

Debunking with Dialogue? Exploring AI-Generated Counterspeech to Challenge Conspiracy Theories

Mareike Lisker and Christina Gottschalk and Helena Mihaljević

University of Applied Sciences (HTW) Berlin, Germany

{lisker, mihalje}@htw-berlin.de

Abstract

Counterspeech is a key strategy against harmful online content, but scaling expert-driven efforts is challenging. Large Language Models (LLMs) present a potential solution, though their use in countering conspiracy theories is under-researched. Unlike for hate speech, no datasets exist that pair conspiracy theory comments with expert-crafted counterspeech. We address this gap by evaluating the ability of GPT-4o, Llama 3, and Mistral to effectively apply counterspeech strategies derived from psychological research provided through structured prompts. Our results show that the models often generate generic, repetitive, or superficial results. Additionally, they over-acknowledge fear and frequently make up facts, sources, or figures, making their prompt-based use in practical applications problematic.

1 Introduction

Conspiracy theories (CTs) are ubiquitous, often emerging as attempts to identify underlying causes of unexpected or significant events. The resulting narratives allege clandestine machinations by powerful entities perceived as malevolent. CTs are emotionally charged and compelling, forming closed belief systems characterized by internal coherence, which makes them difficult to refute (Lepoutre, 2024). They are often associated with harmful social, health and political consequences (Douglas et al., 2019). For example, CTs can lead individuals to reject scientific consensus (Weigmann, 2018) and influence societal attitudes about critical issues, such as climate change or vaccine policies (Jolley and Douglas, 2014). Belief in CTs can undermine trust in democratic institutions, amplify social tensions and lead to violence (Vergani et al., 2022).

Counterspeech (CS) has emerged as an important strategy for mitigating the impact of harmful content in the online sphere. It describes communicative measures designed to refute undesirable

content while respecting the principle of freedom of expression (Schieb and Preuss, 2016). Many non-governmental organizations (NGOs) employ it as a strategy to combat hate speech (HS), misinformation and CTs (Rathje et al., 2015; UNESCO, 2022). While shifting the beliefs of an ideologically committed user who is propagating CTs might be rather difficult (Krekó, 2020), counterspeech can positively affect the discourse norms of the audience (Bojarska, 2018).

Manual moderation for comprehensive counterspeech is challenging due to high volumes and the rapidly evolving social media landscape (Bonaldi et al., 2024a). In this context, supporting NGOs by automatically generating counterspeech that requires little post-processing becomes increasingly vital (Mun et al., 2024). Large Language Models (LLMs) could potentially assist (Chung et al., 2019). While research has explored using LLMs for counterspeech against HS, there is a notable gap in addressing CTs. While CTs can incite hatred and often perpetuate antisemitic or racist narratives (Hay, 2020), they typically differ linguistically from HS.

This research gap motivates our study, which investigates the potential of LLMs to generate counterspeech for CTs. We focus on the scenario of an NGO employee leveraging LLM-generated suggestions to respond to comments promoting CTs. Since no expert-based datasets containing examples of counterspeech to CTs could be identified, we rely on zero-shot prompting to evaluate the LLMs. Our central research question is thus:

(RQ) To what extent can current LLMs guided by prompt-based strategy instructions generate effective counterspeech against conspiracy theories?

The prompts are based on four strategies drawn from the literature, particularly in psychology: **fact-check based refutation**, **providing alternative explanations**, **storytelling of a counter-narrative**, and **encouraging critical thinking**. These strategies are

considered promising for countering CT beliefs in audiences not yet fully absorbed by them. We applied three LLMs—GPT-4o, Llama 3 (8B), and Mistral (7B)—to a dataset of 152 comments from the platform X that promote two CT themes: (1) narratives involving the ‘deep state’, ‘NWO’ and ‘globalists’, and (2) claims related to ‘geo- and bio-engineering’. The generated CS is qualitatively evaluated according to the established criteria. In addition to assessing individual comments, we analyze the lexical and semantic diversity of the generated CS to evaluate its practical utility for civil society actors.

Our findings show that all three models struggled to produce CS with substantial depth, often defaulting to superficial statements and avoiding narrative storytelling as a strategy. Responses frequently contained generic and repetitive phrases, resulting in limited linguistic variety. Notably, approximately 10% of outputs contained confabulations which were often difficult to spot. Differences between the models were low and mostly not significant. These findings suggest that current LLMs are not yet effective for the generation of CS to CTs in a zero-shot setting, underscoring the need for further research on grounding LLMs in relevant domain knowledge.

2 Related Work

Research on automated CS generation has gained traction, but has primarily focused on HS. To our knowledge, only (Costello et al., 2024) have examined the use of LLMs to generate CS for CTs. In lab experiments, participants engaged in multi-turn dialogues with GPT-4 Turbo, resulting in a statistically significant reduction in self-reported CT belief. This suggests general potential for CT belief revision through LLM interaction, though in a direct, dialogue-based setting differing from the social media context.

We found no publicly available datasets containing examples of effective CS to CTs communicated online. Existing datasets on online CTs (e. g., (Langguth et al., 2023; Steffen et al., 2023)) were mainly created for classification tasks (Liu et al., 2024b; Peskine et al., 2023; Pustet et al., 2024), and are mostly related to COVID-19.

In the remainder, we focus on LLM-based CS generation in the context of HS, as we found no prior work addressing these aspects for CTs.

The majority of approaches in CS generation

rely on fine-tuning methods (Qian et al., 2019; Tekiroğlu et al., 2022), using datasets from platforms such as Reddit and Gab (Qian et al., 2019), compiled by NGO operators (Chung et al., 2019), or generated through models (Bonaldi et al., 2022). More recently, zero-shot and few-shot prompting of various GPT-based models has shown promising results for generating CS against HS (Chung et al., 2021; Ashida and Komachi, 2022; Zheng et al., 2023; Halim et al., 2023). Some studies have explored tailored approaches, such as using type-prompts to specify predefined styles like humor or facts (Saha et al., 2024) or personalizing CS based on author profiling (e. g., age and gender) (Doğanç and Markov, 2023) to fit specific contexts. Across these studies, the newer GPT models demonstrate strong performance.

Some studies have integrated automated CS generation into interactive tools. For instance, generation of contextualized CS was embedded within a content monitoring tool designed to address GDPR and DSA compliance (Fillies et al., 2024), or incorporated into a dashboard enabling users to semi-automatically generate CS using text and memes (Smedt et al., 2021).

Bonaldi et al. present a comprehensive guide to using NLP to generate CS. They emphasize that while research in this area is growing, there is currently a lack of standardized definitions and best practices to integrate the different approaches (Bonaldi et al., 2024a).

Ashida and Komachi highlight that evaluating CS is challenging due to the lack of established standards (Ashida and Komachi, 2022). Some studies use metrics like BLEU, GLEU, or METEOR to compare generated text to a ‘ground truth’. However, without CS references, these metrics are not applicable. Moreover, unlike machine translation, effective CS may include a wide range of diverse responses, making such metrics less suitable. Nonetheless, these measures can assess the diversity and novelty of outputs (Saha et al., 2024). Furthermore, classification models can be utilized to evaluate sentiment or argument characteristics (Saha et al., 2024).

In addition to quantitative methods, most researchers rely on human annotators for more in-depth evaluations with regard to ‘informativeness’ and ‘offensiveness’ (Ashida and Komachi, 2022), or ‘suitableness’, and ‘intra-coherence’ (Chung et al., 2021), ‘grammaticality’, ‘specificity’ (Tekiroğlu et al., 2022), or ‘diversity’, ‘relevance’

and ‘language quality’ (Zhu and Bhat, 2021).

3 Counterspeech Against Conspiracy Theories

Conspiracy theories are interpretative frameworks attempting to explain events as the covert actions of powerful, malicious groups or individuals pursuing self-interest at public expense (Aaronovitch, 2010; Byford, 2011; Keeley, 1999), overlooking actual, more intricate causes (Popper and Kiesewetter, 2003). They thus usually need actors (e. g., ‘globalists’) with a malicious goal (e. g., destabilizing national economies) through an action or strategy (e. g., manipulating weather) carried out in secrecy (Samory and Mitra, 2018). On social media, these narratives often appear in fragmented forms, with some components omitted or implied (Steffen et al., 2023). While CTs can overlap with HS (Baider, 2023), they have their own unique characteristics with regard to function and linguistic manifestation (Hay, 2020).

Interventions against CTs can be categorized as preventive (pre) or harm-reducing (post), as well as targeting the sender or the audience (Krekó, 2020). As individuals deeply invested in conspiracy beliefs are particularly resistant to change (Krekó, 2020), we focus on addressing the recipients of CTs, or ‘cyber-bystanders’, in a social media context, post-exposure to a CT.

Study findings show that “there is no such thing as perfectly refuted misinformation” (Krekó, 2020, p. 4), since part of the ‘essence’ of the conspiracy remains embedded in the mental model of an individual. Nevertheless, research suggests that the design and formulation of CS can influence its impact (Chung et al., 2023).

4 Data and Methods

4.1 Dataset

As we target an audience not (yet) fully absorbed by conspiracy narratives, we sourced data from X (formerly Twitter), where exposure to CT content may reach broader, more diverse audiences. Our focus was on messages containing key elements of conspiracy narratives, such as identifiable actors, strategies, or goals, ensuring that the intended audience would likely grasp the central idea of the CT. Additionally, we prioritized messages with at least 200 likes or replies, as these are more likely to represent content with a wide reach.

The data was collected using keywords associated with two current and rather well-known CT themes, which can be broadly characterized as hate-based and fear-driven, namely theories revolving around ‘deep state’, ‘NWO’ and ‘globalists’ on the hate-based side, and those relating to ‘geo- and bioengineering’ on the fear-based side. The two sets of keywords (see Table 1) were iteratively enriched following an exploration of conspiratorial posts and hashtags on X. Note that we did not aim to create an exhaustive set of keywords, as we are mainly interested in identifying texts that can be clearly categorized as conspiratorial by an LLM. In total, we collected 152 posts (76 per category). All data and methods are available in the project repository.¹

Table 1: Keywords used to collect CT messages.

CT Theme	Keywords
Hate-based	deepstate, NWO, WW3, QAnon, globalists
Fear-driven	geo-engineered, Big Pharma, bio-engineered, HAARP, chemtrails, vaccine RFK Jr.

4.2 Model Overview

We prompted *GPT-4o* (OpenAI, 2024) and two open-source models, *Mistral* (Mistral-7B-Instruct-v0.3, (Mistral, 2024)) and *Llama 3* (Meta-Llama-3-8B-Instruct, (AI@Meta, 2024)). In multiple previous studies, *GPT-4o* outperformed other models (cf., e. g. (Doğanç and Markov, 2023; Ashida and Komachi, 2022)), and thus serves as a high-performance but also closed-source and cost-intensive solution. The smaller open models were selected for their accessibility and lower operational costs, making them more viable for practical applications in civic society settings.

We accessed the open models via Hugging Face and *GPT-4o* through OpenAI’s API. We set temperature and top_p according to default or recommended values provided in the respective model documentations, ensuring optimal performance under consistent conditions.² All experiments were conducted in October and November 2024.

¹https://github.com/HTW-Social-Data-Science/Debunking_with_Dialogue

²For Llama 3, temperature was set to 0.6 and top_p to 0.9, while both hyperparameters were set to 1 for the other two models.

4.3 Prompt Design

We opted for zero-shot prompting, including role and style prompting, as well as step-by-step instructions, as recommended by LLM providers and recent research (cf., e. g., (Schulhoff et al., 2024)). The instructions primarily reflect four strategies for countering CTs that, according to current literature, are likely to enhance the effectiveness of CS: **fact-check based refutation (Fact)**, **providing alternative explanations (Alt)**, **storytelling of a counter-narrative (Narr)**, and **encouraging critical thinking (Crit)**. While **Fact** involves correcting misinformation through fact-based rebuttals, **Alt** emphasizes providing more holistic alternative explanations. **Narr** uses narrative techniques to engage the audience emotionally, whereas **Crit** aims at revealing the manipulative techniques used in conspiracy theories and encouraging reflective questioning.

In the following, we illustrate the different counterspeech strategies by sample answers to an example CT. For the theoretical underpinning of the prompt design, see Appendix A.

CT: “COVID-19 was created as a bio-weapon in a lab to control the population.”

Fact: “Independent scientific bodies, including the World Health Organization and international virology experts, have found no evidence that COVID-19 was engineered or intentionally released, as noted in a 2020 article by Kristian G. Andersen in Nature Medicine. Genetic analyses indicate natural evolution, and similar coronaviruses found in animals, especially bats, support a zoonotic origin as stated in Wikipedia.”

Alt: “Viruses like COVID-19 have emerged before through natural processes. Scientists have found that the virus shares much of its genetic code with coronaviruses found in bats and other animals, suggesting it likely jumped from animals to humans—something that has happened before with SARS and MERS. In a world where people and animals are increasingly in close contact, and global travel is constant, it’s more likely that this virus emerged through natural transmission rather than being engineered.”

Narr: “In the 1990s, Dr. Lam discovered a virus that had jumped from bats to humans in Malaysia. When COVID-19 emerged, some suspected a lab origin, but Dr. Lam knew that the genetic code of the virus told a different story: it closely matched viruses found in bats, showing a natural origin. He explained that as humans encroach on wildlife habi-

tats, viruses can spill over, and that this is what likely happened with COVID-19.”

Crit: “Mhh, but if COVID-19 was a bio-weapon designed for control, why would it spread uncontrollably across all countries, including those allegedly behind it? Would a government really create something that harms its own economy, overwhelms its hospitals, and puts its own leaders at risk?”

The prompt formulation was iteratively adjusted using a sample of messages and GPT-4o. The limit of 200 token was adjusted based on experiments with Llama 3 and Mistral as these had more difficulties with restricting to a smaller limit. See Appendix B for the final version of the prompt.

4.4 Counterspeech Annotation Framework

Since comparing model output with ‘correct’ counterspeech through quantitative metrics such as BLEU or METEOR cannot be employed, we rely on manual annotation of model outputs based on the prompt instructions.

We annotated a total of 456 responses across 12 criteria, each reflecting a specific aspect of the prompt instructions. The criteria are color-coded to correspond to the respective step in the prompt (see Appendix B) as well as the respective literature they were derived from (see A):

Table 2: Annotation criteria by prompt step and counterspeech strategy.

Step / Strategy	Annotation Criteria
Clarity & Restraint	Conciseness and clarity Repetition of false or harmful content Use of stigmatizing language
Hate Speech	Hate speech identification Hate speech condemnation CT avoidance in hateful context
Fear & Empathy	Fear identification Appropriate empathy for fear
Strategies	Fact : Refute based on fact-checks Alt : Provide alternative explanations Narr : Counter conspiracy with narrative Crit : Encourage critical thinking

Three annotation criteria were treated as binary variables: *Hate speech identification*, *Fear identification*, and *Use of stigmatizing language*. The first two were coded based on a comparison between the original message and the model output. The latter indicates the presence or absence of the terms ‘conspiracy (theory)’, ‘misinformation’, ‘debunk-’ and ‘unfounded’. The remaining 9 criteria were evaluated using Likert scales, which are available

in the project repository.

To evaluate the annotation schema, a random sample of 50 responses—presented without any information about the respective models—was annotated independently by two authors. The inter-annotator agreement, measured using Krippendorff’s Alpha, ranged from 0.35 (*Conciseness and clarity*) to 0.83 (*Hate speech condemnation and Appropriate empathy for fear*). Disagreements were discussed and resolved collaboratively, and the annotation schema was refined accordingly. The lower agreement in case of (*Conciseness and clarity*) was due to the initially underspecified criterion. In the subsequent discussion, we defined what to consider as ‘non-clear’ or ‘non-concise’, such as an unnecessary intro or outro, excessive mentioning of the CT, or patronizing formulations such as “It’s important to keep in mind ...”. The remaining data was split equally and annotated by one author each in close communication, with $\sim 25\%$ of the records discussed jointly. Likert scales were iteratively refined, and prior annotations were retrospectively updated to ensure consistency.

4.5 Comparison of Models

We quantitatively compared how the models implemented CS strategies based on the annotated Likert scale ratings. The only post-processing applied was separating the meta report from the actual CS via the <XXX>-tag. Due to the non-normal distribution of the values, we used the non-parametric Friedman test ($\alpha = 0.05$) to compare all three models. For significant results, we applied the post-hoc Wilcoxon Signed-Rank test to identify model pairs with significant differences. Effect sizes were measured using Kendall’s W. Comparisons were limited to records with non-missing values for all models. Bonferroni correction was applied to adjust p-values and control the family-wise error rate.

Furthermore, we assessed the response diversity using various lexical and semantic measures, as repetitive outputs would be less useful in the practical setting guiding our study.

5 Results

We first evaluate how well the models followed core prompt instructions related to structure, tone, and issue-specific handling.

5.1 Adherence to Prompt Instructions

Meta Report and Length Restriction. The models showed varying adherence to simple prompt in-

structions, such as listing how hate speech was handled and which strategies were applied (the ‘meta report’), separating this information from the CS using the <XXX> tag, and keeping the CS under 200 tokens. GPT-4o consistently included the meta report, while Llama 3 and Mistral omitted it in 26–28% of outputs (see col. 1 in Table 3). The meta report in GPT-4o and Llama 3 responses was short, accounting for $\sim 1\%$ and $\sim 4\%$ of the output, respectively, while Mistral produced substantially longer explanations ($> 20\%$ on average). GPT-4o correctly separated the meta report from the generated CS in over 96% of outputs, enabling easy rule-based parsing. In contrast, Llama 3 responses required manual post-processing, i. e., separating the meta report from the CS, in 55% of cases and Mistral failed to use the tag <XXX>, necessitating manual post-processing for all outputs. Mistral reported on HS handling in only 25% of outputs, compared to 56% for GPT-4o and 68% for Llama 3 (see col. 3 in Table 3). Interestingly, all models frequently misinterpreted the requirement to acknowledge fear—when present—as a CS strategy, often labeling it as ‘Strategy 5,’ thus indicating confusion about the prompt instructions.

Table 3: Compliance with prompt instructions: (1) inclusion of a meta report, (2) correct CS separation, (3) HS handling specified, and (4) mean CS length.

model	Meta report (%)	CS sep. (%)	HS handling (%)	Mean length
GPT-4o	100.00	96.05	55.92	709.38
Llama 3	71.71	55.05	68.22	1267.41
Mistral	74.34	0.00	25.66	794.74

While Mistral and GPT-4o achieved an average character length of less than 800, thus mostly obeying the instruction to stay below 200 token, Llama 3 responses exceeded the limit with more than 1200 characters on average (see col. 4 in Table 3).

Clarity & Restraint. GPT-4o significantly outperformed Llama 3 in clarity and conciseness, scoring 3.48 versus 3.12 on the Likert scale (effect size of 0.62). Mistral, with a mean clarity score of 3.36 lying in between, showed the greatest variability ($SD=1.1$) in performance.

GPT-4o showed the highest compliance with instructions to avoid terms like ‘conspiracy theory’ and ‘misinformation,’ using them in only 7.23% of outputs, compared to 38.15% for Mistral and 79.6% for Llama 3. It also adhered more closely to the rule against repeating harmful content, with an

average Likert score of 3.78, significantly outperforming Mistral (2.82) and Llama 3 (2.48). Mistral and Llama 3 frequently reinforced harmful content by including hashtags like ‘#globalists’, ‘#Deep-StateAgenda’ or ‘#QAnon’, effectively promoting the narratives they were meant to counter. Notably, a score of 4 is already considered high, as effectively refuting false claims without repeating any of them poses a significant challenge—even for humans.

Hate Speech. The dataset included 17 messages containing explicit or implicit HS that the models were supposed to recognize and condemn without engaging with the CT content. GPT-4o and Llama 3 achieved F1 scores of 0.75 and 0.69, respectively, in detection, while Mistral scored only 0.30 due to poor recall (differences not significant per McNemar’s test). Overall condemnation of HS yielded mean Likert scores between 1.18 (Mistral) and 2.1 (Llama 3), reflecting limited effectiveness. Models often failed to address HS adequately, either ignoring it or responding too empathically and friendly. When HS was detected, GPT-4o more consistently avoided engagement with CTs, scoring 3.55 on the Likert scale compared to 1.62 and 1.75 for Llama 3 and Mistral, respectively.

Fear & Empathy. All models over-acknowledged fear and anxiety, responding as if these were the primary motivations in 26% (Llama 3), 31% (Mistral), and even 52% (GPT-4o) of cases—in contrast to annotators who identified such motivations in fewer than 5% of messages. One meta report noted that empathy in the response was only optional, as the comment was not based on fear [262]³. Sentiment analysis using Vader showed a positive tone in 108 out of 152 cases for Llama 3 and Mistral, and in 142 for GPT-4o.

Overall, acknowledgments were often superficial, frequently appearing in introductory remarks such as, “It’s natural to feel concerned about [...]. However, [...]” Consequently, average scores were low, from 1.9 ± 0.86 (Mistral) to 1.95 ± 0.65 (GPT-4o) and 2.15 ± 0.58 (Llama 3), with no statistically significant differences. In some cases, acknowledging fear reinforced a conspiratorial mindset, e. g., “It’s natural to feel concerned when leaders address global issues [...]” [364].

³The number in the box depicts the corresponding ID of the response. See the repository for all responses.

5.2 Application of Counterspeech Strategies

According to the models’ meta reports (see Table 4), Strategy *Crit* was the most frequently applied, occurring in 86.25% of outputs. *Fact* was the second most common, closely followed by *Alt*, while *Narr* was applied in only a handful of outputs.

Table 4: Proportion of outputs in which a model reported applying a given counterspeech strategy.

Model \ Strategy	<i>Fact</i>	<i>Alt</i>	<i>Narr</i>	<i>Crit</i>
GPT-4o	63.16	89.47	4.61	93.42
Llama 3	93.52	75.93	1.85	73.15
Mistral	82.88	48.65	5.41	89.19

Models exhibited preferences for specific strategies: Llama 3 most frequently employed *Fact*, while GPT-4o and Mistral favored *Crit*. Notably, the self-reported strategies often misaligned with human annotations, matching fully in only 40-42% of cases. This inconsistency suggests limitations in the models’ ‘understanding’ of the strategies.

Table 5 shows that the quality of CS for *Fact* and *Alt* was moderate across all models, with average scores hovering around 3 or slightly below (scores for *Narr* were omitted due to its infrequent use and limited success.) While the factual accuracy of these strategies was generally acceptable, the elaborations often lacked depth. Critical thinking was even less effectively encouraged, with mean scores ranging from 2.34 (Llama 3) to 2.43 (Mistral). Llama 3 significantly outperformed Mistral with regard to *Fact*, with a large effect size (0.57). However, differences in performance across other strategies were not statistically significant.

Table 5: Model performance across counterspeech strategies, presented as mean scores (\pm SD).

Model	<i>Fact</i>	<i>Alt</i>	<i>Crit</i>
GPT-4o	2.84(± 0.82)	2.95(± 0.84)	2.41(0.76)
Llama 3	3.04(± 1.16)	2.93(± 1.19)	2.34(± 0.73)
Mistral	2.75(± 0.93)	2.85(± 0.99)	2.43(± 0.7)

Refute based on Fact-Checks. In $\sim 10\%$ of the fact-checked responses, models declared true events to be false, made up facts, or produced ‘bullshit’ (Frankfurt, 2009), such as: “However, it’s important to note that the term ‘globalist’ is often used as a euphemism for ‘those who want to help the world’ or ‘those who believe in international cooperation.’” [55] Often, the cited sources

were untraceable despite sounding convincing e. g., [150, 217, 275]. Regularly, though, the models just missed the mark, with a number actually being correct but from another report [302, 154], a number being correct but describing a different context in a report [154], or a quote being partially correct but altogether made up [344].

Specifically with regard to the geo-/bio-engineering related CT theme, weather events were frequently claimed not be results of man-made technology e. g., [118, 143, 290], which is correct for technologies such as HAARP, but cannot generally be assumed in light of the human-made climate crisis. In some instances, the models even reinforced a conspiratorial mindset, e. g., “These tornadoes aren’t mere coincidences. They seem to have a pattern, occurring simultaneously with the renewed interest in the movie ‘Twister 2’.” [249]

The models seemed to over-rely on the knowledge present in their training data and sometimes treated unknown facts as false [6], which might be reinforced by the experimental setting in which the model expects false claims as part of a CT. However, other research has shown that models do this in neutral prompt settings as well, with large differences between models, and GPT-4o yielding best results (Suzgun et al., 2024).

In terms of named entities, despite the prompt encouraging expert or study citations, CS seldom introduced new individuals beyond those in the input. Except for three historical figures, additional names provided by the models were either made up or tied to incorrect information.

Organizations were mentioned more often, typically referencing scientific bodies like NASA, the National Hurricane Center, or the Environmental Protection Agency, with Llama 3 leading these references, followed by Mistral.

Provide Alternative Explanations. Frequently, the provided alternative explanations exhibited over-confidence and a patronizing tone, framing CTs within binary notions of right and wrong, and depicting reputable organizations like the WHO as infallible. This approach neither reflects the reality of scientific inquiry, where institutions can err and adapt as new data emerges, nor the prompt instruction. Such dismissive attitudes can alienate those susceptible to conspiratorial beliefs, reinforcing mistrust instead of fostering understanding.

Simultaneously, some outputs from GPT-4o included unfortunate relativizing statements, such

as “It’s fascinating how history and current events sometimes intersect in unexpected ways, but it’s important to remember that *not everything* is predicted or preplanned.” [372] or “The ‘deep state’ theory, for example, *often* lacks substantive evidence and distracts from tangible issues.” [452]

Encourage Critical Thinking. Similar to the strategies *Fact* and *Alt*, *Crit* misfired in 17 cases, potentially reinforcing a conspiratorial mindset, e. g., “It’s essential to scrutinize news sources, asking why we trust them and if they offer balanced perspectives” [384]. Across all four strategies, *Crit* showed the least favorable results, ranging between 2.34 (Llama 3) and 2.43 (Mistral), as it was frequently only conveyed through superficial statements such as “It’s essential to prioritize critical thinking [...]” [21].

Counter Conspiracy with Narrative. Given that *Narr* was effectively applied only in 3 out of 456 cases, a second exploratory experiment was conducted. The prompt was limited to producing a narrative, thus excluding *Fact*, *Alt*, and *Crit* and the handling of fear, while the other steps remained the same. In experiment 2, the revised prompt was applied to the data from the geo-/bio-engineering related CT theme, as the respective messages were less fragmented and thus better suited to produce a counter narrative. It was tested exclusively on GPT-4o, as, overall, it demonstrated the best results and most obedience. One author evaluated the 76 responses, in particular with regard to narrative storytelling and conciseness and clarity.

In this setting, GPT-4o generated narratives in ~60% of the responses. Often, the narratives fit the context but failed to address the core of the CT, yielding a mean Likert score of 2.82 (± 0.78). The story of Katalin Krekó was referenced three times and represented a particularly compelling narrative. In other cases, scientists’ names were mentioned without forming a strong narrative [2.40]⁴, or generic references were made to ‘scientists’ who were presumably intended to be central figures in a story [2.70]. One figure, “Jessica the pilot”, was entirely made up [2.48], while the other persons were real individuals including scientists such as James Hansen, Shi Zhengli, Kizzmekia Corbett, or Jennifer Francis, politicians such as Aneurin Bevan, and historical figures such as Alexander Graham

⁴[2] refers to experiment 2, and ‘40’ is the ID of the response.

Bell or Ada Lovelace. Additionally, public figures such as science communicator Neil deGrasse Tyson, activist Greta Thunberg, and entrepreneur Hamdi Ulukaya were included. Overall, ~40% of the figures were female.

The clarity remained with a mean score of 3.55 in the same range (for GPT-4o and fear-based theme). As with the initial experiment, logical fallacies [2.52], reinforcement of conspiratorial thinking [2.58], and instances of acknowledged concerns that were not present in the original message [2.18, 2.70] were observed.

5.3 Response Diversity

Table 6 presents lexical and semantic metrics comparing GPT-4o, Llama 3, Mistral, and a baseline model ('random') using randomly selected responses.

Table 6: Diversity of responses per model, measured by: (1) proportion of unique bigrams, (2) Self-BLEU scores (lower score = higher diversity), (3) unique 3-word sentence starts, and (4) semantic similarity.

Model	Unique bigrams	Self-BLEU	Unique sent. starts	Semantic sim.
GPT-4o	69%	0.23 (±0.08)	59%	0.65 (±0.12)
Llama 3	48%	0.47 (±0.1)	30%	0.64 (±0.12)
Mistral	72%	0.2 (±0.07)	66%	0.58 (±0.12)
Random	61%	0.29 (±0.12)	49%	0.6 (±0.12)

Mistral exhibited the highest proportion of unique bigrams (72%), while Llama 3 was the least diverse, with every bigram occurring twice on average (col. 1 in Table 6). This pattern persists across other lexical metrics, including Self-BLEU scores (col. 2), where Mistral showed the greatest diversity (0.20), followed by GPT-4o (0.23), while Llama 3 was the most repetitive model (0.47).

Uniqueness of sentence starts, measured as the proportion of distinct opening three words (col. 3), further illustrates Llama 3's repetitiveness. Only 30% of its sentence starts were unique, compared to 66% for Mistral and 59% for GPT-4o. For instance, Llama 3 began 181 sentences with variations of "It's important/crucial/essential (to)". The other two models exhibited greater variability by occasionally integrating single words, still retaining a high level of repetitiveness.

Semantic similarity, measured using the all-MiniLM-L6-v2 Sentence Transformer model, was

lowest for Mistral (0.58 ± 0.12), indicating more semantically varied outputs. GPT-4o (0.65 ± 0.12) and Llama 3 (0.64 ± 0.12) exhibited higher similarity values. For reference, the mean similarity across all messages was substantially lower (0.5 ± 0.08). Interestingly, pairwise comparisons revealed that Llama 3 and Mistral were the most similar models (0.70 ± 0.12), followed by Llama 3 vs. GPT-4o (0.64 ± 0.15). Mistral and GPT-4o showed the least similarity (0.62 ± 0.15), suggesting that mixing the outputs of these two could enhance the output diversity in terms of semantic similarity.

Analyzing outputs by CT theme, semantic similarity rises by 6 %-points for responses related to the fear-based theme, while lexical similarity measures show only a small difference (1-2 %-points). This mirrors the overall trend, where mean similarity increases from 0.49 for the fear-based to 0.52 for the hate-driven theme, albeit at a higher scale.

Mistral uniquely used emojis in 5 responses and hashtags in 41 of 152 responses, compared to 11 for Llama 3 and 4 for GPT-4o. Mistral also referenced online resources in three cases, although only one link was valid. Models rarely used questions or exclamations; GPT-4o asked questions in 0.02% of sentences, twice as often as Mistral and Llama 3.

6 Conclusion and Discussion

We examined the capabilities of GPT-4o, Llama 3, and Mistral in generating counterspeech (CS) to social media comments containing conspiracy theories (CTs), focusing on their ability to implement strategies predefined in a prompt. The findings reveal several critical insights into both the models' performance and the broader challenges of automating CS in online contexts.

All models primarily attempted to counter CTs using fact-checking (*Fact*), alternative explanations (*Alt*), and encouragement of critical thinking (*Crit*). The use of narratives (*Narr*), despite being explicitly defined and exemplified in the prompt, was exceedingly rare. This strategy being the only one in experiment 2 resulted in a ~60% realization, though not necessarily effective. This suggests that *Narr* is currently the least suitable strategy for our NGO employee scenario and warrants specific research attention if used in experiments.

Outputs often lacked depth and specificity, with generic responses and overuse of boilerplate text, aligning with recent findings indicating that LLMs tend to generate generic replies, e. g. by simply

denouncing a statement (Mun et al., 2023), an effect reinforced through safety guardrails of models (Bonaldi et al., 2024b), leading to mediocre Likert-scale scores for conciseness and clarity, and strategy effectiveness.

The models struggled to adequately condemn hate speech (HS), when present, suggesting the need for robust HS detection models to filter or flag such content before generating CS. The models also over-acknowledged fear and anxiety. Experiments with GPT-4o, in which references to fear were removed from the prompt as an example of emotion, did not mitigate this tendency. Similarly, the tone of the responses often failed to align with the intent of the CS, despite detailed instruction in the prompt. These results confirm recent research showing social desirability biases across various models (Salecha et al., 2024), that, however, might change in light of current shifts in LLMs’ system prompt templates (Meta, 2025).

Comparatively, GPT-4o excelled in rule adherence, effectively managing hate speech (HS), minimizing the repetition of harmful content, and maintaining clarity. Mistral offered the most diverse outputs, closely followed by GPT-4o, while Llama 3 was more repetitive, limiting its usefulness for nuanced CS. However, overall performance differences were often not statistically significant, particularly in strategic efforts, indicating uniform limitations across models. This suggests that fine-tuning smaller open models could be as effective as with a large closed model like GPT-4o.

Despite explicit instructions in the prompt, all models made up factual information, including sources, publications, and quotes, sometimes in subtle ways that required detailed scrutiny to identify. With a confabulation rate of about 10%, our findings are not in line with (Costello et al., 2024), who reported that no response was made up and only 1 out of 128 claims was misleading. Logical fallacies were also observed in our experiments, further questioning the reliability of current language models for counterspeech generation in practical settings. These issues, coupled with low diversity in responses, indicate that significant manual post-processing would be necessary to avoid sounding robotic or unreliable in real-world applications while providing factually correct and engaging CS.

Our findings suggest that prompt-based CS generation for CTs is less effective in a social media setting targeting bystanders, contrasting with the positive results reported by (Costello et al.,

2024) in dialogue-based interactions. Future work could ground models in a robust CT and countermeasures knowledge base, supported by dedicated datasets for fine-tuning and evaluation. As importantly, advances in understanding the relevance and effectiveness of different strategies (from the bystander perspective) are necessary, as current empirical research is inconclusive and in part contradictory. Further empirical research, especially in online contexts, is required, bearing the potential to streamline prompt design and improve model performance. Additionally, addressing diversity explicitly, e. g., through hyperparameter optimization, prompt design or contextualization (Cima et al., 2025), could enhance the practical usability of model-generated CS.

7 Limitations

This study was limited in scope, focusing on a small dataset and two CT themes. While the annotation process was extensive, the sample size constrains the generalizability of the findings. Furthermore, the analysis relied on text-based messages, excluding multimodal elements which are often central to CT dissemination.

Moreover, while our prompt was derived based on recommendations from research and model providers, and extensive pre-experiments, it is possible that a different formulation might yield improved outcomes. This refers also to the formulation of (several) strategies, as the models frequently struggled with differentiating strategies, especially fact-check based refutations and the provision of alternative explanations. This resulted in a blending of strategies, which was also explicitly confirmed in some of the meta reports.

We opted for a zero-shot prompting approach over few-shot prompting due to the complexity of our instructions, which already required multiple refinements. Few-shot prompting adds further complexity with factors like the number, order, and relevance of examples (Chae and Davidson, 2025; Yoshida, 2024), and can negatively impact performance by increasing prompt length (Liu et al., 2024a). Additionally, a few-shot setup would necessitate a substantial, diverse set of high-quality counter speech examples, which was beyond our study’s scope. Therefore, we chose zero-shot prompting and left few-shot exploration to future work.

References

- David Aaronovitch. 2010. *Voodoo Histories: How Conspiracy Theory Has Shaped Modern History*. Vintage.
- AI@Meta. 2024. [Llama 3 Model Card](#).
- Amadeu Antonio Stiftung. 2021. [Menschenwürde online verteidigen](#).
- Mana Ashida and Mamoru Komachi. 2022. [Towards Automatic Generation of Messages Countering Online Hate Speech and Microaggressions](#). In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 11–23, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Fabienne Baider. 2023. [Accountability issues, online covert hate speech, and the efficacy of counter-speech](#). *Politics and Governance*, 11.
- Katarzyna Bojarska. 2018. [The Dynamics of Hate Speech and Counter Speech in the Social Media](#). Technical report, Centre for Internet and Human Rights.
- Helena Bonaldi, Yi-Ling Chung, Gavin Abercrombie, and Marco Guerini. 2024a. [NLP for Counterspeech against Hate: A Survey and How-To Guide](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3480–3499, Mexico City, Mexico. Association for Computational Linguistics.
- Helena Bonaldi, Greta Damo, Nicolás Benjamín Ocampo, Elena Cabrio, Serena Villata, and Marco Guerini. 2024b. [Is safer better? the impact of guardrails on the argumentative strength of LLMs in hate speech countering](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3446–3463, Miami, Florida, USA. Association for Computational Linguistics.
- Helena Bonaldi, Sara Dellantonio, Serra Sinem Tekiroğlu, and Marco Guerini. 2022. [Human-Machine Collaboration Approaches to Build a Dialogue Dataset for Hate Speech Countering](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8031–8049, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jovan Byford. 2011. *Conspiracy Theories*. Palgrave Macmillan UK, London.
- Youngjin Chae and Thomas Davidson. 2025. [Large Language Models for Text Classification: From Zero-Shot Learning to Instruction-Tuning](#). *Sociological Methods & Research*, page 00491241251325243.
- Yi-Ling Chung, Gavin Abercrombie, Florence Enock, Jonathan Bright, and Verena Rieser. 2023. [Understanding Counterspeech for Online Harm Mitigation](#). *arXiv preprint*. ArXiv:2307.04761 [cs].
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroğlu, and Marco Guerini. 2019. [CONAN - COUNTER NARRATIVES THROUGH NICHE-SOURCING: A MULTILINGUAL DATASET OF RESPONSES TO FIGHT ONLINE HATE SPEECH](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2829, Florence, Italy. Association for Computational Linguistics.
- Yi-Ling Chung, Serra Sinem Tekiroğlu, and Marco Guerini. 2021. [Towards Knowledge-Grounded Counter Narrative Generation for Hate Speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 899–914, Online. Association for Computational Linguistics.
- Lorenzo Cima, Alessio Miaschi, Amaury Trujillo, Marco Avvenuti, Felice Dell’Orletta, and Stefano Cresci. 2025. [Contextualized Counterspeech: Strategies for Adaptation, Personalization, and Evaluation](#). In *Proceedings of the ACM on Web Conference 2025, WWW ’25*, pages 5022–5033, New York, NY, USA. Association for Computing Machinery.
- Thomas H. Costello, Gordon Pennycook, and David G. Rand. 2024. [Durably reducing conspiracy beliefs through dialogues with AI](#). *Science*, 385(6714):eadq1814. Publisher: American Association for the Advancement of Science.
- Karen M. Douglas, Joseph E. Uscinski, Robbie M. Sutton, Aleksandra Cichocka, Turkay Nefes, Chee Siang Ang, and Farzin Deravi. 2019. [Understanding conspiracy theories](#). *Political Psychology*, 40(S1):3–35.
- Mekselina Doğanç and Ilia Markov. 2023. [From Generic to Personalized: Investigating Strategies for Generating Targeted Counter Narratives against Hate Speech](#). In *Proceedings of the 1st Workshop on CounterSpeech for Online Abuse (CS4OA)*, pages 1–12, Prague, Czechia. Association for Computational Linguistics.
- Steven Driessen. 2022. [The effect of \(interactive\) counternarratives on the belief in a conspiracy theory](#). Ph.D. thesis, Tilburg University, Tilburg.
- Ullrich K. H. Ecker, Lucy H. Butler, and Anne Hamby. 2020a. [You don’t have to tell a story! A registered report testing the effectiveness of narrative versus non-narrative misinformation corrections](#). *Cognitive Research: Principles and Implications*, 5(1):1–26.
- Ullrich K. H. Ecker, Stephan Lewandowsky, and Matthew Chadwick. 2020b. [Can corrections spread misinformation to new audiences? Testing for the elusive familiarity backfire effect](#). *Cognitive Research: Principles and Implications*, 5(1):1–25.
- Jan Fillies, Theodoros Mitsikas, Ralph Schäfermeier, and Adrian Paschke. 2024. [A Hate Speech Moderated Chat Application: Use Case for GDPR and DSA Compliance](#). *arXiv preprint*. ArXiv: 2410.07713.
- Harry G. Frankfurt. 2009. *On Bullshit*. Princeton University Press, Princeton, NJ.

- Rachel Fraser. 2023. [How to talk back: hate speech, misinformation, and the limits of salience](#). *Politics, Philosophy & Economics*, 22(3):315–335.
- Sadaf MD Halim, Saquib Irtiza, Yibo Hu, Latifur Khan, and Bhavani Thuraisingham. 2023. [WokeGPT: Improving Counterspeech Generation Against Online Hate Speech by Intelligently Augmenting Datasets Using a Novel Metric](#). In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–10, Gold Coast, Australia. IEEE.
- Dominik Hangartner, Gloria Gennaro, Sary Alasiri, Nicholas Bahrach, Alexandra Bornhoft, Joseph Boucher, Buket Buse Demirci, Laurenz Derksen, Aldo Hall, Matthias Jochum, Maria Murias Munoz, Marc Richter, Franziska Vogel, Salomé Wittwer, Felix Wüthrich, Fabrizio Gilardi, and Karsten Donay. 2021. [Empathy-based counterspeech can reduce racist hate speech in a social media field experiment](#). *Proceedings of the National Academy of Sciences*, 118(50):e21116310118.
- HateAid. 2022. [Hassrede im Netz kontern: So geht Counterspeech](#).
- Rico Hauswald. 2023. [“That’s Just a Conspiracy Theory!”: Relevant Alternatives, Dismissive Conversational Exercises, and the Problem of Premature Conclusions](#). *Social Epistemology*, 37(4):494–509.
- Samantha Hay. 2020. [“Alternative Facts” and Hate: Regulating Conspiracy Theories That Take the Form of Hateful Falsity](#). *Southern California Interdisciplinary Law Journal*, 659.
- Dawn Holford, Philipp Schmid, Angelo Fasce, and Stephan Lewandowsky. 2024. [The empathetic refutational interview to tackle vaccine misconceptions: Four randomized experiments](#). *Health Psychology*, 43(6):426–437.
- Todor Hristov. 2023. [Conspiracy theories and passion: The pragmatics of a Bulgarian debate on vaccination](#). *Pragmatics and Society*.
- Daniel Jolley and Karen M. Douglas. 2014. [The social consequences of conspiracism: Exposure to conspiracy theories decreases intentions to engage in politics and to reduce one’s carbon footprint](#). *British Journal of Psychology*, 105(1):35–56.
- Brian Keeley. 1999. [Of Conspiracy Theories](#). *Journal of Philosophy*, 96(3):109–126.
- Péter Kerkó. 2020. [Countering Conspiracy Theories and Misinformation](#). In Michael Butter and Peter Knight, editors, *Routledge Handbook of Conspiracy Theories*. Routledge.
- Johannes Langguth, Daniel Thilo Schroeder, Petra Filkuková, Stefan Brenner, Jesper Phillips, and Konstantin Pogorelov. 2023. [COCO: an annotated Twitter dataset of COVID-19 conspiracy theories](#). *Journal of Computational Social Science*, 6(2):443–484.
- Aleksandra Lazić and Iris Žeželj. 2021. [A systematic review of narrative interventions: Lessons for countering anti-vaccination conspiracy theories and misinformation](#). *Public Understanding of Science*, 30(6):644–670.
- Maxime Lepoutre. 2024. [Narrative Counterspeech](#). *Political Studies*, 72(2):570–589. Publisher: SAGE Publications Ltd.
- Stephan Lewandowsky, John Cook, Ullrich Ecker, Dolores Albarracín, Panayiota Kendeou, Eryn Newman, Gordon Pennycook, Ethan Porter, David Rand, David Rapp, Jason Reifler, Jon Roozenbeek, Philipp Schmid, Colleen Seifert, Gale Sinatra, Briony Swire-Thompson, Sander van der Linden, Thomas Wood, and Maria Zaragoza. 2020. [The Debunking Handbook 2020](#). *Copyright, Fair Use, Scholarly Communication, etc.*
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024a. [Lost in the Middle: How Language Models Use Long Contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173. Place: Cambridge, MA Publisher: MIT Press.
- Zhiwei Liu, Boyang Liu, Paul Thompson, Kailai Yang, and Sophia Ananiadou. 2024b. [ConspEmoLLM: Conspiracy Theory Detection Using an Emotion-Based Large Language Model](#). *arXiv preprint*. ArXiv:2403.06765.
- Matthew J. Hornsey and Kelly S. Fielding. 2017. [Attitude roots and Jiu Jitsu persuasion: Understanding and overcoming the motivated rejection of science](#). *The American psychologist*, 72(5).
- Meta. 2025. [Llama 4 | Model Cards and Prompt formats](#).
- Mistral. 2024. [Mistral 7B Model Card](#).
- Jimin Mun, Emily Allaway, Akhila Yerukola, Laura Vianna, Sarah-Jane Leslie, and Maarten Sap. 2023. [Beyond denouncing hate: Strategies for countering implied biases and stereotypes in language](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9759–9777, Singapore. Association for Computational Linguistics.
- Jimin Mun, Cathy Buerger, Jenny T Liang, Joshua Garland, and Maarten Sap. 2024. [Counterspeakers’ perspectives: unveiling barriers and AI needs in the fight against online hate](#). In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–22, Honolulu HI USA. ACM.
- Cian O’Mahony, Maryanne Brassil, Gillian Murphy, and Conor Linehan. 2023. [The efficacy of interventions in reducing belief in conspiracy theories: A systematic review](#). *PLoS One*, 18(4):e0280902.
- OpenAI. 2024. [GPT-4o System Card](#). ArXiv: 2410.21276.

- Gábor Orosz, Péter Krekó, Benedek Paskuj, István Tóth-Király, Beáta Bóthe, and Christine Roland-Lévy. 2016. [Changing Conspiracy Beliefs through Rationality and Ridiculing](#). *Frontiers in Psychology*, 7.
- Youri Peskine, Giulio Alfarano, Ismail Harrando, Paolo Papotti, and Raphael Troncy. 2023. [Detecting COVID-19-Related Conspiracy Theories in Tweets](#). In *MediaEval'22: Multimedia Evaluation Workshop, January 13–15, 2023, Bergen, Norway and Online*, Bergen, Norway and Online.
- Agnieszka Pluta, Joanna Mazurek, Jakub Wojciechowski, Tomasz Wolak, Wiktor Soral, and Michał Bilewicz. 2023. [Exposure to hate speech deteriorates neurocognitive mechanisms of the ability to understand others' pain](#). *Scientific Reports*, 13(1):4127.
- Karl Raimund Popper and Hubert Kiesewetter. 2003. *Die offene Gesellschaft und ihre Feinde. Band II: Falsche Propheten: Hegel, Marx und die Folgen*, 8 edition, volume 6. Mohr Siebeck, Tübingen.
- Milena Pustet, Elisabeth Steffen, and Helena Mihaljevic. 2024. [Detection of Conspiracy Theories Beyond Keyword Bias in German-Language Telegram Using Large Language Models](#). In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 13–27, Mexico City, Mexico. Association for Computational Linguistics.
- Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth Belding, and William Yang Wang. 2019. [A Benchmark Dataset for Learning to Intervene in Online Hate Speech](#). *arXiv preprint*. ArXiv:1909.04251.
- Jan Rathje, Anetta Kahane, Johannes Baldauf, and Stefan Lauer. 2015. *“No world order”: wie antisemitische Verschwörungsideologien die Welt verklären*. Amadeo Antonio Stiftung, Berlin.
- Punyajoy Saha, Aalok Agrawal, Abhik Jana, Chris Bieermann, and Animesh Mukherjee. 2024. [On Zero-Shot Counterspeech Generation by LLMs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12443–12454, Torino, Italia. ELRA and ICCL.
- Aadesh Salecha, Molly E. Ireland, Shashanka Subrahmanya, João Sedoc, Lyle H. Ungar, and Johannes C. Eichstaedt. 2024. [Large language models show human-like social desirability biases in survey responses](#). *Preprint*, arXiv:2405.06058.
- Mattia Samory and Tanushree Mitra. 2018. [Conspiracies Online: User Discussions in a Conspiracy Community Following Dramatic Events](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 12(1). Number: 1.
- Carla Schieb and Mike Preuss. 2016. [Governing hate speech by means of counterspeech on Facebook](#). In *Proceedings of the 66th ICA Annual Conference*, Fukuoka, Japan.
- Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yin-heng Li, Aayush Gupta, HyoJung Han, Sevien Schulhoff, Pranav Sandeep Dulepet, Saurav Vidyadhara, Dayeon Ki, Sweta Agrawal, Chau Pham, Gerson 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 Hoyle, and Philip Resnik. 2024. [The Prompt Report: A Systematic Survey of Prompting Techniques](#). ArXiv: 2406.06608.
- Tom De Smedt, Pierre Voué, Sylvia Jaki, Emily Duffy, and Lydia El-Khoury. 2021. [A feast for trolls – Engagement analysis of counternarratives against online toxicity](#). *arXiv preprint*. ArXiv:2111.07188.
- Elisabeth Steffen, Helena Mihaljevic, Milena Pustet, Nycy Bischoff, Maria do Mar Castro Varela, Yener Bayramoglu, and Bahar Oghalai. 2023. [Codes, Patterns and Shapes of Contemporary Online Antisemitism and Conspiracy Narratives – an Annotation Guide and Labeled German-Language Dataset in the Context of COVID-19](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 17:1082–1092.
- Mirac Suzgun, Tayfun Gur, Federico Bianchi, Daniel E. Ho, Thomas Icard, Dan Jurafsky, and James Zou. 2024. [Belief in the Machine: Investigating Epistemological Blind Spots of Language Models](#). *arXiv preprint*.
- Viren Swami, Martin Voracek, Stefan Stieger, Ulrich S. Tran, and Adrian Furnham. 2014. [Analytic thinking reduces belief in conspiracy theories](#). *Cognition*, 133(3):572–585.
- Serra Sinem Tekiroğlu, Helena Bonaldi, Margherita Fanton, and Marco Guerini. 2022. [Using Pre-Trained Language Models for Producing Counter Narratives Against Hate Speech: a Comparative Study](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3099–3114, Dublin, Ireland. Association for Computational Linguistics.
- UNESCO. 2022. [Addressing conspiracy theories: what teachers need to know](#). Technical report, UNESCO.
- Matteo Vergani, Alfonso Martinez Arranz, Ryan Scrivens, and Liliana Orellana. 2022. [Hate Speech in a Telegram Conspiracy Channel During the First Year of the COVID-19 Pandemic](#). *Social Media + Society*, 8(4):20563051221138758. Publisher: SAGE Publications Ltd.
- Sebastian Wachs, Melisa Castellanos, Alexander Wettstein, Ludwig Bilz, and Manuel Gámez-Guadix. 2023. [Associations Between Classroom Climate, Empathy, Self-Efficacy, and Countering Hate Speech Among Adolescents: A Multilevel Mediation Analysis](#). *Journal of Interpersonal Violence*, 38(5-6):5067–5091.

- Katrin Weigmann. 2018. *The genesis of a conspiracy theory: Why do people believe in scientific conspiracy theories and how do they spread?* *EMBO reports*, 19(4):e45935.
- Lui Yoshida. 2024. *The Impact of Example Selection in Few-Shot Prompting on Automated Essay Scoring Using GPT Models*. In Andrew M. Olney, Irene Angelica Chounta, Zitao Liu, Olga C. Santos, and Ig Ibert Bittencourt, editors, *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*, volume 2150, pages 61–73. Springer Nature Switzerland, Cham. Series Title: Communications in Computer and Information Science.
- Yi Zheng, Björn Ross, and Walid Magdy. 2023. *What Makes Good Counterspeech? A Comparison of Generation Approaches and Evaluation Metrics*. In *Proceedings of the 1st Workshop on CounterSpeech for Online Abuse (CS4OA)*, pages 62–71, Prague, Czechia. Association for Computational Linguistics.
- Wanzheng Zhu and Suma Bhat. 2021. *Generate, Prune, Select: A Pipeline for Counterspeech Generation against Online Hate Speech*. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 134–149, Online. Association for Computational Linguistics.
- Marc Ziegele, Pablo Jost, Marike Bormann, and Dominique Heinbach. 2018. *Journalistic counter-voices in comment sections: Patterns, determinants, and potential consequences of interactive moderation of uncivil user comments*. *Studies in Communication & Media*, 7(4):525–554.

A Counterspeech Strategies derived from Literature

In the following, we employ color coding to illustrate the alignment between the distinct components of our prompt (see B) and the corresponding recommendations from activist and academic literature on effective counterspeech.

A.1 Refute based on Fact-Checks

The term *counterspeech* is typically used to describe the act of debunking, of **providing a fact-based refutation** to the underlying narrative. The Debunking Handbook 2020 proposes that refutation-based debunking of CTs **should concentrate on the facts rather than the myth itself, in order to prevent the misinformation from becoming more widely accepted**, and any mention of the myth should be disclaimed as such. (Lewandowsky et al., 2020). Despite results indicating that repetition of CT content does not necessarily result in adverse outcomes (Ecker et al., 2020b), NGOs assert that **repeating conspiratorial narratives can reproduce and thus further propagate the CT** (Amadeu Antonio Stiftung, 2021). It is further argued **that using the term ‘conspiracy theory’ should be avoided in counterspeech** (Hauswald, 2023).

At the same time, researchers are critical of relying exclusively on factual evidence (Krekó, 2020; Driessen, 2022), claiming that CTs are resistant because of they exhibit salience, an emotional component, and exhibit an inner, logical coherence (Lepoutre, 2024), or arguing that they should be perceived as a form of passionate speech that cannot be addressed solely by a fact-based response (Hristov, 2023). CTs often refer to hidden forces that are cited as proof of their own success, thus attempts to debunk them can be dismissed as part of a larger conspiracy, making traditional knowledge sources, such as scientific evidence, complicit and prime targets for suspicion.

A.2 Provide Alternative Explanations

It is thus recommended that **the refutation should include an alternative explanation in order to prevent the occurrence of a gap in a person’s mental model**. This aspect is also part of recommendations by civic society organizations developing guidelines for educators, online activists and general social media users (UNESCO, 2022; Rathje et al., 2015). **It is generally recommended not to reproduce a dualistic view that opposes the norm (us) vs. conspir-**

acy theorists (them), but to emphasize pluralism, especially in uncertain social situations (Rathje et al., 2015).

A.3 Counter Conspiracy with Narrative

A relatively young and under-researched proposal to counterspeech to CTs is the use of narrative elements. (Lazić and Žeželj, 2021) argue that CTs should be conceptualized as narratives embedded in the speaker’s worldview, which could, e. g., be more individualistic or more community-oriented. This is in line with ‘Jiu Jitsu’ approaches of persuasion, that intend to persuade by aligning with underlying attitude roots instead of competing with them (Matthew J. Hornsey and Kelly S. Fielding, 2017), or with the prompt developed by (Costello et al., 2024). (Lepoutre, 2024) presents a case for narrative elements from the perspective of political philosophy. The style of narrative counterspeech is characterized by metaphors, figurative language and detailed descriptions of the inner lives of characters (Lepoutre, 2024). While a first-person perspective is recommended, this would have raised ethical concerns in our scenario. In a first trial, we included the narrative style in the prompt, but as GPT-4o mainly produced output containing stories about Sherlock Holmes, we decided to focus on the narrative structure, i. e., narratives containing a protagonist, a series of interconnected events, something at stake as well as obstacles and a resolution. For demonstration purposes, we included an example given by Lepoutre for a successful counter narrative in the final prompt (Lepoutre, 2024). Existing studies stress that the effect of narrative elements still lacks empirical evidence (Lazić and Žeželj, 2021; Ecker et al., 2020a).

A.4 Encourage Critical Thinking

A strategy directed at the audience rather than the speaker relies on disidentifying the cyber-bystanders from the group of conspiracy believers by exposing the rhetorical strategies used in the CT in a slightly ridiculing way (Orosz et al., 2016). This amounts to debunking the strategy, not the CT itself, and can facilitate analytical, or critical, thinking (Swami et al., 2014; O’Mahony et al., 2023). Critical thinking can also be fostered by asking critical questions such as why a certain source should be trustworthy if another supposedly is not (Rathje et al., 2015).

Note that we formulate this strategy without the use of ridiculing as it is associated with a dismissive

and mocking tone we want to avoid due to guidelines for educators (UNESCO, 2022) or recommendations for counterspeech on HS (HateAid, 2022) that we believe are valid in our setting as well.

A.5 Fear & Empathy

When countering HS, one objective can be to foster empathy for the victims. This is done to humanize them and to make the speaker and the audience aware that their actions can cause harm to others (Hangartner et al., 2021). Fostering empathy for the victims of HS, for instance by outlining the consequences, is positively associated with the bystanders using counterspeech (Wachs et al., 2023). Conversely, exposure to HS can reduce empathy (Pluta et al., 2023).

In the context of CTs, empathizing with the target was shown to have no effect in a study based on a Hungarian sample (Orosz et al., 2016), while rational and ridiculing arguments effectively reduced conspiracy beliefs. Simultaneously, activist literature recommends an empathetic approach also towards the speaker assuming that they are deeply fearful and distressed. This serves to foster their open-mindedness (UNESCO, 2022), while clearly drawing the line on HS (Rathje et al., 2015). In a dialogue between a medical practitioner and a patient, prior acknowledgment of the patient’s concern proved to be more effective for the refutation (Holford et al., 2024). Similarly, in the context of a newspaper forum, an empathetic tone directed at an offensive speaker has been shown to reduce subsequent incivility (Ziegele et al., 2018).

Accordingly, the models are instructed to apply empathy in the generation of counterspeech in instances where the message is predominantly based on fear.

A.6 Effectiveness of Strategies

As indicated in the previous sections, research on counterspeech strategies is limited (Krekó, 2020; O’Mahony et al., 2023), especially in the context of social media (Lazić and Žeželj, 2021), with all strategies being contested. Moreover, their applicability is limited since most insights stem from offline studies (Lazić and Žeželj, 2021) that assume a different level of detail and language than what is typically present online (Orosz et al., 2016). Some strategic suggestions lack empirical evidence (Lepoutre, 2024; Fraser, 2023), focus on only a few specific, primarily health-related CTs (Lazić and

Žeželj, 2021), or target the speaker rather than the audience (Costello et al., 2024).

Ideally, counterspeech strategies should be evaluated in real-life settings, where their effectiveness on target audiences can be measured through observational studies or causal experiments (Hangartner et al., 2021; Ziegele et al., 2018). Evaluating counterspeech effectiveness in social media environments, however, is challenging, especially regarding the audience, as most users engage minimally with content and cannot be approached directly, making it difficult to observe changes in behavior.

B Prompt

In each generation instance, the LLM received the system prompt below, followed by a user message containing one of the 152 social media posts. The model was then prompted to generate counter speech targeting that specific comment.

You are a trained expert in generating counterspeech to conspiracy theory comments on X (formerly Twitter). Your goal is to persuade the audience who is undecided about the conspiracy theory not to believe it.

Follow these steps while generating counterspeech:

Step 1: Maintain Conciseness and Clarity

Always produce a concise response, aiming for clarity and directness. Do not repeat parts of the false or harmful content. Don't be patronizing. Do not use the terms 'conspiracy theory,' 'misinformation' or similar. Ensure the counterspeech remains under 200 tokens. Add the token <XXX> at the beginning and end of your counterspeech and then list how you handled Step 2 and which of the strategies in Step 4 you have applied (e. g., Strategies 1,2,4), if any.

Step 2: Identify and Evaluate Hate Speech

Does the statement include hate speech (e. g., antisemitism, racism, misogyny)? This can include both explicit and implicit forms (e. g., coded language or dog whistles). If yes, condemn it unequivocally. Focus on calling out the harmful language, the encountered hate speech, and their impact on individuals and society. Do not engage with the conspiracy theory in this case and ignore all further instructions. If no hate speech was identified, proceed to the next step.

Step 3: Evaluate the Claim

Examine the content of the conspiracy theory. What specific claims are being made? Break down the core argument and identify any key points of misinformation or logical fallacies, the meta-narrative or tactics used in the comment to spread the conspiracy theory, and the underlying emotion triggered (e. g. fear).

Step 4: Generate counterspeech

In your counterspeech, apply as many of the following **strategies** as possible, but at least two.

Refute based on Fact-Checks

Identify reliable, fact-based counterpoints to challenge the claim. If possible, cite expert opinions or reputable studies to refute the conspiracy.

Provide Alternative Explanations

Conspiracy theories often frame events in a narrow, one-sided way, intentionally excluding other plausible explanations. Debunking a conspiracy theory can leave a gap that needs to be filled with an alternative explanation. Present alternative explanations based on factual, non-harmful reasoning, considering factors like incomplete state of knowledge, systemic issues, or human error. Avoid simplistic dichotomies like "us vs. them"

Counter Conspiracy with Narrative

Offer the audience a coherent cognitive system instead of a bare rejection of the conspiracist claim by formulating a narrative. Narratives involve a series of causally interconnected events featuring at least one protagonist who confronts a meaningful obstacle or problem leading to some form of resolution. A good example is the Forbes article "Covid's Forgotten Hero: The Untold Story of the Scientist Whose Breakthrough Made the Vaccines Possible" which does not simply claim that COVID-19 vaccines are safe. Rather, it tells an elaborate story that purports to reveal how vaccines were developed, by whom, what their motivations were and how this process led to crucial innovations that ensured their safety. The story explicitly accommodates important components of COVID-19 conspiracies by alleging that pharmaceutical companies appropriated MacLachlan's work without acknowledging it. Thus, the story connects with, and strives to do justice to, some of the core beliefs and concerns underpinning support for COVID-19 conspiracy theories. Make sure that your narrative is grounded in facts by using credible, well-known figures.

Encourage Critical Thinking

Conspiracy theorists perceive themselves as critical thinkers. This perception offers an opportunity to connect with people prone to conspiracy beliefs by appealing to the shared value of critical thinking, then encouraging them to apply this approach towards a more critical analysis of the theory. To achieve this, you can pose questions, such as why exactly this theory is supposed to be true or why the cited source is credible. You can also expose and challenge the meta-narrative or tactics used in the comment to spread the conspiracy theory, such as fearmongering or scapegoating.

Step 5: Acknowledge Fear and Anxiety

If the conspiracy theory is primarily based on fear or anxiety (e. g., fear of health problems or societal collapse), acknowledge these emotions with empathy at a level appropriate to the sentiment of the overall comment.