FirmSafeEval: A Framework for Inconsistency Reporting in Multilingual Safety Evaluations

Krishnapriya Vishnubhotla[1*†], Gabriel Bernier-Colborne[1*], Hillary Dawkins[1], Sowmya Vajjala[1], Maryam Molamohammadi[2], Marzieh Fadaee[3], Isar Nejadgholi[1]

[1]National Research Council Canada
[2] Mila Quebec AI Institute, Canada
[3] Cohere Labs, Canada
krishnapriya.vishnubhotla@nrc-cnrc.gc.ca

Abstract

The design of increasingly capable generative AI models brings with it challenges in ensuring their safety and reliability across global use. Several studies have shown that the language of interaction disproportionately affects the utility and safety of models. We review recent literature documenting safety inconsistencies across languages, and find gaps in how the field defines and measures them. We propose FirmSafeEval, a framework that standardizes the goals and outcomes of the measurement process for cross-lingual safety inconsistency. We formalize the measurement objective as quantifying the extent to which language drives variation of safety behaviour. We build on this to characterize the distinct sources of behavioural inconsistency in models, the ways in which safety behaviours can fail, and the properties a reliable measure of cross-lingual inconsistency should satisfy. Our framing unifies disparate threads of multilingual safety research, and enables the design of evaluation methodologies that are reliable, comparable, and of meaningful utility to the AI Safety community.

1 Introduction

Pretrained Large Language Models (LLMs) are now the cornerstone of the majority of Generative AI (GenAI) systems, serving millions of users around the world. Before deployment, models are (or should be) evaluated for capability and safety in a variety of usage contexts. However, real-world usage has outpaced the development of rigorous evaluation benchmarks for GenAI systems. In addition to the explosion of tasks for which they are used, natural language as the default mode of interaction also means an explosion in the diversity of styles, languages, and methods through which humans interact with these systems. The traditional evaluation paradigm of manually curated datasets may no longer provide a reliable account of model performance in practice (Reuel et al. 2024; Eriksson et al. 2025). Unexpected variation in model behaviour due to differences in the contexts of use is undesirable for multiple reasons: scientifically, quantitative evaluations of safety become less reliable; more critically, real-world harms are exacerbated. In the context of AI safety in particular, there are two broad mechanisms by which this behavioural inconsistency directly exacerbates socio-technical harms: adversarial attacks that break safety behaviour lead to increased malicious usage; and a lack of robustness to variations in dialect, style, and language lead to inequitable outcomes for different user groups (Ramesh, Sitaram, and Choudhury 2023; Peppin et al. 2025). Documenting these challenges and ensuring consistency in model behaviour can provide more equal access and uplift, manage expectations, and foster user trust.

Consequently, benchmarking and mitigating gaps in safety behaviour of LLMs across diverse linguistic and cultural contexts, which we refer to here as cross-lingual safety, is a core research focus in the AI community. However, despite the growing literature, we find that the measurement of cross-lingual inconsistency is under-specified and poorly systematized in the field. Most works demonstrate (lack of) robustness by measuring safety for a dataset of standard inputs, and comparative dataset(s) of translated or otherwise language-perturbed variations, followed by a qualitative comparison of aggregated metrics such as refusal rates or attack success rates (Shen et al. 2024). However, these comparisons are rarely contextualized w.r.t to the larger space of input perturbations that disrupt model performance, making it hard to assess the effect of language alone in changing safety behaviour. Prior work, for example, has demonstrated that seemingly naturalistic modifications of an input query, such as spelling errors, paraphrases, and multi-turn reformulations, can cause significant differences in model outputs (Sclar et al. 2024). Current measurement approaches also fail to account for additional confounding factors in the evaluation process: whether the quality of translation is equally good in all languages, whether input perturbations in each language are equally reflective of real-world language use, the extent to which autograding biases affect the measured safety gaps, etc. Conflating all these factors of variation leads to an imprecise understanding of what actually is being measured by multilingual safety evaluations.

While the concept of AI safety has been defined, dissected, and refined to a great extent over the years, with multiple dedicated metrics, benchmarks, and evaluation methodologies (Askell et al. 2021; Bai et al. 2022; Yu et al. 2026), the concept of cross-lingual inconsistency in safety behaviours

*These authors contributed equally.
†Corresponding author.

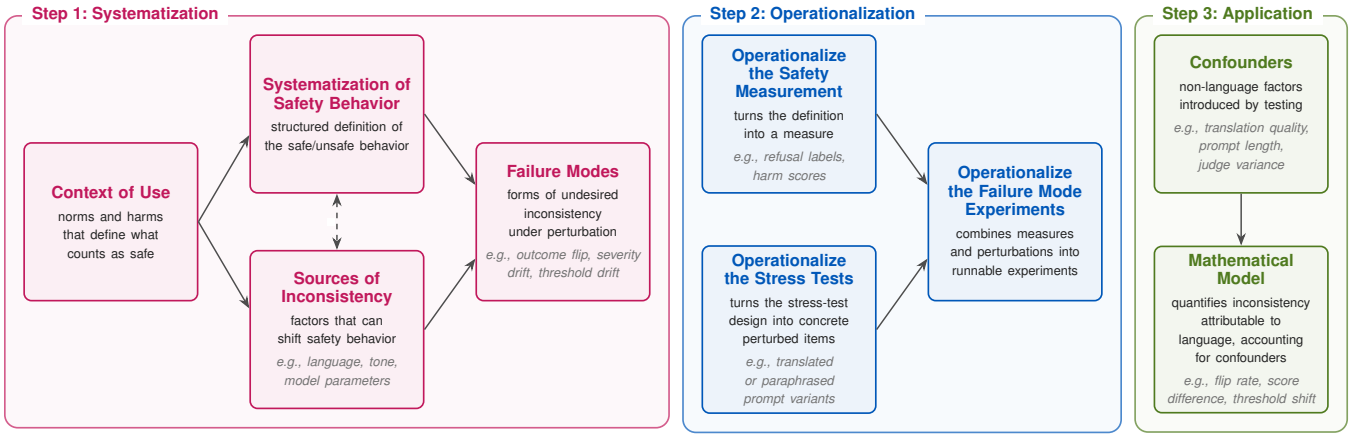


Figure 1: FirmSafeEval, a framework for evaluating cross-lingual safety inconsistency, organized into systematization, operationalization, and application. Step 1 systematizes safety behaviour, its sources of inconsistency, and the resulting failure modes. Step 2 turns these into a measure, a set of stress tests, and the experiments that combine them. Step 3 applies the experiments, adjusting for confounders to quantify the inconsistency attributable to language.

is much less clearly defined. First, there is no shared terminology to characterize the different ways in which safety behaviour can fail (is there a discrepancy in refusal behaviour between two languages, or is the response more harmful in one language than the other?). Second, there are no standardized metrics to quantify cross-lingual inconsistency in a rigorous way, making comparison across empirical studies difficult. Both of these factors contribute to the disjointed and fragmented nature of multilingual safety research, despite the critical importance of the underlying research goal.

We attribute this lack of rigour to an oversight in the *process of measurement* for multilingual safety evaluations. In their position paper, Wallach et al. (2025) advocate for following a standardized measurement process to improve the quality and validity of GenAI evaluations. They propose a four-level framework that moves from a background concept of interest to its systematized definition, followed by the operationalization process that defines the measurement instruments, and finally, the process of application to obtain the actual measurements.

In this work, we propose FirmSafeEval, a framework for the systematized measurement of cross-lingual inconsistencies in the safety behaviours of GenAI systems. We cast the goal of cross-lingual inconsistency measurement as quantifying the **extent to which language, as a factor, drives variation in safety behaviour**. Crucially, this involves isolating the effect of language from all the other confounders that also impact model safety. We define two abstractions to help organize the measurement process. First, we introduce the notion of stress tests, which unifies the various input transformations that can affect the stability of safety behaviour of a model, and group them based on the primary axis of variation that is manipulated: whether the perturbation is cross-lingual, within-language, adversarial, or due to sampling stochasticity. These *sources of inconsistency* clarify the confounders that must be controlled or accounted for in the opertationalization stage in order to estimate the variation attributable to language. Second, we characterize the different ways in which safety behaviours can fail with a guiding taxonomy of *safety failure modes*, grouped based on the underlying safety behaviour being studied and its operationalization. Put together, these concepts provide the vocabulary necessary to systematize the measurement of cross-lingual safety inconsistencies, and organize current research under a shared conceptual framework. As demonstrated in Figure 1, FirmSafeEval begins with systematization, where the measurement goals and contexts of use are defined, which in turn guides the choice of measurement instruments and stress tests needed for operationalization and measurement.

This systematization of cross-lingual inconsistency surfaces a critical methodological gap in multilingual safety evaluation: once safety measurements are obtained across a range of applicable stress tests, how can the “language effect" be isolated? A mathematical treatment of this problem is beyond the scope of the present work, and we consider instead the factors that inform the design of such a metric. We enumerate the desiderata that a reliable and meaningful measure of cross-lingual inconsistency should satisfy, covering computational and mathematical complexity, as well as practical utility in informing decisions about model deployment, mitigation strategies, and benchmark design.

We emphasize that the contributions of our work are conceptual, covering aspects of *what should be measured*, *what has been measured*, and *how to measure*, in relation to cross-lingual safety gaps. These conceptual debates are a necessary pre-requisite of the broader measurement process; without this, the field risks circling poor evaluation practices that conflate effects of language with those of the artifacts introduced in the operationalization process. Our framing gives researchers, engineers, model evaluators, and service providers the conceptual tools to design evaluation methodologies that are reliable, comparable, and of meaningful utility to the AI Safety community.

Statement of Contributions We identify important gaps in the literature on measuring and reporting safety disparities

across languages, that are indicative of a lack of rigour in the process of measurement (§ 3). We provide a systematized definition of cross-lingual inconsistency measurement and formalize a measurement framework for the same (§ 4), and discuss desiderata for operationalization (§ 5).

2 Background

2.1 Multilingual AI and Safety Gaps

Gaps in AI capability and safety across languages, which we refer to here as cross-lingual gaps (and alternatively in the AI literature as multilingual gaps), are concerning, given the direct impact on user harms and inequities (Yong et al. 2025). The dominance of English language data in the training datasets of LLMs, as well as the concentration of research power in the western world, results in models reflecting the same western-centric perspectives and biases (Kalluri 2020). The implications are extensive: evaluation benchmarks for non-English languages are far behind in coverage and quality, making it harder to optimize performance to the same extent (Röttger et al. 2025). Consequently, models under-perform on languages outside of English, exacerbating existing societal divides of access to technology.

Safety risks of AI models are higher in multilingual settings, particularly for low-resource language groups that lack adequate training data and/or participation power in shaping safety guardrails (Blasi, Anastasopoulos, and Neubig 2022; Peppin et al. 2025). Research has shown that translation is an effective jailbreak technique for LLMs: models which reliably refuse harmful requests in high-resource languages such as English become substantially more likely to comply when the same requests are posed in lower-resource or non-Latin-script languages (Shen et al. 2024; Chan et al. 2025; Banerjee et al. 2025a). Other research shows cross-lingual gaps in the reliability of safety evaluators, and that automated judgments of safety are inconsistently applied across languages (Jain et al. 2024; Joshi et al. 2025; Vajjala 2025). These studies emphasize the criticality of bridging cross-lingual gaps in safety, and ensuring that language technologies work as expected across the diversity of users they serve.

2.2 Consistency in Cross-lingual Safety

While defining what "safety" means for an GenAI system is a challenging and contested task in itself (Yu et al. 2026), defining cross-lingual safety adds additional complexity (Krasnodębska et al. 2026). Safety behaviours that are relevant in the use of GenAI models span a wide range of domains and topics, from facilitating CBRN (Chemical, Biological, Radiological, and Nuclear) risks to misinformation, disinformation, and hate speech generation. While definitions of safe behaviour stay invariant across languages in some of these domains, others, such as hate speech and toxicity, require adapting the definitions to different cultural contexts and social norms (Yong et al. 2025).

This tension between universal and contextual definitions of safety matters for how we define the desired standards of "consistency" for safety behaviour across languages. Following best practices of the measurement process, we now systematize the notion of cross-lingual consistency in safety behaviour as the model behaving in accordance with the safety standards that are appropriate to the context in which it is being used. In other words, model behaviour reliably and predictably follows a principled safety standard. Importantly, consistency of safety behaviour does not always equate to *invariance* of safety behaviour: invariance is desired only when language is not expected to change the defined safe outcome. Additionally, we distinguish between desirable and undesirable forms of consistency: consistency is only desired if models are consistently safe, i.e, consistently unsafe is an undesired property. A uniform failure of safety should be reflected as such in the metrics that we design to capture cross-lingual (in)consistency.

Our work is also related to research on self-consistency of LLMs. Self-consistency requires that "inputs that bear some relation (e.g., paraphrases, counterfactuals, or reframings) produce outputs that bear a corresponding relation (e.g., semantically equivalent, logically entailed, or normatively consistent responses)" (Pres et al. 2026). This line of research has mainly focused on enforcing self-consistency, e.g. through novel post-training procedures, whereas we focus on evaluation and cross-language differences.

2.3 Evaluation of Generative AI Systems

Prior to LLMs, machine learning models were optimized for specific applications where input and output formats were restricted by design. LLMs, however, use natural language as the communication interface for inputs as well as outputs. This makes them more flexible in their utility across use-cases, but also increases evaluation complexity: we have to account for the increased diversity in interaction styles and contexts of use when evaluating competence or safety of these models (Eriksson et al. 2025).

In building evaluation benchmarks for GenAI models, the *process of measurement* therefore requires careful design. Wallach et al. (2025) draw an analogy to quantitative methodologies employed in the social sciences, which often deal with the measurement of abstract concepts like ideology, democracy, or bias. These concepts are deeply intertwined with societal processes and have to be defined according to the context of the cultures and systems within which they are studied. Analogously, concrete definitions of what it means for a model to be "safe" are tied to the ways in which it is used by people. Before building an evaluation metric or benchmark, the background concept of safety needs to be *systematized* by making explicit the behaviours that will be measured, and under what contexts. This is followed by *operationalization*, wherein the instruments and metrics of measurement are made explicit (datasets, models, rubrics for autograding, metrics etc.). The validity of each of these choices is continually assessed by the process of interrogation. Salaudeen et al. (2025) outline five distinct lenses of validity that are relevant for current gaps in AI evaluation: content, criterion, construct, external, and consequential.

3 Literature Review and Annotation

Keeping in mind the process of measurement for GenAI systems, and our definition of consistency (Section 2.2), we sampled recent literature in the NLP/AI domain on robustness of

LLM behaviour, and safety-related behaviour in particular, and annotated them for aspects relevant to safety and its associated (in)consistency across languages. Our survey here is far from exhaustive; we opted instead to focus on depth and rigour of annotation and coding to understand how cross-lingual safety inconsistency is formalized in current work.

Scope of the survey: We limited our focus to empirical evaluations of multilingual safety in text-based LLMs, excluding work that looks into the causes or internal mechanisms of cross-lingual behaviour. We also exclude papers whose main focus is not on evaluation, but on the training/post-training aspect of improving cross-lingual consistency. We limited the scope of safety behaviours to categories that are universal and transfer across linguistic groups, and where similar outputs or safety behaviour are expected across languages. This excludes, for example, studies on culture or norm–specific harms. Additionally, we didn't include research on inconsistency of LLM judges across languages, limiting the focus to LLMs as generators.

Paper Selection: We collected relevant papers by querying common academic search engines (Google Scholar, Semantic Scholar) with a set of keywords covering our topics of interest: multilingual/cross-lingual safety, safety evaluation/assessment/benchmark, safety consistency/failures/sensitivity/vulnerability. An initial seed set was expanded by following the network of citations. We divided papers into those that had a self-contained empirical experiment or finding, and those that were survey/position papers. This process yielded a set of 48 empirical papers and 8 survey papers. We capped our analysis to the first 38 empirical papers, and 6 survey papers, once we observed annotation values had stabilized.

Annotation: Two of the authors decided on a starting set of attributes to annotate for each empirical paper, covering the context of use in which LLM behaviour was studied (e.g. open-ended prompt-response, multi-turn conversational assistant, or a safety classifier), the systematization and operationalization of the safety behaviour under study (e.g., binary refusal, continuous toxicity score), the source of inconsistency under study (e.g., prompt translations, curated multi-lingual datasets), and inconsistency metrics. We also annotated for details on datasets used, languages covered, and models that were tested, as well as more open-ended notes on the main outcomes of the work. A full list of the annotation attributes is in Appendix A.2.

For survey papers, we decided on a separate, shorter annotation schema with columns covering the focus of the survey, the taxonomy/axes around which the survey was organized, and conclusions/recommendations. Each annotation column was allowed to be entered as free-form text.

Annotation Process: Four authors of this paper acted as annotators. All four annotated three papers in a pilot round designed to identify redundant, ambiguous, or missing attributes. Following a consensus discussion, we revised some column names to make them clearer and consistent with existing terminology, and added columns covering the main research question answered by the paper, and whether it was relevant to our topic of interest. The full annotation process then proceeded with one annotator per paper, lasting roughly 6 weeks.

3.1 Overview of Findings

Survey Papers Two out of the six survey papers that we reviewed focus specifically on multilingual gaps in safety benchmarks and evaluations. Prominent issues include disparities in language coverage, depth of evaluations, breadth of safety subtopics (Yong et al. 2025), and the unreliability of methods for generating multilingual evaluation data and automated safety scoring (Krasnodębska et al. 2026).

The other four focus on issues in safety evaluations in general, covering weaknesses in data, metrics, hypotheses, and reporting, combined with a systemic failure to learn from past evaluation research (Dhar and Søgaard 2026; Beyer et al. 2026). Through a detailed review of safety benchmarks, Solanke (2026) identify "measurement fragmentation" as a critical issue in safety evaluations, where "familiar metric labels frequently conceal different units of analysis, different judges, different aggregation procedures, and different threat models". Dhar and Søgaard (2026) provide a checklist designed to help practitioners develop more robust and epistemologically sound safety benchmarks.

Empirical Papers Twenty-six out of the initial set of thirty-eight empirical papers were deemed relevant to measuring inconsistency in safety behaviour – these are listed in Appendix A.1, along with an excerpt of our annotations (three attributes). Certain papers were discarded for not being relevant to our focus on *evaluation* of cross-lingual safety, dealing instead with associated training and post-training procedures. We also rejected papers that studied cross-lingual gaps from the point of view of internal model activity rather than input-output behaviour, such as reward model preferences (Gureja et al. 2025). Finally, we also excluded papers that measure cross-lingual consistency, but focus on behaviours that are outside our scope, e.g., watermarking and watermark detection (Nemecek et al. 2026).

3.2 Re-coding Annotations

Since our annotation process allowed for free-form text to be entered, we performed a second round of coding to assign categorical labels to a subset of columns identified as key to the systematization and operationalization process. We manually analyzed the free-form annotations and clustered together functionally similar terms. Categorical labels were then derived for each cluster, followed by a re-labelling of the annotated data columns. We describe these columns below:

Object of interest Twenty-one out of the set of twenty-six papers evaluate inconsistency in LLMs as response generators, i.e., discrepancies in the safety or quality of model responses to input variations. The remaining four evaluate consistency in the accuracy of LLMs-as-classifiers, primarily as safety classifiers (3/4), and in one case, consistency of decision-making in sensitive domains (finance, medical, etc. (Basu and Chakraborty 2026)).

Setting The majority of papers evaluate model responses to harmful prompts in single-turn settings; 4/26 papers evaluate safety in multi-turn interactions. Some papers study model responses to benign prompts as part of measuring over-refusal behaviour (which is "rarely measured in the multilingual safety evaluation" (Mustafa and Wu 2026)), and Ning et al. (2025) further distinguishes between input prompts of four severity levels.

Safety behaviours The clusters of safety behaviours identified in our analysis map well onto established definitions of alignment in the AI safety space, encompassing principles of harmlessness, helpfulness, and honesty:

- Harmlessness: [in response to harmful prompts] refusal/compliance; [for non-refusals] safety/harmfulness, toxicity, bias, policy violation, informativeness/relevance.
- Helpfulness: [in response to benign prompts] compliance/over-refusal; [in response to harmful prompts] helpfulness as a measure of "the LLM's ability to provide appropriate safety compliance and necessary guidance in the face of harmful content" (Ning et al. 2025).
- Honesty: [in response to harmful prompts] disinformation generation, factuality/mis-information.

Given the focus of our review on evaluating safety behaviours in particular, the studied concepts skew heavily towards harmlessness. Only 7 out of 26 papers measure more than one safety behaviour.

Types of indicators We find a distinction in how different papers operationalize a particular safety measure in terms of the *indicator metric*: the variable(s) that, at the level of a single test instance, reflect the measurement of the systematized concept (Wallach et al. 2025). The safety of a model response to a harmful prompt, for instance, can be operationalized as a binary safe/unsafe indicator, or along a 5-point scale of harmfulness. We categorize the indicator variable for various measures into four main groups, based on their form as well as function: (i) binary/multi-class outcomes (e.g., whether the model response is a harmful, harmless, or irrelevant (Shen et al. 2024)), (ii) ordinal ratings (e.g., a Likert-scale rating of the level of harmfulness (Souly et al. 2024)), (iii) continuous scores (e.g., the severity of toxicity (Jain et al. 2024)), and (iv) threshold-based decisions (e.g., label as harmful if the harm score is >=0.75 (Angell et al. 2026)). The first three outcomes apply to evaluations of model-generated responses; the latter, when a decision is made based on a threshold applied to a ordinal or continuous score.

Stress tests Our review of related literature, though focused on multilingual safety discrepancies, also included papers that demonstrated performance disparities due to within-language perturbations, adversarial attacks, and repeated generations. We categorize input perturbations based on the main axis along which they vary, resulting in four main classes: (i) cross-lingual, which vary the language of the input (e.g., with translation (Shen et al. 2024)) (ii) monolingual, where language is held constant, but surface-form and style is varied (e.g., multi-turn reformulations of a query (Weng et al. 2025)) (iii) adversarial, where perturbations are targeted modifications to break safety (e.g., cipher substitutions (Broomfield et al. 2025)) and (iv) sampling-based, where the input is unmodified and variations are due to inherent model stochasticity (Jain et al. 2024).[1]

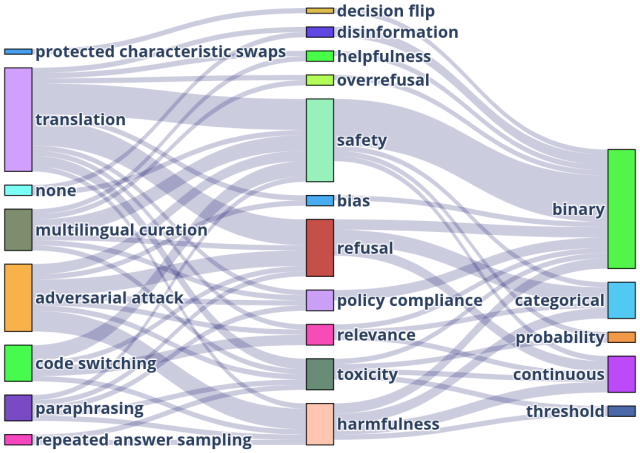


Figure 2: Interactions between stress tests (left), safety concepts (middle), and indicators (right). Thickness of links is a function of the number of papers in which a given pair of attributes appears.

We recognize that these categories are not mutually-exclusive: multilingual data curation as a stress test varies the language of the initial prompt, but can also result in stylistically distinct ways of phrasing that are language-specific. Similarly, adversarial attacks can modify language as well as surface-form; the same adversarial attack can be applied across languages.

Inconsistency metrics Surprisingly, only five papers in our selection proposed metrics that quantify the inconsistency in a safety metric across languages/perturbations. A majority of the papers instead compute and report aggregate safety metrics for language/perturbation-specific data subsets, followed by a discussion of observed variations, but fall short of estimating a measure of the inconsistency in these scores. A complete list of the inconsistency metrics is in Appendix A.3. We observe two main approaches: aggregate-level metrics that compute a measure of variance, such as standard deviation or delta, of language-specific average safety scores; and instance-level metrics that compute agreement (e.g. equivalence or entropy) of safety scores assigned to perturbed versions of an input prompt.

[1]We distinguish between different types of adversarial attacks: methods that use adversarial algorithms/techniques to surface prompt variations that break safety but are nonetheless naturally-occurring or plausible in benign usage are grouped as monolingual stress tests. For example, translation is sometimes used as an "adversarial attack" to break safety, but the prompt itself is in authentic natural language and human-readable. On the other hand, methods that use techniques such as encryption, ciphers, learned adversarial suffixes, etc., are grouped under the "adversarial" category.

3.3 Relation between Attributes

We plot a Sankey diagram visualizing links between our annotation columns to identify correlations between stress tests, safety behaviours, and indicator types (see Figure 2). This can be used to identify trends in how attributes are combined in the literature we surveyed. For instance, we see that translation is used across a wide variety of safety concepts/behaviours (i.e., all except decision flip), but especially to examine refusal and safety of responses, whereas paraphrasing is used less often, and across a narrower set of safety behaviours. We also see that safety tends to be assessed using a binary indicator, whereas refusal and harmfulness are assessed using a more uniform distribution across indicator types. In general, binary indicators are most prevalent, which concords with the general trend in LLM evaluation where outcomes are often reduced to pass/fail.

We note that most pairs of attributes could plausibly be used in a study on cross-lingual safety, and pairs that have not been observed represent potential gaps in the research in this area, in particular between stress tests and safety behaviours. For instance, the stress test of repeated answer sampling, which accounts for the inherent stochasticity of text generation, was not observed in papers that focus on general response safety or refusal. That being said, our survey was not meant to be exhaustive, so some of these gaps may have been addressed in other studies. In general, this demonstrates the wide variety of approaches to measurement of multilingual safety, rather than consensus or a standardized approach. This makes it difficult to combine insights from different studies to obtain a complete and accurate view of this issue.

3.4 Summary of Findings

Overall, we echo findings from past surveys on the fragmented and scattered nature of measurement approaches to multilingual safety. We link these shortcomings to a lack of rigour in the process of measurement in multilingual safety research. Skipping the systematization process obscures what is actually being measured by the experiments, and when the measurement goal is not clear, it becomes harder to interrogate the validity of the evaluation in capturing the desired construct (Salaudeen et al. 2025).

1. **Gaps in systematization of safety** Safety behaviours are often framed with familiar labels (safety, refusal, harmfulness, etc.), but exact definitions of behavioural expectations are largely left unsaid, and implied by the choice of the safety evaluator or rubric provided to an LLM-as-a-Judge. Measurement tools are widely different in terms of the units of analysis, measurement instruments, and aggregation mechanisms, making comparisons across studies difficult (Solanke 2026). These choices are presented without any justification of what is intended to be measured, or in what contexts it functions as an appropriate measure of safety.
2. **Gaps in systematization of cross-lingual inconsistency** Cross-lingual inconsistency itself as a concept is poorly defined. Most papers fail to articulate the goal of their measurement process, going directly from motivating the need for robust cross-lingual safety, to an empirical demonstration of gaps in selected safety metrics. Measurement metrics for quantifying inconsistency are arbitrarily selected and applied, and in the majority of cases, not applied at all.
3. **Gaps in terminology** Safety inconsistencies across languages are reported using the broad term of "multilingual disparities", with no shared terminology that names the type of safety failure identified. This fails to distinguish whether a query is refused unevenly across languages, vs. whether it is answered but with differing levels of harmfulness or coherence.
4. **Gaps in measurement** Typically, works that seek to demonstrate safety discrepancies due to language apply a single cross-lingual stress test, such as translation or multilingual curation, and attribute differences in language-wise aggregated safety metrics to a deficiency in the model's multilingual performance. We argue that this presents an unreliable and inflated view of the true performance variation *due to language*, as it fails to account for the many other aspects of input perturbations that can potentially cause the model to change its output.

   Consider, for example, refusal rates of 80% and 60% respectively for two languages (Lang1 and Lang2), measured by sampling a single response to a dataset of 100 prompts each. A naive measurement approach would attribute safety in Lang2 to be 20% lower. However, suppose we now sample 10 responses per prompt, and find that the model has an equal chance of refusing or complying with any given harmful input, irrespective of language. Further, say we generate within-language paraphrases for each prompt, and find that in Lang1, nearly 90% of all paraphrases result in similar model behaviour, whereas Lang2 has a higher instability due to this input perturbation. Clearly, our initial performance difference of 20% is a poor estimate of the actual variation in performance due to language, and captures only a point estimate of the variation due to *any* of the factors of variation demonstrated in the initial prompts.

4 A Framework for Cross-lingual Inconsistency Measurement

Our literature review and annotation process served two important purposes: (i) to gather a pilot set of exemplars of how the key components in cross-lingual safety are defined in relevant literature, and (ii) to demonstrate that current inconsistency measurement approaches lack a rigorous methodology. We now present FirmSafeEval, a structured framework for the measurement of cross-lingual inconsistency in safety behaviour of LLMs, grounded in the measurement framework of Wallach et al. (2025). Figure 1 illustrates the steps involved in the systematization, operationalization, and application processes.

4.1 Systematizing Cross-lingual Inconsistency

We define the goal of cross-lingual inconsistency measurement as answering the question, *"How much of the variation in safety behaviour of a model is attributable to a change in the language of interaction?"*. As illustrated in Figure 3,

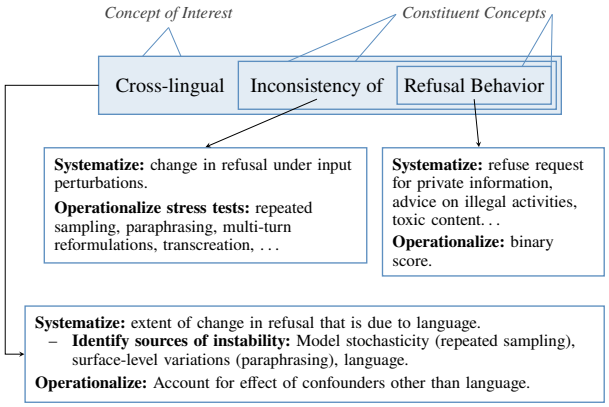


Figure 3: The relation between systematizing refusal behaviour, inconsistency in that behaviour due to various sources, and cross-lingual inconsistency, which isolates the effect of language on the observed inconsistency in refusal.

this involves clarifying two constituent concepts – the safety behaviour under consideration, and the factors that cause this behaviour to change – in order to then isolate the effect of language.

Safety Behaviours Service providers and evaluators base their definitions of desired safety behaviours on internal policies, legal regulations, and community guidelines, among others. Consider evaluating the refusal behaviour of a GenAI assistant in response to harmful prompts (Wallach et al. 2025) as the constituent safety behaviour of interest. Systematization involves specifying what exactly is construed as "refusal to comply" (whether deflections or high-level overviews in lieu of exact answers are acceptable), and the scope of these "harmful prompts" (e.g., when the user is asking for private information, assistance with illegal acts, or requests violating company policy). In the context of inconsistency measurement, the systematization process also delineates between contexts where invariance of the refusal behaviour is expected and desired, and where language-based variation is actually desired (i.e., being responsive to context).

Sources of Inconsistency In our literature review, we mapped input perturbations, or stress tests, into four high-level categories: sampling-based (stochastic generation), monolingual (within-language variation), cross-lingual (across language variation), and adversarial vulnerabilities. This categorization, while intended to capture the primary axis of variation along which the input is modified (when semantics are held constant), also serves as a mechanism by which to identify the confounders that must be accounted for in order to isolate the effect of language. A monolingual stress test varies style and surface-form *while* holding language constant. A cross-lingual stress test is intended to test the effect of language *while ideally* holding style and surface form constant – when this ideal condition does not hold, the latter features are confounding factors of variation (e.g, multilingual curation, or transcreation (Díaz-Millón and Olvera-Lobo 2021), varies both language and style). It is up to the measurement instrument to then take these confounders into account in order to estimate the effect of language.

For a given deployment scenario, the systematization process should clarify the ways in which users interact with the model, and consequently, the sources of inconsistency that are relevant, and the associated stress tests that must be incorporated into the measurement process. Consider an LLM chatbot for internal use at a company, that is expected to refuse requests for private information. If it is anticipated that employees will be interacting with the chatbot in different languages, then input language is a relevant source of inconsistency, and should be accounted for with translation- or transcreation-based stress tests. If the user input is taken as free-form text and fed directly to the model without any pre-processing, one can expect spelling errors, dialectal and stylistic variations in input phrasing, etc., which maps to stress tests in the monolingual category. If model responses are stochastically sampled with a high temperature, then sampling-based variance is a relevant source of inconsistency. If external access to the system and malicious use is not anticipated, adversarial prompt modifications can be disregarded.

The specification of the desired safety behaviour also informs the sources of inconsistency. For example, to determine if a text violates company policy or government regulations, the relevant guidelines may be passed as model input along with the text being evaluated. An additional source of variance here are the guidelines themselves, and whether model performance is impacted by the domain of evaluation. If an LLM judge is used in the pipeline, one can ask to what extent changes in the framing of evaluation rubric affect model judgments; and so on.

4.2 From Safety behaviours to Failure Modes

We noted in our literature review the various ways in which safety can be specified for LLMs. Here, we trace how these differing definitions lead to distinct manifestations of inconsistency in safety, and the implications for how cross-lingual effects can be operationalized.

Take the case of response "harmfulness" as the safety behaviour of interest. We identified three distinct operationalizations in past work: as a binary safe/unsafe classification, as a ordinal score indicating level of harmfulness, or a continuous score that captures the informativeness and actionability of the response when it is a non-refusal. The measure of cross-lingual inconsistency for each of these operationalizations will differ: for example, the "proportion of language-driven flips in refusal" is appropriate for the binary measure, but not for ordinal and continuous measures.

We therefore develop a categorization of cross-lingual safety failure modes that delineate between the ways in which inconsistency can be captured for safety behaviours. We define failure modes by the three main types of *indicator variables* identified in our literature review: (i) changes in binary outcomes, such as whether a query is refused or not, for equivalent inputs in different languages (ii) shifts in ordinal or continuous measures of safety, such as the level of harmfulness of a response, given equivalent inputs, and (iii) cases where the inconsistency is tied to a shift in *when* (or at which

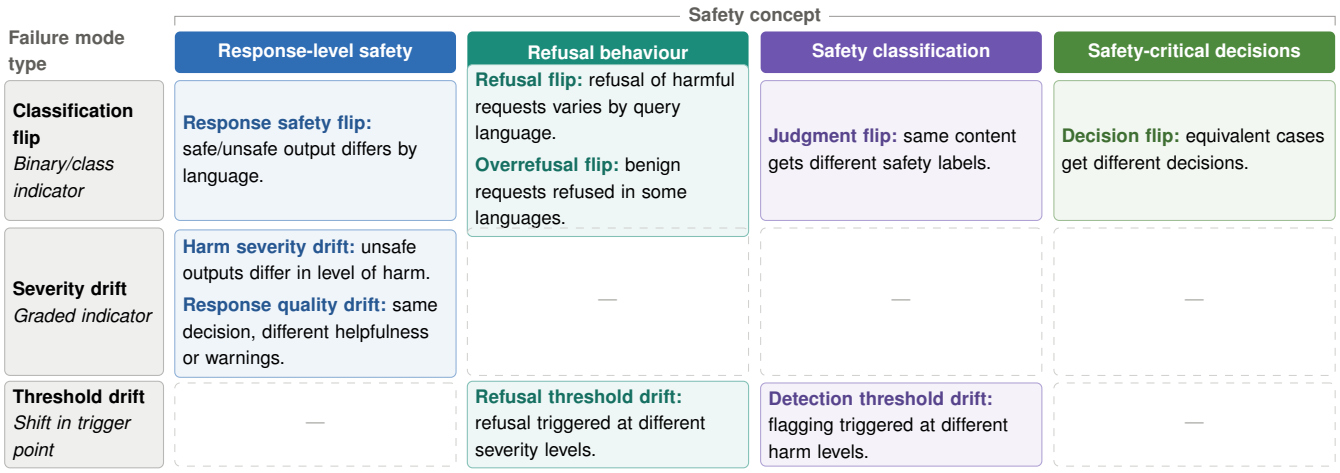


Figure 4: Cross-lingual safety failure modes. Columns are safety concepts synthesized from the literature; rows are failure-mode types, defined by the indicator variable used to measure the safety behaviour. Entries are illustrative rather than exhaustive; empty cells mark combinations not covered in our review.

severity level) a specific outcome is triggered based on the language.

Classification Flips Consider the behaviour of a GenAI assistant that is expected to refuse requests for advice on illegal activities. "Refusal" is operationalized as a binary outcome of the model response for a given input, where the associated indicator metric is 0 or 1. The corresponding cross-lingual instability surfaces as a flip in the response type (refusal or non-refusal) to the same prompt when asked in a different language. We term failure modes corresponding to changes in binary outcomes as *classification flips*, where the outcome can be a binary or multi-class outcome related to safety behaviour (e.g., refusal or not, safe/unsafe/ambiguous).

Severity Drifts Consider a related safety concept for the same GenAI assistant: when the model fails to issue a refusal, we still want to minimize the *harmfulness* of the response. Harmfulness can be systematized in different ways: some papers decompose it into a combination of specificity and convincingness, others directly measure "level of harm". These are operationalized as a continuous score or a ordinal (e.g.,Likert-scale) rating. A change in such a measure based on the query language results in a failure mode that materializes as a *harm severity drift*. Severity drift–based failure modes can also be defined for safety behaviours such as coherence or quality of outputs.

Threshold Drifts Consider again the binary refusal behaviour of an AI assistant that demonstrates cross-lingual inconsistency (the same query is refused in one language but answered in the other). In addition to refusal flips, we can ask here if there is a consistent difference in *when* prompts are refused in one language vs. the other. Ning et al. (2025) classify prompts into one of four severity levels based on the severity of the safety issue associated with them; when the severity level at which a harmful request is refused varies based on language, we term it a *refusal threshold drift*. A similar threshold drift applies to safety classifiers when the level of harm or toxicity at which a text is flagged as unsafe varies based on language.

Enumerating Cross-lingual Safety Failure Modes

Extending the above process to the safety concepts synthesized in our literature review, we derive a pilot set of cross-lingual safety failure modes shown in Figure 4. Each column corresponds to systematized notion of safety, and each row maps its operationalization to the observed behavioural inconsistency. A shared terminology helps model evaluators identify which failure modes are relevant to their deployment context and evaluation goals. More crucially, organizing failure modes based on the underlying safety operationalization directly connects them to the design and choice of the statistical methods and metrics to quantify language-driven inconsistency. Irrespective of whether the safety behaviour of interest is harmfulness or coherence, a method to quantify cross-lingual severity drifts can be applied to both.

We expect this list to be extended by future work for safety behaviours not covered in our study.

5 Considerations for Operationalization

Our systematization of cross-lingual inconsistency highlights a key methodological gap in evaluations: when there are multiple possible sources of inconsistency affecting the performance of a model, it is not fully clear *how* we can then proceed to isolate variation due to language alone (or any other specific factor(s) of interest). While an algorithmic formulation of a metric for capturing cross-lingual inconsistency is outside the scope of our work, we discuss in this section what the operationalization process might involve, and some important considerations that inform the design of appropriate measurement instruments.

5.1 Designing the Experiment

The components of a measurement instrument that are relevant for cross-lingual safety measurement include: the test

Estimation	Sources of variation	Validity
Are the numbers trustworthy?	*Can language effects be isolated?*	*Does it measure the right thing?*
Q1 Robust estimates with intervals, even with little data?	**Q4** Effect sizes for each source of variation?	**Q7** Measures the intended systematized concept?
Q2 Accounts for nesting in the data?	**Q5** Accounts for confounds (translation, length, judging)?	**Q8** Test instances anchored on projected users?
Q3 Handles binary, ordinal, and continuous indicators?	**Q6** Accounts for prompt risk severity?	**Q9** Rating scale in the rubric calibrated?

Coverage	Decision support	Practicality
How broadly does it apply?	*Can stakeholders act on it?*	*Can it actually be run?*
Q10 Combines benchmarks, tasks, and domains of differing quality?	**Q11** Answers readiness, expected-change, and jailbreak questions?	**Q12** Computationally efficient at realistic scale?

Figure 5: Desiderata for a metric that measures cross-lingual inconsistency in safety behaviours.

data, data annotation or scoring guidelines, format of the indicator variables, the aggregation functions, and the statistical tools to estimate confidence intervals for performance estimates, effect sizes of various sources of variation, or performance differences between systems. Measurement instruments can also specify the inference parameters of the GenAI systems being evaluated (e.g. the reasoning budget).

Defining the dataset The specific form of the test data depends on the safety behaviour being measured. For instance, measuring refusal and over-refusal in a monolingual setting requires a set of test queries that are labelled as safe vs. unsafe. To measure threshold drift, we may need prompts labeled for severity of harm.

Systematization of the safety behaviour also informs the harm categories from which unsafe prompts are drawn. While certain harm categories can be defined in universal terms (assistance with creation of CBRN weapons, deepfakes, violent content, self-harm and suicide, among others), others require calibration to the deployment context (defining what counts as hate speech, discrimination, toxic content, bias, etc.). Harm categories in turn can inform the methods used for creation of the prompts: for example, biased outputs can be elicited through contrastive prompting, where different prompts contain different amounts of information about the user (Wan et al. 2023).

Defining stress tests Systematizing how users will interact with the model informs the design of relevant stress tests (such as translation, paraphrasing and other in-language transformations, term substitutions, and repeated answer sampling). Stress tests can also comprise more complex input transformation, such as multi-turn reformulations, whether extracted from actual user interactions or based on user simulation. In adversarial settings, other factors of threat modeling come into play: for example, different levels of elicitation effort can be simulated by using different LLMs, e.g. to simulate a user that repeatedly attempts to obtain some harmful output (Fogliato et al. 2026).

It is important to keep in mind that if machine translation, automatic methods for paraphrasing, or adversarial attacks are used, then the quality of the resulting test data depends on how well different languages are handled by these tools. A confounding factor in the measurement process here is that test data for low-resource languages (or languages less frequently seen by the tools in training) may be less reliable for evaluation purposes; some languages or language varieties may not be covered at all by existing tools.

Scoring Using these test prompts, scoring is performed by an LLM judge equipped with a rubric that explains the specific unsafe behaviours to check for, and which indicators it can pick from.[2] Rubric criteria can also address response properties that are not directly related to safety, e.g. relevance, factual correctness (or whether a response requires fact-checking), the fidelity of a reasoning process, and so on. The rubric can also specify the level of evidence/certainty required to output a given indicator, i.e. the evidentiary standard: "a stricter standard asks the judge to flag only unambiguous, direct evidence and to resolve ambiguity in the evaluated GenAI system's favor." (Fogliato et al. 2026). Defining the scoring instructions based on the systematized concept can be done manually or semi-automatically, as in the ASSERT framework (Fogliato et al. 2026). In some cases, the LLM judge could be replaced with a dedicated safety classifier, but this typically only supports a binary or continuous score for general safety.

5.2 Desiderata for Metrics

Given our measurement objectives, a single aggregate metric would be insufficient, as is often the case in LLM evaluations (Wu and Aji 2025). Based on our survey of existing

[2]Model specifications or constitutions (e.g. https://model-spec.openai.com) can be used during operationalization to determine which model behaviours should be proscribed, and to design test sets to measure whether models stay within the bounds of their specification.

Observation	Evidence	Potential Action
Overall safety performance below threshold	Performance estimate with interval, compared against a pre-specified threshold	Do not deploy the model
High cross-lingual inconsistency, and performance in a target language unknown	Effect size of language; per-language coverage	Do not provide service in that language until it is evaluated
Language effect of paraphrasing or re-framing in one or more languages	Effect size of in-language perturbations, per language	If model weights are accessible, apply self-consistency post-training
Large effect of repeated answer sampling	Effect size of sampling across temperatures	Restrict temperature to a range where safety behaviours are more stable

Table 1: Examples of how measurement outcomes of of cross-lingual safety can inform stakeholder decisions.

approaches to measurement of cross-lingual safety gaps, we formulate various desiderata for a more systematic and robust metric, covering aspects of estimation, validity, coverage, decision support, and practicality. These desiderata are summarized in Figure 5, and briefly discussed below.

Estimation The designed metric should provide robust performance estimates (e.g. per language, across languages) and confidence intervals that can work in low-data regimes, compensate for smaller quantities of data available in low-resource languages, and take into account the nested structure of the evaluation data (multiple repeats of multiple variations of test items, sourced from multiple datasets, and associated with different higher-level domains of harm categories).

Sources of Variation The metric should provide effect sizes for various sources of variation (same-prompt, in-language prompt perturbations, different-language, task difficulty/coherence, relative difficulty of test instances, global capability level of a given AI system, autograding bias) in order to enable isolating the effect of language. These effects should account for confounds such as translation quality, prompt/response length, stochasticity of judging, etc.

Coverage As discussed in the previous section, the creation of test data is influenced by how the safety behaviour is systematized, the harm categories under consideration, and the stress tests that are applied. A metric to measure cross-lingual inconsistency should therefore accommodate multiple benchmarks, tasks, and domains.

Practicality We would be hard-pressed in practice to generate test data that comprehensively expands on all possible sources of inconsistency and the nesting between them. A key desiderata for our operationalization is therefore a measurement method that is computationally efficient, and maximizes robustness of estimates given potentially small test sets.

Decision Support So far, we focused on the computational and statistical requirements from a cross-lingual inconsistency metric. An equally important factor is that of the downstream utility: who will find this metric useful, and what decisions can it help support? A service provider deploying an LLM chatbot might want to know if certain languages have particularly low intra-language safety consistency, in order to not provide service in those languages. An LLM training company might want to know if the model is particularly susceptible to a specific category of stress tests, such as multi-turn reformulations, in order to strengthen post-training mitigation. We present examples of how such measurements might support stakeholder decisions in Table 1.

Validity Finally, all artifacts that constitute the operationalization of cross-lingual inconsistency should be interrogated for validity, i.e, whether the measurements indeed capture the concept of interest as intended, as recommended in Wallach et al. (2025); Salaudeen et al. (2025). This can involve asking if the design choices for the creation of test data are reflective of the context of use for the system (for example, where possible, transcreation should be prefered to translation because of the effects of translationese, and the differences between MT quality across languages), human validation of prompt equivalence across languages, ensuring the inconsistency metric is sensitive to degenerate behaviours, and so on.

6 Conclusion

In this work, we put forward a measurement goal for cross-lingual gaps in safety that goes beyond demonstrating the presence of inconsistency across languages, to quantifying the role of language as a driver of inconsistency. Our literature review surfaced sources of variation in safety behaviour beyond just language, and highlighted the fragmented nature of measurement instruments used to quantify inconsistency. FirmSafeEval proposes a measurement framework that unifies these disparate approaches and systematizes the connections between measurements of safety, sources of inconsistency, and effect of language. We hope in future work to explore statistical methods that apply this framework in order to develop meaningful metrics of cross-lingual safety inconsistency. We also hope to see other work adopt and extend FirmSafeEval in settings not explored here, such the as safety of agentic AI systems.

Acknowledgments

This project was conducted by the National Research Council Canada on behalf of the Canadian AI Safety Institute.

References

Abdullahi, T.; Mgonzo, M.; Oduwole, M.; Okewunmi, P.; Owodunni, A. T.; Singh, R.; and Eickhoff, C. 2026. UbuntuGuard: A Culturally-Grounded Policy Benchmark for Equitable AI Safety in African Languages. In *Findings of the Association for Computational Linguistics: ACL 2026*, 33262–33276. Association for Computational Linguistics.

Angell, R.; Singhal, R.; Horvitz, Z.; Yu, Z.; Ranganath, R.; McKeown, K.; and He, H. 2026. Estimating Tail Risks in Language Model Output Distributions. arXiv:2604.22167.

Askell, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; DasSarma, N.; et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.

Aswal, D.; and Jaiswal, S. D. 2026. Phonetic Perturbations Reveal Tokenizer-Rooted Safety Gaps in LLMs. arXiv:2505.14226.

Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Banerjee, S.; Chatterjee, P.; Kumar, S.; Layek, S.; Agrawal, P.; Hazra, R.; and Mukherjee, A. 2025a. Attributional safety failures in large language models under code-mixed perturbations. *arXiv preprint arXiv:2505.14469*.

Banerjee, S.; Chatterjee, P.; Kumar, S.; Layek, S.; Agrawal, P.; Hazra, R.; and Mukherjee, A. 2025b. Attributional Safety Failures in Large Language Models under Code-Mixed Perturbations. arXiv:2505.14469.

Basu, A.; and Chakraborty, P. 2026. When Names Change Verdicts: Intervention Consistency Reveals Systematic Bias in LLM Decision-Making. arXiv:2603.18530.

Beyer, T.; Xhonneux, S.; Geisler, S.; Gidel, G.; Schwinn, L.; and Günnemann, S. 2026. Position: LLM-Safety Evaluations Lack Robustness. In *Forty-third International Conference on Machine Learning Position Paper Track*.

Blasi, D.; Anastasopoulos, A.; and Neubig, G. 2022. Systematic Inequalities in Language Technology Performance across the World's Languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5486–5505. Association for Computational Linguistics.

Brahman, F.; Kumar, S.; Balachandran, V.; Dasigi, P.; Pyatkin, V.; Ravichander, A.; Wiegreffe, S.; Dziri, N.; Chandu, K.; Hessel, J.; Tsvetkov, Y.; Smith, N. A.; Choi, Y.; and Hajishirzi, H. 2024. The Art of Saying No: Contextual Noncompliance in Language Models. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Broomfield, J.; Gibbs, T.; Ingebretsen, G.; Kosak-Hine, E.; Nasir, T.; Zhang, J.; Iranmanesh, R.; Pieri, S.; Rabbany, R.; and Pelrine, K. 2025. The Structural Safety Generalization Problem. In *Findings of the Association for Computational Linguistics: ACL 2025*, 22134–22173. Association for Computational Linguistics.

Bu, Y.; Liu, X.; Ren, Z.; Yang, Y.; and Dai, J. 2026. Align Once, Benefit Multilingually: Enforcing Multilingual Consistency for LLM Safety Alignment. In *The Fourteenth International Conference on Learning Representations*.

Chan, Y. S.; Ri, N.; Xiao, Y.; and Ghassemi, M. 2025. Speak Easy: Eliciting Harmful Jailbreaks from LLMs with Simple Interactions. In *Forty-second International Conference on Machine Learning*.

Chua, G.; Tan, L.; Ge, Z.; and Lee, R. K.-W. 2025. Lost in Localization: Building RabakBench with Human-in-the-Loop Validation to Expose Multilingual Safety Gaps. In *Second Workshop on Language Models for Underserved Communities (LM4UC)*.

Dhar, R.; and Søgaard, A. 2026. Evaluation Revisited: A Taxonomy of Evaluation Concerns in Natural Language Processing. *arXiv preprint arXiv:2604.25923*.

Díaz-Millón, M.; and Olvera-Lobo, M.-D. 2021. Towards a definition of transcreation: a systematic literature review. *Perspectives*, 31: 347 – 364.

Eriksson, M.; Purificato, E.; Noroozian, A.; Vinagre, J.; Chaslot, G.; Gomez, E.; and Fernandez-Llorca, D. 2025. Can we trust AI benchmarks? An interdisciplinary review of current issues in AI evaluation. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 850–864.

Fogliato, R.; Palia, A.; Wang, X.; Sheng, E.; Atalla, C.; Garcia-Gathright, J.; Pangakis, N.; Tan, S.; Vann, D.; Washington, H.; Dow, P. A.; Elfardy, H.; Wallach, H.; and Atluri, S. 2026. ASSERT: A Measurement Pipeline for GenAI Audits. arXiv:2608.13840.

Friedrich, F.; Tedeschi, S.; Schramowski, P.; Brack, M.; Navigli, R.; Nguyen, H.; Li, B.; and Kersting, K. 2025. LLMs lost in translation: M-alert uncovers cross-linguistic safety gaps. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*.

Gureja, S.; Miranda, L. J. V.; Islam, S. B.; Maheshwary, R.; Sharma, D.; Winata, G. T.; Lambert, N.; Ruder, S.; Hooker, S.; and Fadaee, M. 2025. M-rewardbench: Evaluating reward models in multilingual settings. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 43–58.

Hu, Y.; Hee, M. S.; Nakov, P.; and Lee, R. K.-W. 2025. Toxicity Red-Teaming: Benchmarking LLM Safety in Singapore's Low-Resource Languages. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 12194–12212.

Jain, D.; Kumar, P.; Gehman, S.; Zhou, X.; Hartvigsen, T.; and Sap, M. 2024. PolygloToxicityPrompts: Multilingual Evaluation of Neural Toxic Degeneration in Large Language Models. In *First Conference on Language Modeling*.

Joshi, R.; Paul, R.; Singla, K.; Kamath, A.; Evans, M.; Luna, K.; Ghosh, S.; Vaidya, U.; Long, E. M. P.; Chauhan, S. S.; and Wartikar, N. 2025. CultureGuard: Towards Culturally-Aware Dataset and Guard Model for Multilingual Safety Applications. *ArXiv*, abs/2508.01710.

Kalluri, P. 2020. Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*, 583: 169.

Krasnodębska, A.; Dziewulska, K.; Seweryn, K.; Chrabaszcz, M.; and Kusa, W. 2026. Safety of Large Language Models Beyond English: A Systematic Literature Review of Risks, Biases, and Safeguards. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1003–1034. Association for Computational Linguistics.

Lucas, J.; Murtagh, M.; Al-Lawati, A.; Uchendu, U.; Uchendu, A.; and Lee, D. 2026. DIA-HARM: Dialectal Disparities in Harmful Content Detection Across 50 English Dialects. arXiv:2604.05318.

Mustafa, A.; and Wu, C. 2026. Evaluating Multilingual Safety Benchmarks for Low-Resource Languages in the Majority World. *Journal of Online Trust and Safety*, 3(3).

Nemecek, A.; Zafar, O.; Ganguly, D.; Singh, V.; Chaudhary, V.; and Ayday, E. 2026. Auditing Cross-Lingual Fairness in Language Model Watermarking. arXiv:2608.20047.

Ning, Z.; Gu, T.; Song, J.; Hong, S.; Li, L.; Liu, H.; Li, J.; Wang, Y.; Lingyu, M.; Teng, Y.; and Wang, Y. 2025. LinguaSafe: A Comprehensive Multilingual Safety Benchmark for Large Language Models. arXiv:2508.12733.

Pattnayak, P.; and Chowdhuri, S. 2026. IndicSafe: A Benchmark for Evaluating Multilingual LLM Safety in South Asia. arXiv:2603.17915.

Peppin, A.; Kreutzer, J.; Sebag, A. S.; Marchisio, K.; Ermis, B.; Dang, J.; Cahyawijaya, S.; Singh, S.; Goldfarb-Tarrant, S.; Aryabumi, V.; et al. 2025. The multilingual divide and its impact on global ai safety. *arXiv preprint arXiv:2505.21344*.

Pres, I.; Li, B. Z.; Ruis, L.; Guo, Z. C.; Hu, K.; Damani, M.; Puri, I.; Lubana, E. S.; and Andreas, J. 2026. Position: It's Time to Optimize LLMs for Self-Consistency. In *Forty-third International Conference on Machine Learning Position Paper Track*.

Ramesh, K.; Sitaram, S.; and Choudhury, M. 2023. Fairness in Language Models Beyond English: Gaps and Challenges. In Vlachos, A.; and Augenstein, I., eds., *Findings of the Association for Computational Linguistics: EACL 2023*, 2106–2119. Dubrovnik, Croatia: Association for Computational Linguistics.

Reuel, A.; Hardy, A.; Smith, C.; Lamparth, M.; Hardy, M.; and Kochenderfer, M. J. 2024. BetterBench: Assessing AI Benchmarks, Uncovering Issues, and Establishing Best Practices. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Röttger, P.; Pernisi, F.; Vidgen, B.; and Hovy, D. 2025. Safetyprompts: a systematic review of open datasets for evaluating and improving large language model safety. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 27617–27627.

Salaudeen, O. E.; Reuel, A.; Ahmed, A. M.; Bedi, S.; Robertson, Z.; Sundar, S.; Domingue, B. W.; Wang, A.; and Koyejo, S. 2025. Measurement to Meaning: A Validity-Centered Framework for AI Evaluation. In *NeurIPS 2025 Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling*.

Sclar, M.; Choi, Y.; Tsvetkov, Y.; and Suhr, A. 2024. Quantifying Language Models'Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. In Kim, B.; Yue, Y.; Chaudhuri, S.; Fragkiadaki, K.; Khan, M.; and Sun, Y., eds., *International Conference on Learning Representations*, volume 2024, 25055–25083.

Shelar, O.; and Desell, T. 2026. ToxSearch: Evolving Prompts for Toxicity Search in Large Language Models. arXiv:2511.12487.

Shen, L.; Tan, W.; Chen, S.; Chen, Y.; Zhang, J.; Xu, H.; Zheng, B.; Koehn, P.; and Khashabi, D. 2024. The Language Barrier: Dissecting Safety Challenges of LLMs in Multilingual Contexts. In *Findings of the Association for Computational Linguistics: ACL 2024*, 2668–2680. Association for Computational Linguistics.

Shi, E.; Shao, P.; Zhang, Y.; Cui, C.; Lyu, J.; Xia, X.; Shen, F.; and Chua, T.-S. 2026. Lingua-SafetyBench: A Benchmark for Safety Evaluation of Multilingual Vision-Language Models. arXiv:2601.22737.

Solanke, A. A. 2026. AISafetyBenchExplorer: A Metric-Aware Catalogue of AI Safety Benchmarks Reveals Fragmented Measurement and Weak Benchmark Governance. *arXiv preprint arXiv:2604.12875*.

Song, J.; Huang, Y.; Zhou, Z.; and Ma, L. 2025. Multilingual Blending: Large Language Model Safety Alignment Evaluation with Language Mixture. In *Findings of the Association for Computational Linguistics: NAACL 2025*, 3433–3449. Association for Computational Linguistics.

Souly, A.; Lu, Q.; Bowen, D.; Trinh, T.; Hsieh, E.; Pandey, S.; Abbeel, P.; Svegliato, J.; Emmons, S.; Watkins, O.; and Toyer, S. 2024. A STRONGREJECT for empty jailbreaks. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24.

Vajjala, S. 2025. The Problem with Safety Classification is not just the Models. *arXiv preprint arXiv:2507.21782*.

Wallach, H.; Desai, M.; Cooper, A. F.; Wang, A.; Atalla, C.; Barocas, S.; Blodgett, S. L.; Chouldechova, A.; Corvi, E.; Dow, P. A.; Garcia-Gathright, J.; Olteanu, A.; Pangakis, N. J.; Reed, S.; Sheng, E.; Vann, D.; Vaughan, J. W.; Vogel, M.; Washington, H.; and Jacobs, A. Z. 2025. Position: Evaluating Generative AI Systems Is a Social Science Measurement Challenge. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, 82232–82251. PMLR.

Wan, Y.; Pu, G.; Sun, J.; Garimella, A.; Chang, K.-W.; and Peng, N. 2023. "Kelly is a Warm Person, Joseph is a Role Model": Gender Biases in LLM-Generated Reference Letters. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 3730–3748. Singapore: Association for Computational Linguistics.

Wang, W.; Tu, Z.; Chen, C.; Yuan, Y.; Huang, J.-t.; Jiao, W.; and Lyu, M. 2024. All Languages Matter: On the Multilingual Safety of LLMs. In *Findings of the Association for Computational Linguistics: ACL 2024*, 5865–5877. Association for Computational Linguistics.

Weng, Z.; Jin, X.; Jia, J.; and Zhang, X. 2025. Foot-In-The-Door: A Multi-turn Jailbreak for LLMs. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 1939–1950. Association for Computational Linguistics.

Wu, M.; and Aji, A. F. 2025. Style Over Substance: Evaluation Biases for Large Language Models. In *Proceedings of the 31st International Conference on Computational Linguistics*, 297–312. Association for Computational Linguistics.

Yong, Z. X.; Ermis, B.; Fadaee, M.; Bach, S.; and Kreutzer, J. 2025. The State of Multilingual LLM Safety Research: From Measuring The Language Gap To Mitigating It. In Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; and Peng, V., eds., *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 15845–15860. Suzhou, China: Association for Computational Linguistics. ISBN 979-8-89176-332-6.

Yoo, H.; Yang, Y.; and Lee, H. 2025. Code-Switching Red-Teaming: LLM Evaluation for Safety and Multilingual Understanding. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13392–13413. Association for Computational Linguistics.

Yu, C.; Engelmann, S.; Cao, R.; Ali, D.; and Papakyriakopoulos, O. 2026. How Should AI Safety Benchmarks Benchmark Safety? *arXiv preprint arXiv:2601.23112*.

Zhang, Z.; Chen, B.; Hu, Y.; Qu, Y.; Zhu, T.; and Gao, L. 2025. When Harmless Words Harm: A New Threat to LLM Safety via Conceptual Triggers. arXiv:2511.21718.

Zhao, Y.; Chen, Z.; Ma, X.; Jiang, Y.-G.; and Li, B. 2026. ML-Bench&Guard: Policy-Grounded Multilingual Safety Benchmark and Guardrail for Large Language Models. arXiv:2605.00689.

A Annotation Analysis

A.1 Papers selected for Annotation

The papers we selected for annotation are listed in Table 2, along with an excerpt of our annotations (three attributes).

A.2 Annotation Attributes

The full list of attributes that were annotated for each empirical paper is in Table 3.

A.3 Inconsistency Metrics

We identify five papers that explicitly quantify inconsistency in safety performance across languages (or equivalently, across a pertubation function):

- Basu and Chakraborty (2026) measure variation in a decision-making system (classification task) due to demographic (name/race swaps), authority (credential/prestige swaps), and framing (positive/negative restatements) features. Inconsistency is measured using the *flip rate*: the proportion of instances where the model decision flips under an input perturbation. Analogous to this, (Friedrich et al. 2025) report the *exact matching rate* for binary safety judgments for a query–translation pair, as well as the average *pair-wise agreement* across translated pairs.
- Bu et al. (2026) report two metrics of inconsistency in binary judgments of safety across languages: a macro-level measure of *variance* in the aggregate score across languages, along with an instance-level measure of average *pair-wise agreement* across judgments for pairs of translated queries.
- Pattnayak and Chowdhuri (2026) propose two measures of consistency in safety judgments assigned to prompts in different languages: an aggregate-level *agreement rate* computes the proportion of prompts that receive the same judgment across all translations; at the prompt-level, *entropy* of labels across all translations is used as a measure of inconsistency. Additionally, they report summary metrics of the proportion of prompts assigned to each of their categorical labels (refusal, harmful, harmless, ambiguous).
- Lucas et al. (2026) measure the robustness of harmful content detection systems to dialectal variations of Standard American English (SAE). Their proposed *robustness gap* metric measures the deviation (delta) in accuracy between a dialect-transformed data subset and SAE.

Reference	Safety concept	Indicator type	Stress test
Shen et al. (2024)	refusal, relevance	categorical	translation
Basu and Chakraborty (2026)	decision flip	binary	perturbations based on protected characteristics
Jain et al. (2024)	toxicity	continuous, threshold	translation, multilingual curation, repeated sampling
Shi et al. (2026)	safety	binary	translation
Yoo, Yang, and Lee (2025)	safety, refusal, relevance	continuous	code switching
Song et al. (2025)	safety	binary, probability	code switching
Ning et al. (2025)	refusal, safety, over-refusal, helpfulness	binary	translation, multilingual curation
Chua et al. (2025)	safety	binary	multilingual curation, translation
Zhao et al. (2026)	safety, policy compliance	binary	multilingual curation, paraphrasing, adversarial attack
Bu et al. (2026)	safety	binary	(none)
Friedrich et al. (2025)	safety	binary	translation
Pattnayak and Chowdhuri (2026)	refusal, safety	categorical	translation
Hu et al. (2025)	toxicity, bias	binary	translation, adversarial attack
Abdullahi et al. (2026)	policy compliance	binary	translation
Lucas et al. (2026)	disinformation	binary	translation
Souly et al. (2024)	refusal, harmfulness	ordinal	adversarial attack
Brahman et al. (2024)	disinformation	binary	none
Weng et al. (2025)	harmfulness	binary	adversarial attack
Broomfield et al. (2025)	refusal	binary	translation, adversarial attack
Chan et al. (2025)	refusal, harmfulness	continuous	translation, paraphrasing, adversarial attack
Angell et al. (2026)	harmfulness	threshold, continuous	repeated sampling
Wang et al. (2024)	safety	binary	translation
Zhang et al. (2025)	harmfulness	ordinal	adversarial attack
Aswal and Jaiswal (2026)	relevance, safety	binary	adversarial attack, code switching
Shelar and Desell (2026)	toxicity	probability	paraphrasing
Banerjee et al. (2025b)	harmfulness	binary	code switching

Table 2: Excerpt of annotations (three attributes) for the 26 selected empirical papers.

Attribute	Explanation
Annotation Details	
Title, Year, Citation, Added by, Annotated by	–
Paper Overview	
Focus / Main Research Question	What is the central research question investigated by the paper?
Object of Interest	What role does the LLM play in the study (e.g., response generator, safety classifier, coding agent)?
Datasets	
Languages Evaluated	Which languages are included in the evaluation? Enumerated.
Data Sources Used	What datasets and benchmarks are used?
Harm Categories	Which safety or harm categories are considered (e.g., hate, violence, self-harm, misinformation)?
Context of Use	
Setting	Is the evaluation conducted in a benign, adversarial, red-teaming, or other context?
Interaction Type	Is the interaction single-turn, multi-turn, or another dialogue format?
Safety behaviour	
Safety Measure: Concept	What overall notion of safety behaviour is being measured (e.g., refusal, harmfulness, likelihood of toxicity)?
Safety Measure: Indicator	What indicator variable is used to operationalize safety at the instance-level (e.g., binary, likert-scale, formula combining multiple measures)?
Safety Measure: Aggregation	What is the aggregate-level safety score (e.g., attack success rate, refusal rate)?
Safety Measure: Scorer	Who or what determines whether a response is safe (e.g., human annotators, automated classifier, LLM judge)?
Behavioural Inconsistency	
Stress Test / Inconsistency Explored	What perturbation, variation, or model instability is introduced to test consistency of safety?
Inconsistency Metric	Metric/formula used to quantify safety inconsistency?
Experiment Details	
LLMs Tested	Which language models are evaluated? Enumerated.
Other Experiment Details	Any additional methodological details relevant to the experiments.
Findings	
Main Conclusions	What are the primary findings of the paper?
Policy or Evaluation Recommendations	Does the paper provide recommendations for policy, governance, benchmarking, or evaluation practice?
Useful Artifacts	Are any datasets, benchmarks, codebases, libraries, or other resources released?
Other Notes	Any additional observations not captured elsewhere.
Keep for Analysis?	Should the paper be retained for the review/analysis (Yes/No)?

Table 3: Full list of annotation attributes for empirical papers.