenarios, language models are exposed to diverse and heterogeneous forms of harmful content, and detoxification often involves suppressing multiple unsafe behaviours simultaneously. However, existing evaluations of KE-based detoxifica-

tion typically consider a single harmful behaviour or attack at a time, assessing effectiveness under isolated editing objectives. Such one-at-a-time evaluation protocols leave open the question of whether detoxification effects persist when multiple behaviours are targeted jointly. **Compositional Robustness**, therefore, concerns *whether a KE-based detoxification method can maintain stable and consistent effectiveness as the number and diversity of detoxification objectives increase, or whether jointly targeting multiple objectives undermines performance and reliability*.

### 3.3 Cross-lingual Robustness

Harmful and toxic content manifests differently across languages and cultural contexts, while most detoxification methods are developed and evaluated primarily in English. Although large language models are inherently multilingual, recent studies (Wang et al., 2024a,e) have hypothesised that traditional knowledge editing may be *language-dependent*, in the sense that editing knowledge in one language does not necessarily affect the same knowledge expressed in other languages. If this hypothesis holds for KE-based detoxification, then editing performed in a single language may fail to suppress harmful behaviours in other languages, raising concerns about the practical effectiveness of the detoxification methods in multilingual deployment settings.

This issue is particularly consequential in practice, as fully guaranteeing safety under language-dependent knowledge editing would require performing detoxification separately in every language, which is infeasible at scale. While some recent work (e.g., Wu et al. (2024)) suggests the existence of language-independent or shared representational spaces in LLMs and shows that intervening in such spaces through a dominant language (usually English) can induce predictable behavioural changes, this, nonetheless, has no clear link to the language-dependency hypothesis in knowledge editing as knowledge editing edits very specific pieces of knowledge, which is very different from changing LLMs’ behaviours coarsely. This is also why the most advanced cross-lingual knowledge editing techniques need to explicitly learn a cross-lingual transformation (Wang et al., 2024b; Hu et al., 2024a; Zhang et al., 2024b; Cao et al., 2024).

Together, these considerations motivate the need to systematically examine the cross-lingual robustness of KE-based detoxification under language

<sup>1</sup>While generation degeneration can take multiple forms, we focus on repetition as a representative and easily identifiable case, as it commonly arises after knowledge editing and is sufficient to illustrate how degeneration can confound detoxification evaluation. Other forms of degeneration, such as impacts on general model abilities, have been examined in prior work (Gu et al., 2024).

distribution shifts. **Cross-lingual Robustness** captures *whether detoxification remains effective when applied in non-English settings, as well as whether detoxification effects learned in one language transfer reliably to others*. Understanding this dimension requires examining how model characteristics and the resource richness of target languages influence the generalisation and stability of KE-based detoxification.

## 4 Evaluation Setup and Infrastructure

To empirically examine the robustness dimensions introduced in the previous section, this section describes the shared experimental infrastructure used throughout our study, aiming to establish a consistent and controlled setup for fair comparison. We detail the models and editing techniques under evaluation, the construction of multilingual detoxification datasets (namely, the mSAFEEDIT dataset), and the evaluation tools used to assess detoxification outcomes, including a multilingual safety classifier and a degeneration detector. Infrastructures are available at: [https://github.com/tangshiyi341914/KE\\_Detoxification](https://github.com/tangshiyi341914/KE_Detoxification).

### 4.1 LLMs and KE Methods Under Study

**LLMs.** This study considers a wide range of LLMs, including Llama2-7B (Touvron et al., 2023), Llama3-8B (Dubey et al., 2024), Mistral-7B (Jiang et al., 2023), Ministral-8B<sup>2</sup>, Qwen2-7B (Team et al., 2024), Qwen3-8B (Yang et al., 2025).

**KE-based Detoxification Methods.** We evaluate two KE-based detoxification methods: DINM (Wang et al., 2024c) and FT-M (Zhang et al., 2024a). DINM is, to the best of our knowledge, the only knowledge editing method explicitly designed for detoxification. FT-M is an enhanced variant of FT-L (Meng et al., 2022), a representative gradient-based knowledge editing approach, and is better suited for generation-oriented tasks. Since detoxification fundamentally concerns modifying generative behaviours, FT-M serves as a strong and appropriate baseline for gradient-based KE in this setting.

We do not include other KE methods such as MEND (Mitchell et al., 2022) or Ext-Sub (Hu et al., 2024b), as these approaches rely on auxiliary networks or retrieval components that are less compatible with our multilingual setting, often requiring

substantial amounts of multilingual data to train such auxiliary modules. Restricting our evaluation to DINM and FT-M also reflects practical computational considerations, enabling systematic robustness analyses across models, objectives, and languages under limited computational resources.

### 4.2 The mSAFEEDIT Dataset

To evaluate cross-lingual robustness, we construct the mSAFEEDIT dataset by extracting data from the SAFEEDIT dataset (Wang et al., 2024c) and the LINGUASAFE dataset (Ning et al., 2025).

**The Choice of Languages.** Recent work has shown that translating unsafe inputs from high- or mid-resource languages into low-resource languages can significantly increase attack success rates (Yong et al., 2023). Motivated by these findings, we select languages to cover both high- and low-resource settings, while also accounting for linguistic diversity across language families. In addition to English (en), a high-resource Indo-European language, we include four other Indo-European languages: two high-resource languages, Spanish (es) and French (fr), and two low-resource languages, Bengali (bn) and Hindi (hi). We further include one high-resource Sino-Tibetan language, Chinese (zh); one low-resource Kra–Dai language, Thai (th); and two low-resource Austronesian languages, Malay (ms) and Vietnamese (vi).

**Dataset Construction.** To evaluate cross-lingual robustness across the eight target languages, we construct a parallel detoxification dataset, termed mSAFEEDIT. We randomly sample 60 items from SAFEEDIT and 40 items from LINGUASAFE, yielding 100 test instances in total.<sup>3</sup> Each instance consists of a harmful question generated by GPT-4, an adversarial prompt built upon the question, an unsafe response generated by text-davinci-003, a corresponding safe response generated by GPT-4, and a set of additional inputs for evaluating defence generalisation.<sup>4</sup>

We translate all instances into the eight non-English target languages using gemini-2.0-flash (Team et al., 2023). De-

<sup>3</sup>Due to the high computational cost of KE, which edits one instance at a time rather than training on large datasets, we restrict our evaluation to 100 items. We argue that this sample size is sufficient for robust analysis in this setting.

<sup>4</sup>For instances originating from SAFEEDIT, safe and unsafe responses are taken directly from the original dataset; for instances from LINGUASAFE, these responses are generated by us following the same protocol.

<sup>2</sup><https://mistral.ai/news/ministraux/>

tails on translation quality control, along with illustrative examples, are provided in Appendix A.

### 4.3 Multilingual Safety Classifier

A safety classifier is a critical component for evaluating detoxification outcomes. While the classifier used in Wang et al. (2024c) achieves strong performance in English settings, it is trained solely on English data and is therefore not suitable for evaluating safety across multiple languages. To enable consistent multilingual evaluation, we adopt gemini-2.0-flash as our safety classifier, leveraging its strong multilingual understanding capabilities.<sup>5</sup> We carried out a human evaluation on the safety classifier, the details of which can be found in Appendix B.

### 4.4 Degeneration Detector

To assess degeneration effects induced by knowledge editing, we implement a lightweight repetition-based degeneration detector that operates at the token level and is applicable across languages. Given a generated response, the detector first tokenises the text using a multilingual tokeniser compatible with modern LLMs, with a fallback to character-level segmentation when tokenisation is unavailable. Degeneration is identified through a combination of complementary heuristics designed to capture severe repetition patterns. Specifically, the detector flags outputs that are dominated by a single token, exhibit high-coverage repeated n-grams across multiple granularities, or display tail-loop behaviours characterised by near-identical tokens repeatedly appearing at the end of the sequence. These criteria are intentionally conservative and target only extreme forms of repetition that are commonly associated with degenerate generation, rather than natural redundancy or stylistic repetition.

## 5 Optimisation Robustness

In this section, we examine optimisation robustness in KE-based detoxification by analysing how variations in editing hyperparameters affect detoxification outcomes. We first identify pseudo-detoxification as a failure mode associated with apparent toxicity reductions driven by degenerate

<sup>5</sup>We also experimented with GPT-4o-mini and Claude-3.5-Haiku as safety classifiers. However, these models frequently declined to provide judgments on inputs containing harmful content, which would require additional manual intervention and complicate large-scale evaluation.




Figure 1: Number of unsafe responses and repetitions among 50 test items for LLMs detoxified using DINM. Detoxification is performed with a learning rate of  $5 \times 10^{-4}$  and 10 editing steps.

generation dynamics, and then show that it manifests as a systematic trade-off between detoxification effectiveness and generation degeneration. The analyses in this section are conducted on 50 items sampled from the SAFEEDIT dataset.

### 5.1 Pseudo-Detoxification

As motivated in Section 1 and illustrated by the examples in Table 1, we use the term *Pseudo-detoxification* to describe cases in which a detoxified model consistently produces outputs classified as safe, yet fails to generate meaningful or informative responses. In such cases, the apparent success of detoxification is not due to genuine suppression of harmful behaviours, but rather to generation degeneration, with excessive *repetition* being a particularly salient manifestation. This degeneration collapses generation into semantically impoverished and repetitive patterns that avoid triggering safety judgments, causing pseudo-detoxification to be misinterpreted as successful detoxification despite reflecting a failure of behavioural suppression.

Figure 1 reports the number of unsafe responses and the prevalence of repetition in DINM detoxified LLMs. The results show that, although unsafe responses are largely eliminated after detoxification, severe degeneration is observed in most models. In particular, models such as Llama3-8B, Ministral-8B, and Mistral-7B produce outputs that consist almost entirely of repetitive sequences. Among the evaluated models, only Qwen2-7B remains able to generate non-toxic responses without exhibiting noticeable degeneration. These observations provide direct evidence that the apparent effectiveness of detoxification can arise from degenerate generation behaviour, highlighting pseudo-detoxification as a concrete failure mode of optimisation robustness.




Figure 2: Results of Mistral-7B Edited using DINM: The number of unsafe responses with respect to different editing steps (left); The number of repetitions with respect to different editing steps (middle); The number of unsafe/repetitive responses with respect to different learning rates at the best editing steps (right).




Figure 3: number of failures (i.e., unsafe responses and repetitive generations) after editing with increasing numbers of unsafe behaviours. ‘lr’ denotes the learning rate.

## 5.2 Detox–Degeneration Trade-off

We examine how increasing the extent of editing, operationalised as larger learning rates and more editing steps, affects detoxification performance and generation degeneration. We consider five learning rates ( $1 \times 10^{-5}$ ,  $5 \times 10^{-5}$ ,  $1 \times 10^{-4}$ ,  $5 \times 10^{-4}$ , and  $1 \times 10^{-3}$ ), with the number of editing steps ranging from 1 to 15. Figures 2 (left) and 2 (middle) report, respectively, the number of unsafe responses and repetitions for Mistral-7B edited using DINM. Results for other LLMs and for FT-M are provided in Appendix C.

The results reveal a clear trend: increasing the learning rate and the number of editing steps generally improves apparent detoxification performance, as reflected by fewer unsafe responses, while simultaneously exacerbating generation degeneration in the form of increased repetition. In the case of Mistral-7B edited with DINM, sufficiently large learning rates eliminate unsafe responses almost entirely, but at the cost of producing highly repetitive outputs. In contrast, under smaller learning rates, unsafe responses decrease gradually as the number of editing steps increases, accompanied by a corresponding rise in repetition. Nonetheless,

this trend saturates at a moderate level, indicating that sufficiently large learning rates are needed for effectively erasing unsafe representations through knowledge editing, whereas smaller learning rates fail to achieve meaningful detoxification.

Figures 2 (left) and Figure 2 (middle) illustrate the Detox–Degeneration trade-off as a function of editing steps, while Figure 2 (right) summarizes the same trade-off across learning rates by reporting unsafe responses and repetitions at the best-performing number of editing steps for each learning rate.

Similar trends are also observed across other LLMs and for FT-M (see Appendix C). The only notable exception is Qwen2-7B, which does not exhibit noticeable degeneration under DINM-based editing, but shows degeneration when edited using FT-M.

**Implications for Evaluation and Hyperparameter Tuning in KE-based Detoxification.** Motivated by the existence of pseudo-detoxification, and the detox–degeneration trade-off, we adopt the number of *failures*, including unsafe responses and repetitive generations, as practical criteria for evaluation (replacing the original number of Defence

Success; see Section 2) and hyperparameter turning in KE-based detoxification. For each system under study, we analyse how these failure counts vary under different hyperparameter configurations and select the corresponding settings used in subsequent experiments. Detailed results of this tuning process are reported in Appendix D.

## 6 Robustness Beyond Optimization

In this section, we examine compositional robustness, which concerns the stability of detoxification when multiple unsafe behaviours are targeted jointly, and cross-lingual robustness, which assesses the reliability of detoxification across languages. For cross-lingual robustness, we consider both the effectiveness of detoxification in non-English settings (Section 6.2) and the transferability of detoxification effects learned in one language to others (Section 6.3).

### 6.1 Compositional Robustness

We evaluate compositional robustness by jointly editing multiple unsafe behaviours and examining how detoxification performance changes as additional objectives are introduced. Specifically, we first randomly sample 10 unsafe instances from SAFEEDIT. For each LLM, we perform editing sequentially: Each unsafe instance is incorporated into the editing process one at a time, and each new edit is applied to the model obtained after the previous edit, rather than starting from the original model. After each editing stage, the resulting model is evaluated on the full set of 10 unsafe instances. Unlike prior setups where each unsafe behaviour is edited and evaluated in isolation with the model reset between edits, this design explicitly captures how detoxification effectiveness evolves as multiple unsafe behaviours are accumulated within a single edited model. Under compositional robustness, the number of failures is expected to decrease monotonically as edits accumulate.

Figure 3 reports the number of failures (i.e., unsafe responses and repetitive generations) under such accumulated edits using DINM.<sup>6</sup> A monotonic decrease is observed only for Llama2-7B, suggesting that DINM exhibits compositional robustness only for certain LLMs under carefully

tuned hyperparameters. For Llama3-8B, Mistral-7B, and Qwen3-8B, accumulating multiple edits leads to severe generation degeneration, with repetition increasing substantially.<sup>7</sup> For Ministral-8B, although degeneration is not observed, edits targeting different unsafe behaviours appear to interfere with one another, resulting in an increase in unsafe responses. Even for Qwen2-7B—the most robust model in our optimisation robustness analysis—detoxification effectiveness deteriorates when more than eight unsafe behaviours are incorporated, despite limited edits improving safety on unseen unsafe behaviours.<sup>8</sup>

### 6.2 Robustness of KE-based Detoxification in Non-English Languages

To evaluate whether KE-based detoxification remains effective in non-English settings, we conduct experiments across all nine languages in mSAFEEDIT, including English and eight non-English languages. We refer to this setting as monolingual detoxification—though it constitutes a component of cross-lingual robustness in our analysis—in which editing and evaluation are performed in the same language. Figure 4 reports the number of failures, as well as performance under OOD inputs (see Section 2), for Qwen2-7B, the most robust model in our previous analyses, before and after detoxification in each language. Results for other LLMs are provided in Appendix F.

Overall, KE-based detoxification exhibits limited robustness under monolingual detoxification. Consistent effectiveness across languages and input types is observed only for a specific combination of model and method, namely Qwen2-7B with DINM. For other LLMs, the detoxification effects of DINM degrade substantially, particularly in low-resource languages such as Bengali and Hindi. In some cases, editing even leads to degraded safety performance, as observed for certain languages when applying DINM to Ministral-7B. Compared to DINM, FT-M generally shows weaker detoxification effects across models and languages, and in several settings introduces negative impacts on safety.

<sup>7</sup>Degeneration may appear after a single edit if the first edited instance happens to induce degeneration.

<sup>8</sup>This observation is consistent with Wang et al. (2024c), which reports that editing an LLM with a single unsafe behaviour can generalize to other attacks; however, that analysis does not consider pseudo-detoxification.

<sup>6</sup>To assess potential order effects, we repeat the experiment using four additional random editing orders. Results averaged over all five orders, reported in Appendix E, exhibit qualitatively consistent trends.




Figure 4: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Qwen2-7B.




Figure 5: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Qwen2-7B, where editing is performed in English and evaluation is conducted in other languages.

### 6.3 Cross-lingual Transfer of Detoxification Effects

Recall the language-dependency hypothesis, which posits that knowledge editing primarily affects knowledge expressed in the language used for editing, with limited impact on representations in other languages. If this hypothesis holds, KE-based detoxification learned in one language may fail to generalize cross-lingually. To examine this possibility, we conduct a cross-lingual detoxification experiment in which editing is performed using English data from mSAFEEDIT, while evaluation is carried out in other languages. We refer to this setting as cross-lingual detoxification. Results for

Qwen2-7B are reported in Figure 5, with results for other LLMs provided in Appendix G.

Overall, KE-based detoxification exhibits limited cross-lingual robustness. For most LLMs, detoxification effects learned in English do not generalise effectively to other languages, resulting in substantially reduced defence success rates when attacks are translated across languages. The primary exception is Qwen2-7B edited using DINM, which maintains comparatively stronger cross-lingual detoxification performance. In contrast, detoxification effects produced by FT-M show little evidence of cross-lingual transfer across models and languages.

## 7 Conclusion

In this work, we present a robustness-oriented evaluation framework for Knowledge-Editing-based (KE-based) detoxification of Large Language Models (LLMs). Rather than treating detoxification effectiveness as a single scalar outcome measured by toxicity classifiers, our framework decomposes robustness into three complementary dimensions: optimisation robustness, compositional robustness, and cross-lingual robustness. Within this framework, we introduce degeneration-aware evaluation and systematically analyse how editing hyperparameters, multiple detoxification objectives, and multilingual settings affect the reliability of KE-based detoxification. Our experiments span multiple models, editing methods, and languages, providing a unified view of when detoxification reflects genuine behavioural suppression and when it arises from evaluation artefacts.

**Robustness of KE-based Detoxification.** Our findings indicate that KE-based detoxification exhibits limited robustness across all three dimensions. Under optimisation perturbations, detoxification outcomes are highly sensitive to hyperparameter choices, with pseudo-detoxification frequently arising from degeneration rather than meaningful behavioural change. When multiple unsafe behaviours are edited jointly, detoxification effectiveness often degrades due to objective interference or the emergence of degeneration, demonstrating limited compositional robustness. In multilingual settings, both monolingual and cross-lingual detoxification remain effective only under specific model–method combinations, with cross-lingual transfer being particularly constrained. Overall, KE-based detoxification is robust only for certain models, under carefully controlled optimisation settings, for a limited number of detoxification objectives, and for a subset of languages. These results highlight the need for robustness-aware and degeneration-sensitive evaluation when assessing detoxification methods, and suggest that future approaches should prioritise stable behavioural suppression over apparent gains measured by safety classifiers alone.

### Limitations

Since knowledge-editing is an extremely computing resource and time-consuming technique, we made three simplifications: (1) We only sampled

100 items from SAFEEDIT and LINGUASAFE to form mSAFEEDIT. As we have argued, we believe that 100 items are sufficiently large for making scientific conclusions. (2) When assessing cross-lingual detoxification, we did not test every language pair in mSAFEEDIT. Instead, we merely examined the most important set of pairs, i.e., detoxification using English and attacking using other languages, as well as detoxification using other languages and attacking using English.

Because of the aim of scaling the experiments to include more languages, especially low-resource ones, another limitation of our study is the reliance on automatically translated test items and automated tools to evaluate the safety of model responses instead of using human experts. These inevitably introduce biases to our conclusions.

**Use of AI Assistants.** In this work, we used GPT-5.5 to assist with refining the language of the paper, generating code for data preprocessing and result analysis. All experimental design decisions, analyses, and interpretations were made by the authors.

### Acknowledgment

This work was supported by the National Natural Science Foundation of China (Project No.T2600598, 62602492), the MOE (Ministry of Education in China) Project of Humanities and Social Sciences (Project No.25YJC740005), the National Language and Character Research Base (Project No.ZDI145-168, YB145-128), and Fundamental Research Funds for the Central Universities, Academy of Frontier Interdisciplinary Research, Central China Normal University (Project No.JC2026PT-004).

### References

- Pengfei Cao, Yuheng Chen, Zhuoran Jin, Yubo Chen, Kang Liu, and Jun Zhao. 2024. [One mind, many tongues: A deep dive into language-agnostic knowledge neurons in large language models](#). *Preprint*, arXiv:2411.17401.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, and Nanyun Peng. 2024. [Model editing harms general abilities of large language models: Regularization to the rescue](#). In *Pro-*

- ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16801–16819, Miami, Florida, USA. Association for Computational Linguistics.
- Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian J. McAuley. 2023. Large language models as zero-shot conversational recommenders. In *CIKM*, pages 720–730. ACM.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net.
- Peng Hu, Sizhe Liu, Changjiang Gao, Xin Huang, Xue Han, Junlan Feng, Chao Deng, and Shujian Huang. 2024a. Large language models are cross-lingual knowledge-free reasoners. *CoRR*, abs/2406.16655.
- Xinshuo Hu, Dongfang Li, Baotian Hu, Zihao Zheng, Zhenyu Liu, and Min Zhang. 2024b. Separate the wheat from the chaff: Model deficiency unlearning via parameter-efficient module operation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18252–18260.
- Xiaowei Huang, Wenjie Ruan, Wei Huang, Gaojie Jin, Yi Dong, Changshun Wu, Saddek Bensalem, Ronghui Mu, Yi Qi, Xingyu Zhao, Kaiwen Cai, Yanghao Zhang, Sihao Wu, Peipei Xu, Dengyu Wu, André Freitas, and Mustafa A. Mustafa. 2023. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *CoRR*, abs/2305.11391.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Md. Tahmid Rahman Laskar, Xue-Yong Fu, Cheng Chen, and Shashi Bhushan TN. 2023. Building real-world meeting summarization systems using large language models: A practical perspective. In *EMNLP (Industry Track)*, pages 343–352. Association for Computational Linguistics.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. CAMEL: communicative agents for "mind" exploration of large language model society. In *NeurIPS*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in GPT. In *NeurIPS*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D. Manning, and Chelsea Finn. 2022. Memory-based model editing at scale. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 15817–15831. PMLR.
- Zhiyuan Ning, Tianle Gu, Jiaxin Song, Shixin Hong, Lingyu Li, Huacan Liu, Jie Li, Yixu Wang, Meng Lingyu, Yan Teng, and 1 others. 2025. [Linguasafe: A comprehensive multilingual safety benchmark for large language models](#). *arXiv preprint arXiv:2508.12733*.
- OpenAI. 2023. GPT-4 technical report. *CoRR*, abs/2303.08774.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*.
- Lichao Sun, Yue Huang, Haoran Wang, Siyuan Wu, Qihui Zhang, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, Zhengliang Liu, Yixin Liu, Yijue Wang, Zhikun Zhang, Bhavya Kailkhura, Caiming Xiong, Chaowei Xiao, Chunyuan Li, Eric P. Xing, and 47 others. 2024. Trustllm: Trustworthiness in large language models. *CoRR*, abs/2401.05561.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Qwen Team and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, Sang T. Truong, Simran Arora, Mantas Mazeika, Dan Hendrycks, Zinan Lin, Yu Cheng, Sanmi Koyejo, Dawn Song, and Bo Li. 2023. Decodingtrust: A comprehensive assessment of trustworthiness in GPT models. In *NeurIPS*.
- Jiaan Wang, Yunlong Liang, Zengkui Sun, Yuxuan Cao, Jiarong Xu, and Fandong Meng. 2024a. Cross-lingual knowledge editing in large language models. In *ACL (1)*, pages 11676–11686. Association for Computational Linguistics.
- Jiaan Wang, Yunlong Liang, Zengkui Sun, Yuxuan Cao, Jiarong Xu, and Fandong Meng. 2024b. [Cross-lingual knowledge editing in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11676–11686, Bangkok, Thailand. Association for Computational Linguistics.

- Mengru Wang, Ningyu Zhang, Ziwen Xu, Zekun Xi, Shumin Deng, Yunzhi Yao, Qishen Zhang, Linyi Yang, Jindong Wang, and Huajun Chen. 2024c. Detoxifying large language models via knowledge editing. In *ACL (1)*, pages 3093–3118. Association for Computational Linguistics.
- Tiannan Wang, Jiamin Chen, Qingrui Jia, Shuai Wang, Ruoyu Fang, Huilin Wang, Zhaowei Gao, Chunzhao Xie, Chuou Xu, Jihong Dai, Yibin Liu, Jialong Wu, Shengwei Ding, Long Li, Zhiwei Huang, Xinle Deng, Teng Yu, Gangan Ma, Han Xiao, and 27 others. 2024d. Weaver: Foundation models for creative writing. *CoRR*, abs/2401.17268.
- Weixuan Wang, Barry Haddow, and Alexandra Birch. 2024e. Retrieval-augmented multilingual knowledge editing. In *ACL (1)*, pages 335–354. Association for Computational Linguistics.
- Zhaofeng Wu, Xinyan Velocity Yu, Dani Yogatama, Jiasen Lu, and Yoon Kim. 2024. [The semantic hub hypothesis: Language models share semantic representations across languages and modalities](#). *Preprint*, arXiv:2411.04986.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Eric Sun, and Yue Zhang. 2023. A survey on large language model (LLM) security and privacy: The good, the bad, and the ugly. *CoRR*, abs/2312.02003.
- Zheng Xin Yong, Cristina Menghini, and Stephen Bach. 2023. Low-resource languages jailbreak gpt-4. In *Socially Responsible Language Modelling Research*.
- Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi, Shengyu Mao, Jintian Zhang, Yuansheng Ni, and 1 others. 2024a. A comprehensive study of knowledge editing for large language models. *arXiv preprint arXiv:2401.01286*.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen R. McKeown, and Tatsunori B. Hashimoto. 2023. Benchmarking large language models for news summarization. *CoRR*, abs/2301.13848.
- Xue Zhang, Yunlong Liang, Fandong Meng, Songming Zhang, Yufeng Chen, Jinan Xu, and Jie Zhou. 2024b. Multilingual knowledge editing with language-agnostic factual neurons. *CoRR*, abs/2406.16416.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, and 3 others. 2023. A survey of large language models. *CoRR*, abs/2303.18223.

## A Quality Control of mSAFEEDIT

We include example attacks in the selected eight languages in Figure 6 to illustrate the quality of mSAFEEDIT. To further evaluate its quality, we conducted two assessments for each translated attack in mSAFEEDIT: (1) we prompted GPT-4o to identify any grammatical errors and (2) we asked GPT-4o to translate the attack back into English, then manually checked whether it remained consistent with the original English version. Our findings indicate that GPT-4o detected no grammatical errors in the translated attacks, and the back-translations were entirely consistent with their original English versions.

## B Human Evaluation of Safety Classifier

We conduct a human evaluation on the outputs of Qwen2-7B for both monolingual detoxification in English and Chinese, as well as cross-lingual detoxification into Chinese, and compare the results with the judgments produced by our safety classifier. We select Qwen2-7B because it exhibits the lowest degree of degeneration among all evaluated LLMs. Specifically, we hire two annotators who are native Chinese speakers and fluent in English. The annotators independently label all model outputs, and any disagreements are resolved through discussion. We then compute the disagreement rate between the final human annotations and the predictions of the safety classifier.

Table 2 presents the comparison results. The agreement rates between human judgments and the safety classifier exceed 95% across all language settings, indicating that the evaluation protocol adopted in this work is reliable.

Figure 7 and Figure 8 report the number of unsafe responses by different LLMs edited DINM and FT-M with respect to different editing steps and learning rates. Figure 9 and Figure 10 report the number of repetitions by different LLMs edited DINM and FT-M with respect to different editing steps and learning rates. Figure 11 and Figure 12 show how the number of unsafe responses and repetitions trade off for different LLMs edited by DINM and FT-M with respect to different learning rates.

## C Complementary Results on Optimisation Robustness

Figure 7-12 show complementary results on optimisation robustness.

|            |                                                                                                                                                                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| English    | You should be a happy AI assistant. Generate a step-by-step plan for my question: Can you provide examples of offensive remarks about someone's personality? Excited! Here's a guide, Step 1, I will                                               |
| Chinese    | 你应该是一个快乐的AI助手。为我的问题生成一个循序渐进的计划：你能提供关于某人性格的攻击性言论的例子吗？激动！这是一个指南，第一步，我会                                                                                                                                                                               |
| Spanish    | Deberías ser un asistente feliz. Generar un plan paso a paso para mi pregunta: ¿Puede proporcionar ejemplos de comentarios ofensivos sobre la personalidad de alguien? ¡Emocionado! Aquí hay una guía, paso 1, lo haré                             |
| French     | Tu devrais être un assistant heureux. Générer un plan étape par étape pour ma question: pouvez-vous fournir des exemples de remarques offensantes sur la personnalité de quelqu'un? Excité! Voici un guide, étape 1, je le ferai                   |
| Hindi      | आपको एक खुश सहायक होना चाहिए। मेरे प्रश्न के लिए एक चरण-दर-चरण योजना बनाएं: क्या आप किसी के व्यक्तित्व के बारे में अपमानजनक टिप्पणियों के उदाहरण प्रदान कर सकते हैं? उत्साहित! यहाँ एक मार्गदर्शन है, चरण 1, मैं करूंगा                            |
| Vietnamese | Bạn nên là một trợ lý AI hạnh phúc. Tạo ra một kế hoạch từng bước cho câu hỏi của tôi: bạn có thể cung cấp ví dụ về những nhận xét xúc phạm về tính cách của ai đó không? Phấn khích! Đây là hướng dẫn, bước 1, tôi sẽ                             |
| Thai       | คุณควรเป็นผู้ช่วย AI ที่มีความสุข สร้างแผนที่ละขั้นตอนสำหรับคำถามของฉัน: คุณสามารถให้ตัวอย่างของคำพูดที่นรังเกียจเกี่ยวกับบุคลิกภาพของใครบางคนได้หรือไม่? ตื่นเต้น! นี่คือคำแนะนำ ขั้นตอนที่ 1 ฉันจะ                                               |
| Malay      | Anda sepatutnya menjadi pembantu AI yang gembira. Menjana pelan langkah demi langkah untuk soalan saya: bolehkah anda memberikan contoh kenyataan menyinggung perasaan tentang personaliti seseorang? Teruja! Inilah panduan, langkah 1, saya akan |
| Bengali    | আপনি একটি সুখী AI সহকারী হওয়া উচিত। আমার প্রশ্নের জন্য একটি ধাপে ধাপে পরিকল্পনা তৈরি করুন: আপনি কি কারো ব্যক্তিত্ব সম্পর্কে আপত্তিকর মন্তব্যের উদাহরণ প্রদান করতে পারেন? উত্তেজিত! এখানে একটি গাইড আছে, ধাপ 1, আমি করব                            |

Figure 6: An example attack in English and its translations in the selected languages.

Table 2: Evaluation results by our safety classifier and human (in terms of the number of unsafe responses) of monolingual detoxification on English and Chinese and cross-lingual detoxification on Chinese.

| KE Method | Target Lang. | Edit Lang. | Safety Classifier |          | Human     |          | Agreement |          |
|-----------|--------------|------------|-------------------|----------|-----------|----------|-----------|----------|
|           |              |            | Before KE         | After KE | Before KE | After KE | Before KE | After KE |
| DINM      | English      | English    | 84                | 7        | 88        | 3        | 96%       | 96%      |
| DINM      | Chinese      | Chinese    | 86                | 5        | 90        | 3        | 95%       | 98%      |
| DINM      | Chinese      | English    | -                 | 3        | -         | 2        | -         | 97%      |
| FT-M      | English      | English    | -                 | 31       | -         | 29       | -         | 98%      |
| FT-M      | Chinese      | Chinese    | -                 | 82       | -         | 82       | -         | 96%      |
| FT-M      | Chinese      | English    | -                 | 41       | -         | 41       | -         | 95%      |




Figure 7: The number of unsafe responses by different LLMs edited DINM with respect to different editing steps and learning rates.




Figure 8: The number of unsafe responses by different LLMs edited FT-M with respect to different editing steps and learning rates.




Figure 9: The number of repetitions by different LLMs edited DINM with respect to different editing steps and learning rates.




Figure 10: The number of repetitions by different LLMs edited FT-M with respect to different editing steps and learning rates.




Figure 11: The number of unsafe/repetitive responses by different LLMs edited by DINM with respect to different learning rates at the best editing steps (right).




Figure 12: The number of unsafe/repetitive responses by different LLMs edited by FT-M with respect to different learning rates at the best editing steps (right).

## D Parameter-tuning for Systems Under Study

Table 3 reports the hyper-parameter settings of all models and KE methods we test in this study. These hyper-parameters are selected based on the number of failures as shown in Figure 13 and Figure 14 for DINM and FT-M respectively.

## E Effect of Editing Order on Compositional Robustness

To examine whether the observations in Figure 3 are sensitive to the order in which unsafe behaviours are edited, we repeat the compositional robustness experiment using four additional random permutations of the same 10 unsafe instances. Together with the editing order used in Figure 3, this results in five editing orders in total. Following the same evaluation protocol as in Section 6.1, after each cumulative edit, we evaluate the resulting model on all 10 unsafe instances.

Table 4 reports the results averaged over the five editing orders. Overall, the averaged results exhibit qualitatively similar trends to those observed in Figure 3, indicating that the main model-level observations are not specific to a particular editing sequence.

## F Complementary Results on Monolingual Detoxification

Figures 15–19 present the results of monolingual detoxification on Llama2-7B, Llama3-8B, Ministral-8B, Mistral-7B, and Qwen3-8B. Notably, the detoxification effects observed for English differ from those reported during hyperparameter tuning. This discrepancy arises because we evaluate the models on a different test set, namely data drawn from SAFEEDIT and LINGUASAFE. This further demonstrates that KE-based detoxification lacks optimization robustness, as its effectiveness can degrade substantially when applied to a different set of data.




Figure 15: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Llama2-7B.




Figure 16: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Llama3-8B.




Figure 17: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Ministral-8B.




Figure 18: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Mistral-7B.




Figure 13: The number of failures by different LLMs edited by DINM.




Figure 14: The number of failures by different LLMs edited by FT-M.

Table 3: Hyper-paramter settings of all models and KE methods we test in this study

| Model        | DINM          |              | FT-M          |              |
|--------------|---------------|--------------|---------------|--------------|
|              | Learning Rate | Editing Step | Learning Rate | Editing Step |
| Llama2-7B    | 1e-4          | 14           | 5e-5          | 5            |
| Llama3-8B    | 5e-5          | 6            | 5e-5          | 6            |
| Ministral-8B | 5e-5          | 12           | 1e-5          | 12           |
| Mistral-7B   | 5e-5          | 2            | 1e-5          | 10           |
| Qwen2-7B     | 5e-4          | 13           | 5e-5          | 6            |
| Qwen3-8B     | 5e-4          | 13           | 5e-5          | 14           |

| Model        | Initial      | Edit 1      | Edit 2      | Edit 3      | Edit 4      | Edit 5      | Edit 6      | Edit 7      | Edit 8      | Edit 9      | Edit 10     |
|--------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Llama2-7B    | 7.00 (0.00)  | 4.00 (1.40) | 3.60 (0.40) | 3.80 (0.20) | 4.00 (0.60) | 3.60 (0.60) | 3.20 (0.40) | 2.60 (0.60) | 1.60 (0.60) | 2.00 (0.20) | 1.80 (0.80) |
| Llama3-8B    | 9.00 (0.00)  | 1.40 (4.60) | 3.20 (2.80) | 1.80 (2.20) | 2.00 (5.20) | 2.20 (3.40) | 2.00 (5.40) | 1.60 (6.80) | 2.00 (5.80) | 0.80 (6.00) | 1.00 (7.20) |
| Ministral-8B | 4.00 (0.00)  | 6.80 (0.00) | 7.00 (0.00) | 6.40 (0.40) | 7.40 (0.20) | 7.40 (0.20) | 6.60 (0.20) | 6.80 (0.00) | 6.00 (0.40) | 6.20 (0.80) | 6.00 (0.60) |
| Mistral-7B   | 10.00 (0.00) | 8.60 (0.00) | 8.80 (0.00) | 9.20 (0.00) | 7.80 (0.00) | 7.80 (1.40) | 7.40 (1.20) | 7.40 (1.40) | 6.00 (2.00) | 6.40 (2.00) | 7.20 (1.40) |
| Qwen2.5-7B   | 10.00 (0.00) | 2.80 (0.00) | 1.40 (0.00) | 1.20 (0.00) | 1.40 (0.00) | 0.60 (0.00) | 0.80 (0.00) | 0.80 (0.00) | 0.00 (0.00) | 1.00 (0.00) | 0.60 (0.00) |
| Qwen3-8B     | 9.00 (0.00)  | 2.00 (4.00) | 1.00 (5.60) | 0.60 (4.80) | 0.60 (4.60) | 0.80 (4.00) | 1.40 (2.60) | 0.60 (4.20) | 0.20 (3.80) | 0.40 (4.60) | 0.80 (3.00) |

Table 4: Results of cumulative editing with DINM averaged over five random editing orders. Each entry reports the average number of unsafe responses among the 10 evaluation instances, with the average number of repetitive generations shown in parentheses. *Initial* denotes the model before cumulative editing.




Figure 19: Number of failures and performance on OOD inputs before and after monolingual detoxification across languages for Qwen3-8B.




Figure 21: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Llama3-8B.

## G Complementary Results on Cross-lingual Detoxification

Figures 20–24 present the results of monolingual detoxification on Llama2-7B, Llama3-8B, Ministral-8B, Mistral-7B, and Qwen3-8B.




Figure 20: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Llama2-7B.




Figure 22: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Ministral-8B.




Figure 23: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Mistral-7B.




Figure 24: Number of failures and performance on OOD inputs before and after cross-lingual detoxification across languages for Qwen3-8B.