



PDF Download
3715275.3732054.pdf
24 January 2026
Total Citations: 5
Total Downloads: 2300

Latest updates: <https://dl.acm.org/doi/10.1145/3715275.3732054>

RESEARCH-ARTICLE

Policy-as-Prompt: Rethinking Content Moderation in the Age of Large Language Models

KONSTANTINA PALLA, Spotify AB, Stockholm, Stockholm, Sweden

JOSÉ LUIS GARCÍA, Spotify AB, Stockholm, Stockholm, Sweden

CLAUDIA HAUFF, Spotify AB, Stockholm, Stockholm, Sweden

FRANCESCO FABBRI, Spotify AB, Stockholm, Stockholm, Sweden

ANDREAS C DAMIANOU, Spotify AB, Stockholm, Stockholm, Sweden

HENRIK LINDSTRÖM, Spotify AB, Stockholm, Stockholm, Sweden

[View all](#)

Open Access Support provided by:

[Spotify AB](#)

[Spotify USA Inc](#)

Published: 23 June 2025

[Citation in BibTeX format](#)

FACCT '25: The 2025 ACM Conference
on Fairness, Accountability, and
Transparency
June 23 - 26, 2025
Athens, Greece

Policy-as-Prompt: Rethinking Content Moderation in the Age of Large Language Models

Konstantina Palla
Spotify
London, United Kingdom
konstantina.palla@gmail.com

José Luis Redondo García
Spotify
Madrid, Spain
joseluisr@spotify.com

Claudia Hauff
Spotify
Amsterdam, Netherlands
claudiah@spotify.com

Francesco Fabbri
Spotify
Barcelona, Spain
francescof@spotify.com

Andreas Damianou
Spotify
London, United Kingdom
andreasd@spotify.com

Henrik Lindström
Spotify
Stockholm, Sweden
henok@spotify.com

Dan Taber
Spotify
Seattle, USA
dtaber@spotify.com

Mounia Lalmas
Spotify
London, United Kingdom
mounial@spotify.com

Abstract

Content moderation plays a critical role in shaping safe and inclusive online environments, balancing platform standards, user expectations, and regulatory frameworks. Traditionally, this process involves operationalising policies into guidelines, which are then used by downstream human moderators for enforcement, or to further annotate datasets for training machine learning moderation models. However, recent advancements in large language models (LLMs) are transforming this landscape. These models can now interpret policies directly as textual inputs, eliminating the need for extensive data curation. This approach offers unprecedented flexibility, as moderation can be dynamically adjusted through natural language interactions. This paradigm shift raises important questions about how policies are operationalised and the implications for content moderation practices. In this paper, we formalise the emerging policy-as-prompt framework and identify five key challenges across four domains: **Technical Implementation** (1. translating policy to prompts, 2. sensitivity to prompt structure and formatting), **Sociotechnical** (3. the risk of technological determinism in policy formation), **Organisational** (4. evolving roles between policy and machine learning teams), and **Governance** (5. model governance and accountability). Through analysing these challenges across technical, sociotechnical, organisational, and governance dimensions, we discuss potential mitigation approaches. This research provides actionable insights for practitioners and lays the groundwork for future exploration of scalable and adaptive content moderation systems in digital ecosystems.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
FAccT '25, Athens, Greece

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1482-5/25/06
<https://doi.org/10.1145/3715275.3732054>

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Social and professional topics** → **Professional topics**.

Keywords

content moderation, automated decision making, reliable machine learning

ACM Reference Format:

Konstantina Palla, José Luis Redondo García, Claudia Hauff, Francesco Fabbri, Andreas Damianou, Henrik Lindström, Dan Taber, and Mounia Lalmas. 2025. Policy-as-Prompt: Rethinking Content Moderation in the Age of Large Language Models. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3715275.3732054>

1 Introduction

Content moderation involves the systematic monitoring and regulation of user-generated content in online platforms. This critical practice fosters safe, inclusive, and respectful digital spaces while aligning with the values of hosting organisations and adhering to regulatory requirements. As digital ecosystems grow and content creation surges, the demand for scalable and adaptive moderation systems has become increasingly urgent. Online platforms now face the dual challenge of managing vast, dynamic content streams and meeting evolving societal and regulatory expectations. With the internet continuing to transform how we access information, entertainment, and services, content moderation has become essential for effectively managing online platforms, bridging local and global perspectives to address societal needs and foster community trust.

The architecture of moderation pipelines varies significantly across digital services, reflecting their unique operational needs. For instance, video streaming platforms often prioritise automated detection of copyrighted or harmful content, while community forums like Reddit [52] emphasise moderating user interactions and discussions, often requiring human oversight alongside automated

tools [19]. Despite these platform-specific differences, successful content moderation fundamentally relies on two interconnected processes: *operationalisation* and *policy enforcement* [60].

Operationalisation involves transforming abstract policy objectives into actionable protocols for consistent implementation by human moderators and algorithms. This process includes several critical components: developing comprehensive annotation guidelines, curating high-quality training datasets, setting precise model predictions thresholds, and conducting thorough annotator training. Through operationalisation, abstract policies evolve into specific rules supported by robust workflows and systems. Content moderation policies are implemented through various documented forms, which we refer to as policy artifacts. These artifacts operate at different levels of granularity within an organization to cover its moderation needs. At the highest level, platform-wide rules establish foundational principles that guide moderation strategies. These principles then inform more granular artifacts, such as product-specific policies and detailed annotation guidelines, which address domain-specific requirements. This hierarchical structure ensures policies remain broadly applicable while allowing contextual adaptability.

Enforcement puts these artifacts into practice, guiding decision-making processes for both human reviewers and algorithmic systems. During enforcement, content is systematically evaluated against context-appropriate policy artifacts, translating the groundwork laid during operationalisation into concrete moderation actions. Operationalisation provides the foundation for consistent, clear enforcement and works in tandem with the enforcement process. The interplay between operationalisation and enforcement is dynamic and interdependent, with changes or inconsistencies in one often impacting the other. Policy artifacts are regularly updated to address emerging content scenarios, refine enforcement nuances, or establish additional boundaries. For example, the appearance of new content types or edge cases may require adjustments to annotation guidelines, ensuring both moderators and algorithms remain aligned with overarching principles.

Technology plays a crucial role in advancing both operationalisation and enforcement, enabling these processes to adapt and remain effective in response to evolving challenges. As the exponential growth of content volumes has driven unprecedented demands on moderation systems, significant technological advancements have emerged to meet these needs. Platforms increasingly rely on automated solutions, with machine learning algorithms becoming central to the moderation pipeline [50]. This work focuses on a transformative development in AI-assisted enforcement: the ability to encode content moderation policies directly as natural language prompts in LLMs, a paradigm we term “**policy-as-prompt**” and show in Figure 1b. By eliminating the need for extensive manual annotation pipelines, this approach offers unparalleled flexibility and scalability for moderation systems, enabling rapid policy iteration and fine-grained control over enforcement decisions. This emerging approach is gaining significant traction - enforcement strategies continue to vary across platforms, but the adoption of AI-driven methods, particularly those leveraging LLMs with injectable guidelines, is rapidly redefining the content moderation landscape [23, 36].

Table 1: Challenges across different areas in the ‘policy-as-prompt’ implementation.

| Area & Challenges |
|--|
| Technical Implementation 1 Converting policies to prompts 2 Prompt structure and format sensitivity |
| Sociotechnical 3 Technological determinism in policy formation |
| Organizational 4 Evolving policy-ML team roles |
| Governance 5 Model governance and accountability |

To the best of our knowledge, this work presents the first comprehensive exploration of the policy-as-prompt paradigm, examining its key challenges and implications across technical, sociotechnical, organisational, and regulatory domains. Beyond identifying these challenges, we systematically analyse their interdependencies and implications, providing a foundation for further inquiry into this transformative approach. In Section 2 we examine the technological evolution of content moderation, tracing its progression from basic pattern matching to modern LLM-driven approaches. We explore how this progression has fundamentally altered moderation architectures and introduced the emerging policy-as-prompt. In Sections 3 to 6 we break down five key challenges across four domains as outlined in Table 1; **Technical Implementation** (focusing on the challenges of 1. converting policies to prompts and 2. addressing prompt structure and format sensitivity, which we support with empirical findings), **Sociotechnical** (3. highlighting the impact of technological determinism in policy formation), **Organisational** (4. examining the evolving roles of policy-ML teams), and **Governance** (considering the 5. model governance and accountability). We provide a list of recommendations towards mitigating the challenges in Section 7. We discuss the limitations of this work along with concluding remarks in Section 8.

2 The Evolution of Content Moderation Technology

The evolution of content moderation, particularly within the text domain, has been closely tied to the exponential growth of user-generated content, driven by the rapid advancement of algorithmic capabilities. Early moderation efforts relied on basic automation, functioning like a digital “find and replace” tool. Methods such as keyword filters and hash-matching were used to identify and remove prohibited content [14]. While these first-generation tools were effective at detecting exact matches, they were easily circumvented by minor alterations, underscoring their limitations in adapting to the dynamic nature of online expression [51].

The rapid increase in user-generated content demanded more robust and scalable solutions, driving a shift towards machine learning-based approaches [50]. Early word embedding models like Word2Vec [38] and GloVe [47] enabled moderation approaches to capture semantic relationships between words, enhancing their

ability to detect variations in harmful phrases and toxic language. This marked a pivotal step in moving beyond simple pattern matching to understanding the meaning behind online content.

However, the true breakthrough came with the introduction of contextual embedding models [49] and the transformative Transformer architecture [65]. Models like BERT [12] enhanced content moderation by analysing content in context, enabling systems to improve detection of nuances such as sarcasm, coded language, and implicit bias that were previously undetectable. This advancement significantly improved the ability to understand the intent and impact of online communication, specifically in the realm of textual content [7].

Today, advanced language models like GPT series [43], LLaMA series [63] and Claude series [3] are redefining the boundaries of content moderation. These models depart from traditional supervised learning—where systems learn from training examples—and are able to interpret desired behaviour directly from textual instructions. Their natural language understanding and reasoning capabilities allow them to go beyond simply identifying harmful content—they can assess its context and potential impact. They showcase capabilities such as applying complex policy guidelines, engaging in nuanced dialogue with users, and adapting to evolving language and online trends [1, 5, 11, 26, 27, 41, 45]. Rather than relying solely on rigid classification models, content moderation now incorporates general-purpose reasoning systems that provide more flexible and context-sensitive oversight. This approach allows for greater adaptability in addressing emerging challenges, making moderation more effective and responsive to the complexities of online discourse.

By enabling the direct integration of policy guidelines into prompts, LLMs are reshaping the structure of moderation workflows. In the next section, we explore the details of the policy-as-prompt setup, examining its impact on workflow design and comparing it with traditional frameworks to provide a clearer context for this paradigm shift.

2.1 The Traditional Algorithmic Moderation Pipeline

The adoption of the ML-based approaches required the development of extensive infrastructures for data collection, labelling, model training, and deployment. These requirements fundamentally transformed how platforms approached content moderation at an architectural level. A standard framework has emerged that integrates human expertise with machine learning models, creating a scalable and widely adopted industry (see Figure 1a). At the core of this framework are human annotators and machine learning models trained on annotated datasets [17, 19]. Human experts label content based on predefined guidelines, producing datasets that encapsulate these guidelines. These annotated datasets are then used to train machine learning models, enabling them to assess whether content violates platform policies. During training, the model adjusts its internal parameters—known as weights—to better align with the patterns identified in the annotated data. Once trained, the model can be deployed to evaluate new, real-world content, autonomously flagging potential policy violations.

This process is inherently iterative, with models continuously refined based on feedback and new data. However, real-world applications bring additional complexities. Annotators may disagree on how specific content should be labelled, leading to inconsistencies or ambiguities within the dataset [28, 55]. Furthermore, guideline may undergo revisions, but retrospective relabelling of existing data is often impractical, complicating the training process. Despite these challenges, this integration of human expertise and machine learning has become a foundational approach to scalable content moderation.

2.2 Policy-as-prompt

While the annotated dataset approach has long been the industry standard in moderation pipelines with an algorithmic component, the emergence of LLMs has reimaged how guidelines can be integrated into the moderation process. LLMs build their knowledge through extensive pretraining on web-based text, leveraging the transformer architecture [65]. This process employs self-supervised learning, where models predict subsequent words in incomplete sentences, enabling them to associate words and phrases with their typical contexts. As a result, LLMs acquire an internal knowledge of language that enables them to generate coherent and contextually relevant responses across diverse scenarios. Interaction with LLMs relies on prompts—input text that guides the model’s response. LLMs encode the prompt into a high-dimensional vector space, preserving semantic relationships between words and phrases. These representations inform the model’s output, drawing on patterns learned during pretraining.

The quality of an LLM’s response is profoundly influenced by the prompt. A well-crafted prompt ensures that the model’s output aligns with the intended outcome, highlighting the prompt as a pivotal element in effective LLM interaction. Factors such as phrasing, specificity, and context within the prompt play a critical role in shaping the response. As a result, *prompt engineering* has emerged as a systematic practice to extract knowledge from LLMs. This involves designing precise, task-specific instructions, that guide models toward accurate, relevant, and coherent outputs without retraining or altering the model’s parameters (weights) [4]. Techniques range from simple instruction-based prompts to advanced methods like in-context learning, where illustrative examples are included in the prompt to steer the model towards the desired output [54, 56] (see Figure 2). Other sophisticated techniques, such as Chain-of-Thought prompting (CoT), encourage the model to break down complex reasoning into step-by-step deductions [68], while Reason and Act (ReAct) prompts guide the model to interact with external tools or information sources to enhance its responses [71].

In content moderation, prompts enable quick adaptation of LLMs to changing policies or criteria, offering a flexible approach that avoids the need for retraining (see Figure 1b). By re-engineering the prompt with updated policy guidelines, LLMs can swiftly respond to evolving societal norms, emerging risks, or platform-specific needs. In this setting, policy guidelines become an integral part of the prompt, allowing for seamless updates and adjustments to content moderation practices.

LLMs in content moderation currently operate in hybrid setups where human moderators collaborate with and oversee LLM-driven

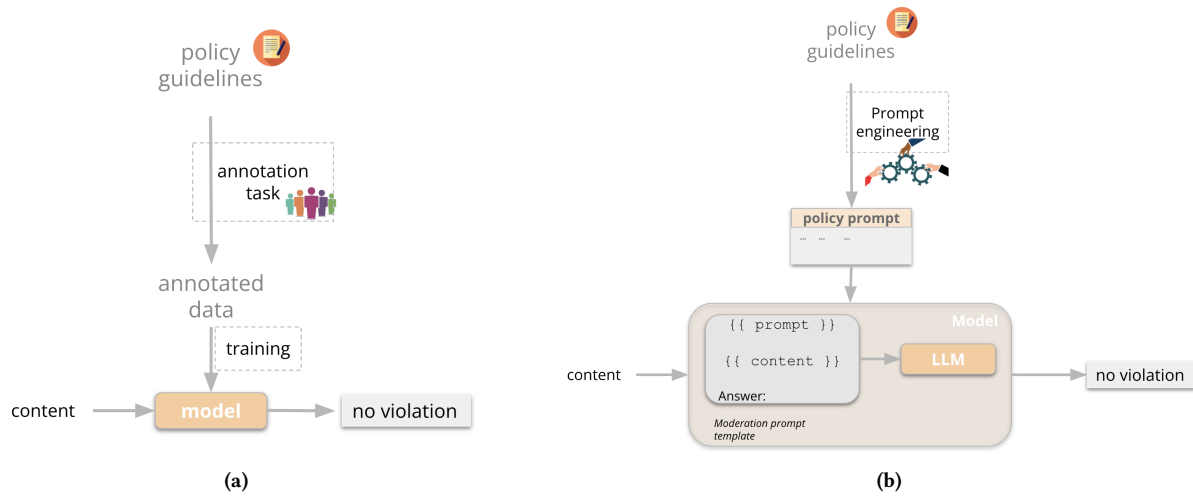


Figure 1: Approaches to content moderation. (a) Traditional pipeline: policy guidelines inform human annotation, which produces training data for models. (b) Policy-as-prompt: policy guidelines are encoded directly as prompts, enabling LLMs to perform moderation without explicit annotation datasets.

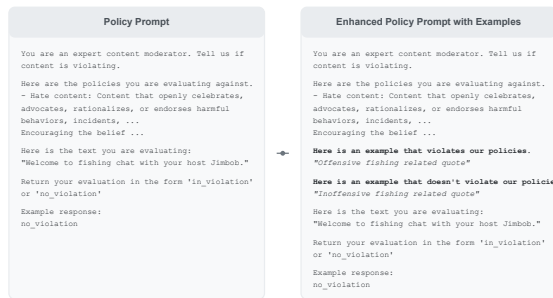


Figure 2: Example prompts demonstrating two approaches for algorithmic content moderation using policy guidelines. Left: A basic prompt where the policy text is provided to the model for direct review of content. Right: An enhanced prompt, following the ‘in-context learning’ technique, that includes both the policy text and specific examples of violating and non-violative content, aiding the model in contextualizing its decisions.

enforcement, but they have the potential to transform content moderation if given greater autonomy. To ensure this does not lead to harm, we need research to understand the complexities of this future. As organizations adopt this paradigm, they must navigate challenges spanning technical, sociotechnical, organisational, and governance dimensions as we show in Table 1. **Technical implementation** challenges emerge from the fundamental change in how policy guidelines are operationalised—moving from trained models based on annotated datasets to dynamic prompts that directly guide LLM behaviour. This shift introduces concerns around accurately converting policy to prompts and the sensitivity of LLMs

to variations in prompt structure and formatting. At the **sociotechnical** level, challenges stem from the limitations of LLMs, which can constrain human judgment and social adaptability. Embedding policies directly into LLM prompts risks reinforcing technological determinism, where policies are rigidly interpreted based on the model’s design, and potentially overlooking the nuanced interplay between technical systems and the human contexts in which they operate. On an **organisational** front, this shift demands changes to workflows and expertise. Traditional content moderation relied on distinct roles between policy formulation and operationalisation, where collaboration was often limited. Policy experts defined guidelines, and machine learning teams operationalised them through annotated datasets. In contrast, policy-as-prompt approaches necessitate hybrid expertise and foster closer collaboration between policy experts and machine learning practitioners. Finally, the model **governance** domain encompasses challenges related to transparency and accountability. Ambiguities in decision attribution, coupled with difficulties in tracking prompt modifications, hinder auditability and complicate responsible deployment.

3 Technical Implementation Challenges

3.1 Converting Policies to Prompts

Traditionally, policy writing has been a human-centric endeavour, where subject matter experts meticulously crafted guidelines designed for clarity and human interpretation. These guidelines were refined through iterative feedback from human reviewers, including annotators, moderators, and domain experts, who brought contextual understanding, nuanced judgment, and collective experience to the process. Ambiguities were addressed through deliberation and testing. For instance, annotators would extensively test policies by creating hypothetical scenarios, exploring edge cases, and providing qualitative feedback, enhancing the robustness of the

guidelines. The primary metric of success was the policy’s clarity and actionability for human practitioners.

In the policy-as-prompt paradigm, the policy prompt becomes a distinct artifact. Policy writing now serves a dual purpose: it must remain accessible to human understanding, while also be supplemented with representations optimised for machine interpretation. This requires translating a policy’s intent and rules into a format that is optimised for machine processing while faithfully preserving its meaning. However, verifying whether a prompt accurately conveys the policy’s intent to the model is challenging. Such verification often relies on other signals, such as comparing LLM moderation decisions against human-labelled ground truth or expert judgments, to assess whether the model applies the policy as intended.

In traditional supervised learning approach, policies were operationalised through labelled training data, embedding nuance across many examples in the dataset. While this approach had its own challenges—such as human annotators occasionally misunderstanding policies or models generalised imperfectly—it benefited from redundancy in the dataset, which allowed for statistical smoothing through the sheer volume of labelled examples. In contrast, policy-as-prompt encodes intent directly into the prompt, requiring the prompt to fully capture the nuances of the policy. Unlike human-centric policy development and enforcement, which relies on human interpretation, writing for machines relies on machine interpretation. For example, a policy stating “remove content promoting illegal activities” may be interpreted too broadly by a model, resulting in the removal of educational content discussing historical prohibition, or too narrowly, by failing to catch euphemistic or coded language. Without explicit and diverse test scenarios, such discrepancies between policy intent and model behaviour may go undetected.

Additionally, as discussed in Section 3.2, the process is highly sensitive to formatting and phrasing. Seemingly minor changes—such as using bullet points instead of paragraphs or altering punctuation—can significantly influence model outputs. Similarly, prompts that encode abstract policy principles (e.g., “respect user dignity”) require precise operationalisation to ensure models act consistently with human expectations. Yet, LLMs often struggle with abstraction unless such principles are made concrete through examples or clarifying constraints. Our empirical findings (Figure 5) highlight this sensitivity: we observe that augmenting policy prompts with additional facts and scenarios measurably alters model outputs, indicating that the level of detail included during prompt formulation significantly impacts enforcement behaviour [32].

The internal workings of LLMs remain largely opaque, leaving prompt engineering an experimental process guided by trial and error [8]. This challenge is further compounded by the lack of shared contextual understanding between humans and machines; LLMs often misinterpret nuanced or implicit cues that are easily understood by human readers [2, 70]. Lastly, unlike traditional approaches, there is no intermediary dataset to more exhaustively mitigate these sensitivities, making the translation of policies into prompts, a non-trivial and error-prone task.

3.2 Prompt Structure and Format Sensitivity

In LLM-based content moderation systems, the structure and format of policy guidelines embedded within prompts are critical factors influencing performance. The presentation and organisation of policy text can directly affect moderation outcomes. Recent empirical studies highlight the significant, yet often overlooked, impact of prompt structure on LLM performance. For instance, Levy et al. [30] found that increasing input length can negatively impact performance, with significant drops occurring well before reaching the models’ maximum input capacity. Similarly, [34] showed that performance drops significantly when relevant information is repositioned, highlighting a lack of robustness in long-context processing. He et al. [21] demonstrated that seemingly minor formatting choices—such as using bullet points instead of paragraphs—can substantially influence model performance across tasks. Perhaps most concerning, Sclar et al. [57] showed that LLMs exhibit unexpected sensitivities to superficial formatting elements like capitalisation and whitespacing, challenging common assumptions about their robustness.

The precise details of why certain prompt structures work better than others is still an active area of research in AI. Related works have observed that these sensitivities likely emerge from fundamental aspects of LLM architecture and pretraining. The models process text through self-attention mechanisms that compute relationships between all tokens in the input sequence [65]. While theoretically capable of handling arbitrary input structures, the effectiveness of these attention patterns is heavily influenced by how information is organised in the prompt. Zhang et al. [73] and Li et al. [31] demonstrated that different prompt structures lead to distinctly different attention patterns, affecting how information flows through the model’s layers. Additionally, Kazemnejad et al. [25] showed that the positional encoding schemes used in transformer architectures can create inherent biases in how models process information at different positions in the sequence, potentially explaining the observed sensitivity to information positioning.

These findings underscore that the structural choices in presenting information to LLMs are far from superficial. Instead, they represent fundamental parameters that can significantly impact model behavior and reliability. This sensitivity has critical implications for applications like content moderation, where consistent and reliable performance is essential to ensure fairness, accuracy, and trust in automated decisions. Further, this structural sensitivity parallels broader concerns in machine learning regarding how seemingly minor design choices can produce significant downstream effects. Recent research shows that variations in machine learning pipeline components, even when they do not substantially affect overall accuracy metrics, can lead to distinct error patterns or divergent outcomes on specific data subsets. This phenomenon, known as *predictive multiplicity*, has been observed across various classification tasks and raises particular concerns in high-stakes applications like content moderation [18, 58]. Similarly, in the policy-as-prompt setting, variations in prompt formatted may result in inconsistent moderation decisions on specific pieces of content, even when overall performance metrics appear unchanged.

3.2.1 Demonstration. To support our discussion, we empirically investigate how the structure, style and the format of policy prompts influence LLM-based content moderation performance. We also

further investigate whether predictive multiplicity emerges in this scenario.

Dataset. We collaborated with moderators from a large-scale music platform to create a dataset of 2,115 text descriptions labelled across sixteen different moderation categories. We report in this study six categories, which are widely used in industry for content moderation [60]: ‘not violating’, ‘Harm’, ‘Self Harm’, ‘Hate’, ‘Sexual Content’ and ‘jailbreak’. The dataset comprises user-generated requests for playlist creation, offering real-world data for studying moderation. Each entry was labelled as not violating or falling under one specific policy category from the list of considered ones.

Setup. Our experimental architecture follows the policy-as-prompt framework depicted in Figure 1b. For the LLM component, we use GPT-4o-mini [44], for its balance between performance and cost while maintaining characteristics representative of industry-standard models. The model’s `{ prompt }` included a dictionary with pairs of policy label and policy descriptions for each of the categories considered, alongside the preamble instruction: ‘You are a helpful content moderation assistant. You categorise user input into one of these categories depending on whether the input violates the category’. We use constrained decoding to prevent the model from deviating from the set of considered categories. While this architecture reflects common industry practices, our specific implementation choices—such as model selection, prompt design, and policy structures—were tailored for research purposes and are not representative of production systems.

First, we examine how subtle variations in prompt structure and style influence model performance. To investigate this, in Figure 3 we focus on cases when semantic differences between templates and the original policy text (the *Baseline*) are minimal. The variations include: *Punctuated*, altered punctuation to assess sensitivity to grammatical changes; *Structured*, used bullet points and sections to enhance clarity and facilitate programmatic processing; *Concise*, streamlined phrasing to test the impact of brevity; *Verbose*, detailed and elaborate descriptions to examine the effect of information density; and *Annotator-Friendly*, tailored for human comprehension, emphasising simplicity and clarity. Detailed examples of these prompt variations are provided in Figure 7 in Appendix Section A. We further explore how differences in policy prompt formatting—such as plain text, XML, YAML and JSON—affect model performance in Figure 4.

Figure 3a reveals variability in accuracy across different prompt types, despite the fact that all these variations include similar information. Structured prompts demonstrated the highest accuracy, showing model’s preference towards clarity and organisation. In contrast, verbose prompts performed the worst, indicating that excessive detail may hinder the model’s understanding. ‘Punctuation’ and ‘Concise’ prompts showed comparable performance, with average accuracies of 0.7890 and 0.7897, respectively (p -value = 0.631). Figure 3b compares accuracy across policy categories for the two prompt types, revealing evidence of predictive multiplicity; despite similar overall performance, seemingly identical prompts produced different predictions for the same samples. The results in Figure 4 confirm that variations in policy formats lead to differences in performance..

Next, we explore how increasing the amount of policy text in the prompt impacts model performance. To do this, we incrementally add information in the policy descriptions, such as extra facts and examples. This process resulted in sequential snapshots of policies with varying levels of information, with which we analysed for model performance, as illustrated in Figure 5a. This approach mirrors the evolution of policy guidelines on digital platforms, where trends and new boundaries emerge over time. On average, model performance across all policy categories improves as the policy prompt becomes more detailed (greater semantic coverage). However, when examining the model’s performance for individual policy categories at each snapshot (Figure 5b), we observe variability, suggesting that different categories are influenced to varying degrees by the increase in prompt information.

Discussion. These experiments reveal that the policy-as-prompt approach, while pragmatic and scalable, can introduce significant variability in model behaviour. Our findings highlight that even subtle changes in prompt phrasing, structure, or format, such as adding punctuation, altering layout, or translating policies into structured formats, can measurably affect moderation outcomes. Crucially, this variability has practical implications for real-world deployments. First, it raises concerns about *moderation inconsistency*: differences in prompt structure can result in identical content being flagged in one instance but allowed in another, leading to unpredictable enforcement. Second, it introduces a *risk of under- or over-moderation*: depending on how policies are translated into prompts, harmful content may be inadequately flagged, or benign content may be erroneously removed. Third is the potential of *unfairness and inequity*: if prompt sensitivity leads to inconsistent enforcement, certain user groups or content types may be disproportionately affected.

4 Sociotechnical Challenge: Technological Determinism in Policy Formation

Another challenge we identify in the policy-as-prompt approach is that of technological determinism. Technology significantly influences societal values and structures. The concept of technological determinism posits that technology, along with its design and inherent capabilities, can shape social structures, cultural norms, and even political systems [22, 69]. Algorithmic systems, particularly those used in content moderation, have long been susceptible to technological determinism due to limitations such as model bias and undertraining [19]. Traditionally, human oversight in content moderation served as a vital feedback loop. Human annotators played a key role in shaping training data, ensuring that machine learning models were informed by social contexts and values [17]. This interaction mitigated the influence of technology, enabling the development of nuanced rules that accounted for context and intent, incorporated ethical considerations, and provided mechanisms for human appeals to challenge automated decisions [53].

The advent of LLMs, particularly proprietary and opaque third-party models, is now shifting this balance. By enabling the direct use of policy guidelines in moderation workflows, thus bypassing human annotation, LLMs amplify the risks of technological determinism. As discussed in Section 3, LLMs excel at processing structured, algorithmic-friendly guidelines. However, this capability may inadvertently push experts to prioritise machine clarity




Figure 3: Effect of prompt design on accuracy: (a) Variation in accuracy across prompt types, averaging over five runs for each type and, (b) Performance differences across policy categories for the ‘Punctuation’ and ‘Concise’ prompt types.


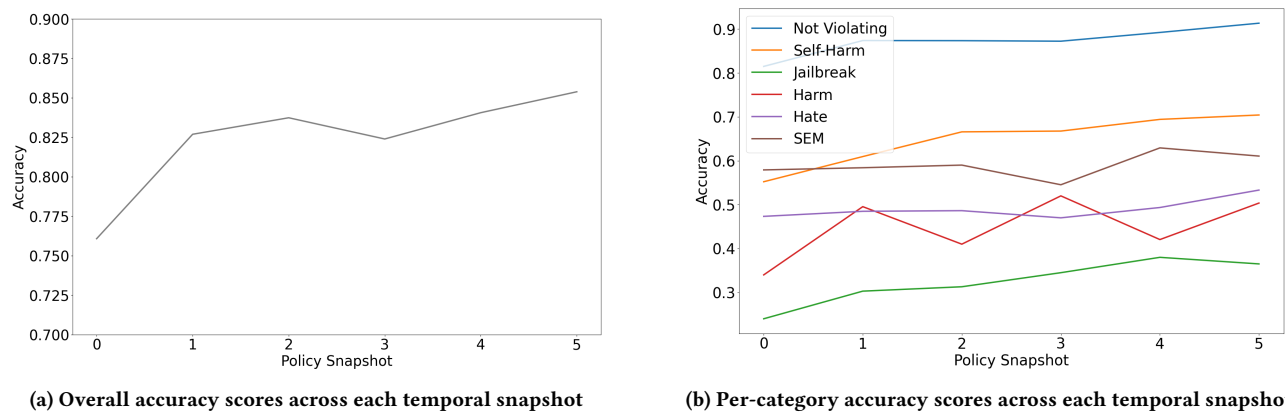


Figure 4: Performance spread (accuracy) for modifications in the format in which the policy is plugged into the prompt. ‘Baseline’ refers to the plain text format.



(a) Overall accuracy scores across each temporal snapshot

(b) Per-category accuracy scores across each temporal snapshot

Figure 5: Analysing both overall and per-category performance of the model when injecting different temporal snapshots of the policies

over the nuance and flexibility required for context-dependent rules. Structured formulations might be favoured even when they do not fully align with underlying policy objectives, creating an environment where complex social issues are reduced to simplistic,

binary judgments. This echoes Lessig’s “code as law” principle, where technical architecture becomes a primary regulator of online behaviour [29]. Similarly, Yeung [72] describes “algorithmic regulation” as the process by which designing and deploying algorithms

transforms into a form of governance, shaping both individual and collective behaviour. Finally, van Dijck [64] further argues that platform governance is increasingly dictated by technical affordances rather than democratic deliberation. Consequently, there is a growing risk that content moderation will be driven primarily by what LLMs can efficiently process, rather than by what best serves the diverse needs of online communities.

A broader concern is that LLM capabilities could become the dominant force shaping policy decisions, overshadowing thoughtful considerations of platform values, user needs, and societal impact. The difficulty of encoding complex ethical considerations into LLM-interpretable rules may push platforms towards simpler, more easily enforceable policies, resulting in the homogenisation of content moderation practices. If LLMs struggle with highly nuanced or context-sensitive rules, platforms may opt for standardised policies to ensure consistent enforcement. This trend recalls the Television Code era of the 1950s and 1960s, when standardised broadcast guidelines homogenised content across local stations, forcing diverse communities to conform to mainstream cultural norms rather than addressing their unique need [6, 35]. Similarly, LLM-driven content moderation risks prioritizing technological efficiency over the diverse and context-specific needs of global online communities, potentially undermining efforts to serve the varied requirements of different cultures and contexts.

5 Organisational Challenge: Converging Roles, Policy Authors and ML Practitioners

The integration of promptable LLMs into content moderation workflows has introduced a significant shift in roles and skill sets for both policy experts and machine learning practitioners. As LLMs are used to operationalise guidelines, policy authors are increasingly required to broaden their expertise beyond traditional policy development to include a working knowledge of machine learning principles, such as prompt engineering and the nuances of algorithmic behaviour, including bias, accuracy, and explainability [8]. This expanded skill set enables policy experts to anticipate potential issues, assess LLM performance, and contribute to the creation of more robust and fair moderation systems. Conversely, machine learning practitioners are stepping into domains traditionally governed by policy experts. By designing and refining the prompts that guide LLMs, they play a pivotal role in shaping how policies are interpreted and enforced. However, without formal training in policy development or regulatory frameworks, their contributions may inadvertently introduce systemic biases or misinterpretations that diverge from the policy's original intent and potentially overlook critical ethical considerations.

These emerging workflows are not simply technical translations of policy but complex “contested spaces” where different professional epistemologies and interpretive frameworks converge [20]. This intersection highlights the growing need for cross-disciplinary collaboration, where machine learning practitioners develop a deeper understanding of policy nuances, and policy experts acquire foundational machine learning literacy. In this evolving paradigm, content moderation would no longer rely on distinct, sequential contributions from policy experts and machine learning practitioners.

Instead, it would require hybrid expertise and continuous collaboration.

A notable example of this shift is the adoption of red teaming in LLM evaluation [15, 48]. Traditional evaluation methods, which rely on bespoke testing procedures, often fall short in capturing the full spectrum of behaviours and risks associated with increasingly autonomous and capable LLMs. Red teaming addresses this challenge by borrowing concepts from security practices, where deliberate attack strategies are employed to stress-test systems and uncover vulnerabilities. This approach requires not only technical expertise to understand and exploit the model's intricacies but also insights from fields such as social science, ethics, and policy to identify risks beyond technical performance, such as misuse or harmful societal impacts.

This shift underscores the need for hybrid knowledge systems that blend technical, ethical, and policy expertise. In the future, content moderation may give rise to roles such as “AI policy translators”—professionals skilled in bridging the gap between technical and policy teams. These individuals would play a pivotal role in ensuring that automated systems align with policy goals while leveraging the capabilities of LLMs. By fostering cross-disciplinary collaboration, they would contribute to the development of more robust, adaptable, and ethically sound moderation practices.

6 Governance Challenges

The rapid adoption of LLMs in automated decision-making systems has intensified discussions around model governance, focusing on issues such as transparency, accountability, and fairness in the deployment of AI-driven systems [9, 16, 42]. Rather than revisiting existing discussions on governance for automated decision-making systems, we specifically examine the implications unique to the policy-as-prompt configuration.

Decision attribution poses a significant accountability challenge in the policy-as-prompt setup. The lack of clear system boundaries create confusion when unwanted outcomes occur, making it difficult to pinpoint their exact cause—whether they stem from the policy language in the prompt, the LLM's interpretation of that language, or the complex interaction between the two. This ambiguity directly undermines auditability – reconstructing the decision-making process and documenting the influence of various components becomes incredibly difficult. The involvement of multiple stakeholders, including policy writers, prompt engineers, and model providers, further complicates attribution by obscuring how their contributions influence these outcomes [13].

A further challenge is determining which prompt adjustments warrant formal documentation. As discussed in Section 3, changes to enforcement can be made by modifying prompts, often with the semantic meaning of the underlying policies remaining unchanged. However, the sensitivity to prompt structure can lead to significant shifts in model behaviour, making it difficult to track how enforcement evolves over time. If every minor change requires extensive documentation, it adds unnecessary overhead. Yet, if changes are not sufficiently documented, it becomes increasingly difficult to monitor and understand the impact of these modifications on model performance.

7 Recommendations

In this section, we build on the systematic breakdown of challenges presented earlier (see Table 1) to propose strategic directions for addressing them. These strategies are informed by the insights gained from analysing the technical, sociotechnical, organisational and governance-related aspects of the challenges. Rather than offering a definitive set of solutions, this section emphasises high-level, actionable pathways that can guide future research and practical efforts. By framing these strategies within the context of our earlier analysis, we aim to provide a foundational perspective for tackling the complexities associated with this technology.

7.1 Enhanced Evaluation

A key mitigation strategy across several challenges is to do rigorous evaluation of the policy-as-prompt implementation. This evaluation must extend beyond traditional accuracy metrics to assess performance across critical dimensions.

To address *technical* sensitivity to prompt structure and formatting (see Section 3.2), the evaluation process should include comprehensive sensitivity analysis during the prompt development phase. This would involve studying the impact of formatting, phrasing, and structural changes on model outputs. Stress testing with diverse policy prompts ensures that small adjustments, such as punctuation or formatting changes, do not unintentionally lead to misinterpretation by the model. Sclar et al. [57] recommend reporting performance across a range of plausible prompt formats and styles rather than focusing solely on the performance of a single prompt format. Sensitivity analysis should also address subtle inconsistencies, such as those revealed by predictive multiplicity. For instance, Rashomon sets can identify cases where semantically similar policy phrasing elicit divergent model responses, such as inconsistent handling of edge cases, despite achieving comparable accuracy metrics [18, 37]. By analysing these variations, developers can identify policy formulations that are both effective and robust.

At the *sociotechnical* level, evaluation should move beyond conventional accuracy metrics to include measures of societal readiness and adaptability. Metrics like demographic fairness [39, 66] can help identifying whether specific prompt formulations create disparities across demographic groups, ensuring equitable policy application. Additionally, case libraries can serve as a valuable tool for addressing potential simplifying tendencies of LLMs. These libraries, containing nuanced, real-world examples of moderation edge cases—such as region-specific cultural references or satire—can test how well a policy-as-prompt system manages societal complexity. Specifically for the prompted guidelines, case libraries become crucial in revealing where binary or reductive prompt designs break down when confronted with complex real-world scenarios. Insights from this edge case analysis can guide iterative refinements of prompts, enabling policy makers and practitioners to design policy-as-prompt configurations that balance technical clarity with societal needs.

7.2 Collaborative Prompt Engineering

In Section 3.1 we discussed the complexity of converting policy guidelines to prompt, requiring careful attention to the ambiguities that can arise due to the inherent differences between human and

machine interpretability of the policy prompt. To minimise machine misinterpretations, strategies must focus on crafting prompts that address multi-faceted aspects of content and encourage diverse perspectives. Already techniques like chain-of-thought reasoning [68] have shown promise in crafting prompts that encourage step-by-step reasoning, allowing models to consider multiple facets of a problem before arriving at a conclusion.

Techniques like meta-prompting can dynamically integrate contextual ethical reasoning [62], and multi-persona prompting [67] can embed varied policy perspectives, enabling models to better navigate nuanced contexts. Additionally, fostering collaborative environments where LLMs—trained on diverse data and architectures—contribute to policy interpretation helps mitigate biases and refine prompt designs. These LLMs can collaboratively suggest rewrites and identify gaps creating a feedback loop for continuous improvement [33]. By treating prompt writing as an iterative, collaborative process of continuous refinement, practitioners can develop more nuanced and adaptive approaches to capturing the complex contextual understanding required in content moderation.

7.3 Traceability of Prompt Modifications

At the model *governance* level accountability concerns can be addressed by implementing mechanisms that enhance the transparency and traceability of prompt modifications. A “prompt genealogy” could record changes to prompt structures, formatting, and phrasing, along with the rationale behind each modification. Similar to version control systems used in software development [61] and drawing inspiration from tools like DVC (Data Version Control) [24] and Pachyderm [46], which provide versioning and lineage tracking for datasets and machine learning pipelines, a prompt genealogy would extend these principles to prompt engineering. For instance, practitioners could generate contextual logs or metadata documenting key aspects of the system’s behaviour, such as the inputs provided, the policy components referenced in the prompts, and the resulting outputs. While this approach does not reveal the internal workings of how the LLM processes or interprets prompts internally—a subject of ongoing research [59]—it offers a structured way to analyse how prompt changes or input variations correlate with differences in outcomes. This transparent record would also provide organisations with a comprehensible audit trail, supporting accountability requirements and reproducibility. It would also allow organisations to demonstrate alignment with intended policy goals and, where necessary, facilitate external reviews or assessments.

This mitigation strategy also aids in the management of *prompt sensitivity* (Section 3.2). By implementing a system for tracking changes to prompt structures, formatting, and phrasing, such as the prompt genealogy, it becomes possible to log and monitor how specific modifications affect model behaviour. Finally, a traceability tool can complement the sensitivity analysis idea mentioned in Section 7.1 by providing data for offline tests.

7.4 Bridging Organisational Silos

As discussed in Section 5, integrating policy authors and machine learning practitioners into a cohesive workflow is essential. However, the transition may encounter practical challenges, similar to historical precedents in other technological paradigm shifts. For

example, in the early 2000s the rise of data science in business required close collaboration between domain experts and data specialists [10]. Similarly, the emergence of bioinformatics in the late 20th century necessitated collaboration between biologists and computer scientists before the field gave rise to specialised bioinformaticians [40].

In the short-term, structured collaboration frameworks should be prioritised. These framework can include regular joint working sessions, shared documentation practices, and established feedback loops, ensuring that policy and machine learning teams work in tandem rather than sequentially. While the long-term goal may be the development of unified roles, these interim measures focus on creating better interfaces between existing specialist teams. This approach enables organisations to maintain effective content moderation while gradually building toward more integrated expertise.

8 Conclusions

In this paper, we examined the approach of incorporating policy text directly within LLM prompts, introducing “policy-as-prompt” as a conceptual framework. Our focus is the adoption of this approach for policy enforcement in content moderation. It represents an emerging trend [5, 23, 27, 36] with disruptive potential for content moderation systems. Our contribution is not in introducing a new method (i.e. prompting technique) but, rather, formally characterising this emerging framework, providing the first comprehensive analysis of its application in content moderation, generating novel insights into how policy integration transforms moderation practices, identifying challenges and suggesting mitigations. We address a significant gap in the literature, where the effects of natural language policy integration were previously underexplored or examined only in isolation. Studying this trend early allows us to contribute meaningfully to the evolving discussion about content moderation’s future. Lastly, given the above focus, a direct comparison with other works that focus on specific aspects of LLM prompting is outside the scope of this work.

We explored the complexities of the policy-as-prompt paradigm in content moderation through four dimensions: technical, sociotechnical, organizational, and governance. By presenting a structured overview of the challenges practitioners (would) face when transitioning from “traditional” machine learning or rule-based systems to the policy-as-prompt paradigm, we aim to support decision-making and underscore the importance of holistically considering the interplay between the technical capabilities of large language models, social and organizational contexts, and governance frameworks. Additionally, we propose potential strategies to mitigate some of these challenges, while acknowledging that many open questions remain regarding the implementation of effective content moderation systems under this paradigm.

Our discussion highlights challenges across these four dimensions without claiming to be exhaustive, and readers may identify additional issues beyond those covered in this paper. Furthermore, while we emphasise the transformative potential of the policy-as-prompt paradigm, it is important to note that such systems are not currently implemented in a fully autonomous manner. Instead, they typically operate within hybrid setups, where human moderators work alongside and oversee LLM-driven enforcement mechanisms.

However, even within this hybrid framework, the increasing adoption of the policy-as-prompt configuration introduces considerable complexities and potential risks that warrant further investigation. Addressing these challenges proactively will be critical to developing more effective hybrid moderation systems that thoughtfully integrate LLM capabilities with human oversight.

Empirically, we focused on one specific issue—model sensitivity to prompt structure—to demonstrate the brittleness of current LLMs are. We recognise, however, that model performance is influenced by a broader set of parameters beyond prompt structure, including the extent of pre-training, the choice of fine-tuning techniques, model architecture and size, and other factors that were fixed in our experiments. A more comprehensive analysis that examines how these parameters interact would yield deeper insights into their collective impact on model behaviour and reliability in content moderation tasks.

Going forward, as LLM capabilities continue to improve, the policy-as-prompt has the potential to overcome many of its current limitations, enabling organizations to dynamically align enforcement with nuanced and evolving policy objectives. Sustained exploration of these advancements will be essential in shaping the future of content moderation systems.

References

- [1] Nouar AlDahoul, Myles Joshua Toledo Tan, Harishwar Reddy Kasireddy, and Yasir Zaki. 2024. Advancing Content Moderation: Evaluating Large Language Models for Detecting Sensitive Content Across Text, Images, and Videos. arXiv:2411.17123 [cs.CV] <https://arxiv.org/abs/2411.17123>
- [2] Maryam Amirizani, Elias Martin, Maryna Sivachenko, Afra Mashhadi, and Chirag Shah. 2024. Do LLMs Exhibit Human-Like Reasoning? Evaluating Theory of Mind in LLMs for Open-Ended Responses. arXiv:2406.05659 [cs.CL] <https://arxiv.org/abs/2406.05659>
- [3] Anthropic. 2023. Introducing Claude. <https://www.anthropic.com/index/introducing-claude>. Accessed: January 9, 2025.
- [4] Anthropic. 2024. Prompt Engineering Overview. <https://docs.anthropic.com/en/docs/build-with-claude/prompt-engineering/overview>. Online; accessed 2024.
- [5] Anthropic. 2025. Use Case Guide: Content Moderation with Claude. <https://docs.anthropic.com/en/docs/about-claude/use-case-guides/content-moderation>. Accessed: 2025-01-11.
- [6] Erik Barnouw. 1970. *The Image Empire: A History of Broadcasting in the United States, Volume III-from 1953*. Oxford University Press, New York.
- [7] Mattia Belloni. 2023. Multilingual message content moderation at scale. <https://medium.com/bumble-tech/multilingual-message-content-moderation-at-scale-ddd0da1e23ed>. Accessed: January 9, 2025.
- [8] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kudithipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Muniyika, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs.LG] <https://arxiv.org/abs/2108.07258>

- [9] European Commission. 2020. White Paper on Artificial Intelligence - A European approach to excellence and trust. https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en.
- [10] Thomas H. Davenport and D.J. Patil. 2012. Data Scientist: The Sexiest Job of the 21st Century. *Harvard Business Review* 90, 10 (2012), 70–76. <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>
- [11] Akshar Prabhu Desai, Tejasvi Ravi, Mohammad Luqman, Mohit Sharma, Nithya Kota, and Pranjul Yadav. 2024. Gen-AI for User Safety: A Survey. arXiv:2411.06606 [cs.AI] <https://arxiv.org/abs/2411.06606>
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. doi:10.18653/v1/N19-1423
- [13] Nicholas Diakopoulos. 2016. Accountability in algorithmic decision making. *Commun. ACM* 59, 2 (Jan. 2016), 56–62. doi:10.1145/2844110
- [14] Evan Engstrom and Nick Feamster. 2017. *The Limits of Filtering: A Look at the Functionality & Shortcomings of Content Detection Tools*. Technical Report. Engine. <https://www.engine.is/the-limits-of-filtering>
- [15] Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, Andy Jones, Sam Bowman, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Nelson Elhage, Sheer El-Showk, Stanislav Fort, Zac Hatfield-Dodds, Tom Henighan, Danny Hernandez, Tristan Hume, Josh Jacobson, Scott Johnston, Shauna Kravec, Catherine Olsson, Sam Ringer, Eli Tran-Johnson, Dario Amodei, Tom Brown, Nicholas Joseph, Sam McCandlish, Chris Olah, Jared Kaplan, and Jack Clark. 2022. Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned. arXiv:2209.07858 [cs.CL] <https://arxiv.org/abs/2209.07858>
- [16] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. *Commun. ACM* 64, 12 (December 2021), 86–92. <https://www.microsoft.com/en-us/research/publication/datasheets-for-datasets/>
- [17] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. 1–288 pages. doi:10.12987/9780300235029
- [18] Juan Felipe Gomez, Caio Machado, Lucas Monteiro Paes, and Flavio Calmon. 2024. Algorithmic Arbitrariness in Content Moderation. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (Rio de Janeiro, Brazil) (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 2234–2253. doi:10.1145/3630106.3659036
- [19] Robert Gorwa, Reuben Binns, and Christian Katzenbach. 2020. Algorithmic content moderation: Technical and political challenges in the automation of platform governance. *Big Data & Society* 7, 1 (2020), 2053951719897945. doi:10.1177/2053951719897945 arXiv:<https://doi.org/10.1177/2053951719897945>
- [20] Ben Green. 2021. The Contestation of Tech Ethics: A Sociotechnical Approach to Ethics and Technology in Action. *SSRN Electronic Journal* (01 2021). doi:10.2139/ssrn.3859358
- [21] Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadiq Hasan. 2024. Does Prompt Formatting Have Any Impact on LLM Performance? arXiv:2411.10541 [cs.CL] <https://arxiv.org/abs/2411.10541>
- [22] John Howells. 1994. M. R. Smith and L. Marx (eds): "Does Technology Drive History? - The Dilemma of Technological Determinism. *Prometheus* (1994). Haves ikke HHÅ.
- [23] Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabza. 2023. Llama Guard: LLM-based Input-Output Safeguard for Human-AI Conversations. arXiv:2312.06674 [cs.CL]
- [24] Iterative.ai. 2020. Data Version Control (DVC): Open-source Version Control System for Machine Learning Projects. <https://dvc.org>.
- [25] Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan, Payel Das, and Siva Reddy. 2023. The Impact of Positional Encoding on Length Generalization in Transformers. In *Thirty-seventh Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=Drrl2gcjzl>
- [26] Mahi Kolla, Siddharth Salunkhe, Eshwar Chandrasekharan, and Koustuv Saha. 2024. LLM-Mod: Can Large Language Models Assist Content Moderation?. In *CHI 2024 - Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (Conference on Human Factors in Computing Systems - Proceedings)*. Association for Computing Machinery, United States. doi:10.1145/3613905.3658028 Publisher Copyright: © 2024 Association for Computing Machinery. All rights reserved.; 2024 CHI Conference on Human Factors in Computing Systems, CHI EA 2024 ; Conference date: 11-05-2024 Through 16-05-2024.
- [27] Deepak Kumar, Yousef Abubashem, and Zakir Durumeric. 2024. Watch Your Language: Investigating Content Moderation with Large Language Models. *Proceedings of the International AAAI Conference on Web and Social Media* 18 (05 2024), 865–878. doi:10.1609/icwsm.v18i1.31358
- [28] Savannah Larimore, Ian Kennedy, Breon Haskett, and Alina Arseniev-Koehler. 2021. Reconsidering Annotator Disagreement about Racist Language: Noise or Signal?. In *Proceedings of the Ninth International Workshop on Natural Language Processing for Social Media*, Lun-Wei Ku and Cheng-Te Li (Eds.). Association for Computational Linguistics, Online, 81–90. doi:10.18653/v1/2021.socialnlp-1.7
- [29] Lawrence Lessig. 1999. *Code and Other Laws of Cyberspace*. Basic Books.
- [30] Mosh Levy, Alon Jacoby, and Yoav Goldberg. 2024. Same Task, More Tokens: the Impact of Input Length on the Reasoning Performance of Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 15339–15353. doi:10.18653/v1/2024.acl-long.818
- [31] Chengzhengxu Li, Xiaoming Liu, Zhaohan Zhang, Yichen Wang, Chen Liu, Yu Lan, and Chao Shen. 2024. Concentrate Attention: Towards Domain-Generalizable Prompt Optimization for Language Models. arXiv:2406.10584 [cs.CL] <https://arxiv.org/abs/2406.10584>
- [32] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Hendersson, Qian Huang, Ryan Chi, Sang Michael Xie, Shihani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Ko-reeda. 2023. Holistic Evaluation of Language Models. arXiv:2211.09110 [cs.CL] <https://arxiv.org/abs/2211.09110>
- [33] Xuechen Liang, Meiling Tao, Yinghui Xia, Tianyu Shi, Jun Wang, and JingSong Yang. 2024. CMAT: A Multi-Agent Collaboration Tuning Framework for Enhancing Small Language Models. arXiv:2404.01663 [cs.CL] <https://arxiv.org/abs/2404.01663>
- [34] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173. doi:10.1162/tacl_a_00638
- [35] J. Fred MacDonald. 1990. *One Nation Under Television: The Rise and Decline of Network TV*. Pantheon Books, New York.
- [36] Todor Markov, Chong Zhang, Sandhini Agarwal, Florentine Eloundou Nekoul, Theodore Lee, Steven Adler, Angela Jiang, and Lilian Weng. 2023. A holistic approach to undesired content detection in the real world (AAAI'23/IAAI'23/EAAI'23). AAAI Press, Article 1683, 10 pages. doi:10.1609/aaai.v37i12.26752
- [37] Charles T. Marx, Flavio Du Pin Calmon, and Berk Ustun. 2020. Predictive multiplicity in classification. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. JMLR.org, Article 628, 10 pages.
- [38] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. <http://arxiv.org/abs/1301.3781>
- [39] Shira Mitchell, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. 2021. Algorithmic Fairness: Choices, Assumptions, and Definitions. *Annual Review of Statistics and Its Application* 8, 1 (March 2021), 141–163. doi:10.1146/annurev-statistics-042720-125902
- [40] Debasis Mitra, Debanjan Mitra, Mohamed Sabri Bensaad, Somya Sinha, Kumud Pant, Manu Pant, Ankita Priyadarshini, Pallavi Singh, Saliha Dassamiour, Leila Hambaba, Periyasamy Panneerselvam, and Pradeep K. Das Mohapatra. 2022. Evolution of bioinformatics and its impact on modern bio-science in the twenty-first century: Special attention to pharmacology, plant science and drug discovery. *Computational Toxicology* 24, Article 100248 (Nov. 2022), 100248 pages. doi:10.1016/j.comtox.2022.100248
- [41] Sankha Subhra Mullick, Mohan Bhamhani, Suhit Sinha, Akshat Mathur, Somya Gupta, and Jidnya Shah. 2023. Content Moderation for Evolving Policies using Binary Question Answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, Sunayana Sitaram, Beata Beigman Klebanov, and Jason D Williams (Eds.). Association for Computational Linguistics, Toronto, Canada, 561–573. doi:10.18653/v1/2023.acl-industry.54
- [42] OECD. 2019. Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- [43] OpenAI. 2023. GPT-4 Technical Report. <https://cdn.openai.com/papers/gpt-4.pdf>. Accessed: January 9, 2025.
- [44] OpenAI. 2024. GPT-4o Mini: Advancing Cost-Efficient Intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Accessed: 2025-01-13.
- [45] OpenAI. 2024. Using GPT-4 for Content Moderation. <https://openai.com/index/using-gpt-4-for-content-moderation/>. Accessed: 2025-01-11.

- [46] Inc. Pachyderm. 2021. Pachyderm: Version Control for Data Science and Machine Learning Pipelines. <https://www.pachyderm.com>.
- [47] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Alessandro Moschitti, Bo Pang, and Walter Daelemans (Eds.). Association for Computational Linguistics, Doha, Qatar, 1532–1543. doi:10.3115/v1/D14-1162
- [48] Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. 2022. Red Teaming Language Models with Language Models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3419–3448. doi:10.18653/v1/2022.emnlp-main.225
- [49] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Marilyn Walker, Heng Ji, and Amanda Stent (Eds.). Association for Computational Linguistics, New Orleans, Louisiana, 2227–2237. doi:10.18653/v1/N18-1202
- [50] Erich Prem and Brigitte Krenn. 2024. *On Algorithmic Content Moderation*. Springer Nature Switzerland, Cham, 481–493. doi:10.1007/978-3-031-45304-5_30
- [51] Felix Reda. 2017. When filters fail: These cases show we can't trust algorithms to clean up the internet. <https://felixreda.eu/2017/09/when-filters-fail/> Accessed: 2025-01-09.
- [52] Reddit, Inc. 2025. Reddit Content Policy. <https://redd.itinc.com/policies/content-policy> Accessed: 2025-01-11.
- [53] Sarah T. Roberts. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press. <http://www.jstor.org/stable/j.ctvhrzc20v>
- [54] Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications. arXiv:2402.07927 [cs.AI] <https://arxiv.org/abs/2402.07927>
- [55] Marta Sandri, Elisa Leonardelli, Sara Tonelli, and Elisabetta Jezek. 2023. Why Don't You Do It Right? Analysing Annotators' Disagreement in Subjective Tasks. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, Andreas Vlachos and Isabelle Augenstein (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 2428–2441. doi:10.18653/v1/2023.eacl-main.178
- [56] Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yinheng Li, Aayush Gupta, Hyojung Han, Sevien Schulhoff, Pranav Sandeep Dulepet, Saurav Vidyadhara, Dayeon Ki, Sweta Agrawal, Chau Pham, Gerson C. Kroiz, Feileen Li, Hudson Tao, Ashay Srivastava, Hevander Da Costa, Saloni Gupta, Megan L. Rogers, Inna Goncareenco, Giuseppe Sarli, Igor Galynker, Denis Peskoff, Marine Carpuat, Jules White, Shyamal Anadkat, Alexander Miserlis Hoyle, and Philip Resnik. 2024. The Prompt Report: A Systematic Survey of Prompting Techniques. *CoRR* abs/2406.06608 (2024). <https://doi.org/10.48550/arXiv.2406.06608>
- [57] Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=Rlu5lyNXjT>
- [58] Jan Simson, Florian Pfisterer, and Christoph Kern. 2024. One Model Many Scores: Using Multiverse Analysis to Prevent Fairness Hacking and Evaluate the Influence of Model Design Decisions. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (Rio de Janeiro, Brazil) (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 1305–1320. doi:10.1145/3630106.3658974
- [59] Chandan Singh, Jeevana Priya Inala, Michel Galley, Rich Caruana, and Jianfeng Gao. 2024. Rethinking Interpretability in the Era of Large Language Models. *arXiv preprint arXiv:2402.01761* (2024).
- [60] Mohit Singhal, Chen Ling, Nihal Kumarswamy, Gianluca Stringhini, and Shirin Nilizadeh. 2023. SoK: Content Moderation in Social Media, from Guidelines to Enforcement, and Research to Practice. *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)* (2023), 868–895.
- [61] Diomidis Spinellis. 2005. Version control systems. *IEEE Software* 22, 5 (2005), 108–109. doi:10.1109/MS.2005.145
- [62] Mirac Suzgun and Adam Tauman Kalai. 2024. Meta-Prompting: Enhancing Language Models with Task-Agnostic Scaffolding. arXiv:2401.12954 [cs.CL] <https://arxiv.org/abs/2401.12954>
- [63] Hugo Touvron, Louis Martin, Kevin Stone, Patrick Albert, Amjad Almahairi, Yasaman Babaei, ..., and Daniel Jurafsky. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv preprint arXiv:2307.09288* (2023).
- [64] José van Dijk. 2013. *The Culture of Connectivity: A Critical History of Social Media*. Oxford University Press, Oxford.
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [66] Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*. IEEE, 1–7.
- [67] Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2024. Unleashing the Emergent Cognitive Synergy in Large Language Models: A Task-Solving Agent through Multi-Persona Self-Collaboration. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 257–279. doi:10.18653/v1/2024.naacl-long.15
- [68] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '22)*. Red Hook, NY, USA, Article 1800, 14 pages.
- [69] Langdon Winner. 1980. Do Artifacts Have Politics? *Daedalus* 109, 1 (1980), 121–136.
- [70] Blaise Agüera y Arcas. 2022. Do Large Language Models Understand Us? *Daedalus* 151, 2 (05 2022), 183–197. doi:10.1162/daed_a_01909
- [71] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, and Karthik Narasimhan. 2023. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*.
- [72] Karen Yeung. 2019. Algorithmic Regulation: A Critical Interrogation. *Regulation & Governance* 13, 1 (2019), 94–114. doi:10.1111/rego.12158
- [73] Meiru Zhang, Zaiqiao Meng, and Nigel Collier. 2024. Attention Instruction: Amplifying Attention in the Middle via Prompting. arXiv:2406.17095 [cs.CL] <https://arxiv.org/abs/2406.17095>




Figure 6: Key challenges (outer nodes) categorized by area in the implementation of ‘policy-as-prompt’.

A Experiments

In this section we provide some details supportive to the experimental subsection 3.2 in the main paper.

A.1 Prompt style variations

We observe that model performance, as measured by the F-beta score, fluctuates across textual prompt variations (see Figure 7 for an example of the template variations and Figure 8a for the results). Categories such as Not Violating and Self-Harm maintain relatively high and stable scores, while more contentious categories like Jailbreak, Harm, and Illegal Goods display marked sensitivity. This disparity highlights that model performance is prone to degradation when subjected to structural changes in the policy prompt. This directly supports the thesis of Section 3.2, highlighting the role of prompt design in shaping moderation outcomes. The variations in F-beta scores (Figure 8a) and the high standard deviation (Figure 8b) demonstrate predictive multiplicity in action. Even though the overall performance across prompt variations may appear statistically comparable, the conflicting behaviours observed in key categories (e.g., Jailbreak, Harm) point to a deeper underlying issue: different prompt structures lead to divergent model predictions for the same input samples.

A.2 Policy Evolution Snapshots

To examine the role of policy specificity, we created temporal snapshots of policy descriptions, starting with minimal drafts (one-line descriptions) and progressively adding detail to arrive at comprehensive guidelines (as seen in Table 2 reversely).

A.3 Model Size

To gain insight how models with different size respond to variations of the prompt, we ran experiments using models with different size

from different providers. In Figure 9a we plot the performance distribution and use ‘P#’ to indicate different model provider series. Middle sized models are between the 32B and the 64B parameters mark, while big ones surpass the 100B. Overall, while we observe that the sensitivity to prompt structure and style decreases with model size, it is still present but to a lesser degree. We also observe that different providers - thus different pre-training strategies- affect the model sensitivity. Still, even for the most capable models, predictive multiplicity arises as can be seen in 9b, where the variants *Structured* and *Annotator-Friendly* present significant performance divergences for certain categories, despite negligible differences in averaged accuracy. For example, accuracy scores for Hate are around five points higher for the *Annotator-Friendly* variation.

| | |
|-----------------------------|--|
| <p>1. Punctuated</p> | <p>Prompts that incite violence or hatred towards a person or group, of people based on race religion gender identity or expression sex, ethnicity, nationality sexual orientation veteran, status age disability or other characteristics associated with systemic discrimination or marginalization includes but may not be limited to: praising supporting, or calling for violence or exclusion, against a person or group of people based on the characteristics listed above, dehumanizing statements...</p> |
| <p>2. Structured</p> | <p>- Prompts that incite violence or hatred towards individuals or groups are prohibited. - These prompts target characteristics such as:</p> <ul style="list-style-type: none"> • race • religion • gender identity or expression • sex • ethnicity • nationality • sexual orientation • veteran status • age • disability • other characteristics linked with systemic discrimination or marginalization <p>- Prohibited actions include, but are not limited to:</p> <ul style="list-style-type: none"> • Praising, supporting, or calling for violence or exclusion against individuals or groups based on the aforementioned characteristics. • Making dehumanizing statements... |
| <p>3. Concise</p> | <p>Policy that prohibits content inciting violence or hatred on grounds of race, religion, gender identity, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability, or other marginalized traits includes, but is not limited to: endorsing or advocating violence or exclusion, dehumanizing statements...</p> |
| <p>4. Verbose</p> | <p>The policy is dedicated to addressing and regulating prompts that have the potential to incite or provoke acts of violence or expressions of hatred directed toward an individual or a group of individuals. This regulation specifically targets instances where such behavior is based on categorical aspects such as race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability, or other attributes which are commonly linked to systemic discrimination or marginalization. The policy encompasses, but is not strictly limited to, behaviors such as: endorsing, supporting, or advocating for acts of violence or exclusion toward an individual or groups of individuals based on any of the aforementioned characteristics, making statements that dehumanize an...</p> |
| <p>5. Annotator</p> | <p>If a message encourages violence or hatred towards someone because of their race, religion, gender identity, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability, or other similar traits, it's not allowed. This includes actions like praising, supporting, or demanding violence or exclusion of these individuals. It also involves saying dehumanizing things...</p> |

Figure 7: The illustration of different prompt styles for hate speech policy.

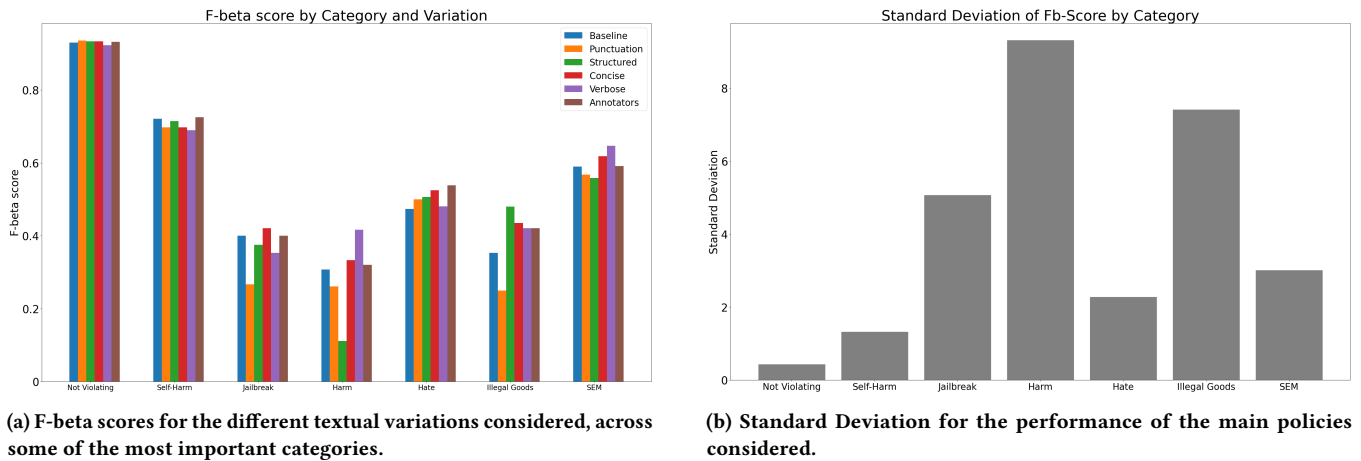


Figure 8: Comparison of performance metrics across categories and their standard deviations.

Prompts that incite violence or hatred towards a person or group of people based on race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability or other characteristics associated with systemic discrimination or marginalization includes, but may not be limited to: - praising, supporting, or calling for violence or exclusion against a person or group of people based on the characteristics listed above - dehumanizing statements about a person or group based on the protected characteristics listed above - using hateful slurs when targeting someone/a group of people based on their protected characteristic promoting or glorifying hate groups/hate figures and their associated images, and/or symbols.

Prompts that incite violence or hatred towards a person or group of people based on race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability or other characteristics associated with systemic discrimination or marginalization includes, but may not be limited to: - praising, supporting, or calling for violence or exclusion against a person or group of people based on the characteristics listed above - dehumanizing statements about a person or group based on the protected characteristics listed above.

Prompts that incite violence or hatred towards a person or group of people based on race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability or other characteristics associated with systemic discrimination or marginalization includes, but may not be limited to: - praising, supporting, or calling for violence against a person or group of people based on the characteristics listed above.

Prompts that incite violence or hatred towards a person or group of people based on race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability or other characteristics associated with systemic discrimination or marginalization includes using hateful slurs when targeting someone/a group of people based on their protected characteristic.

Prompts that incite violence or hatred towards a person or group of people based on race, religion, gender identity or expression, sex, ethnicity, nationality, sexual orientation, veteran status, age, disability or other characteristics associated with systemic discrimination or marginalization.

Snapshot 0

Snapshot 1

Snapshot 2

Snapshot 3

Snapshot 4

Table 2: Different temporal snapshots for the evolution of policy category: Hate.

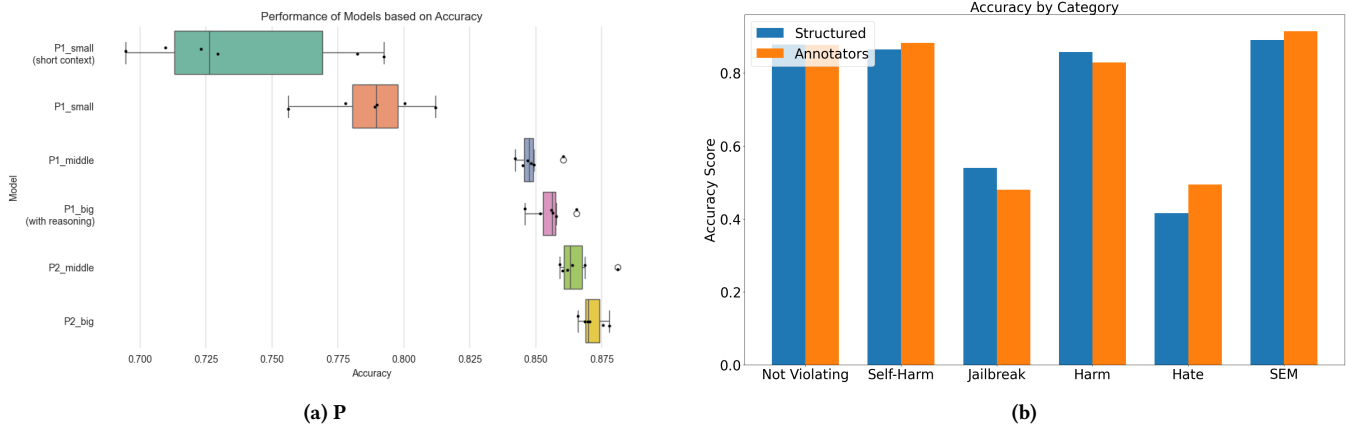


Figure 9: Analysing the effect of model provider and size in sensitivity and predictive multiplicity. a) The performance distribution of different model sizes for the six prompt template variations. b) Accuracy across policy categories for the ‘Structured’ and ‘Annotators’ prompt types.