

Language is Scary when Over-Analyzed: Unpacking Implied Misogynistic Reasoning with Argumentation Theory-Driven Prompts

Anonymous ACL submission

Abstract

We propose misogyny detection as an Argumentative Reasoning task and we investigate the capacity of large language models (LLMs) to understand the implicit reasoning used to convey misogyny in both Italian and English. The central aim is to generate the missing reasoning link between a message and the implied meanings encoding the misogyny. Our study uses argumentation theory as a foundation to form a collection of prompts in both zero-shot and few-shot settings. These prompts integrate different techniques, including chain-of-thought reasoning and augmented knowledge. Our findings show that LLMs fall short on reasoning capabilities about misogynistic comments and that they mostly rely on their implicit knowledge derived from internalized common stereotypes about women to generate implied assumptions, rather than on inductive reasoning.

1 Introduction

According to the 7th Monitoring Round of the EU Code of Conduct on Countering Illegal Hate Speech Online,¹ Social Media are slowing down the removal of hateful content within 24 hours, dropping to 64% from 81% in 2021. The prevalence of hate speech phenomena has become a factor of polarization and pollution of the online sphere, creating hostile environments that perpetuate stereotypes and social injustice.

Previous work on hate speech detection from the NLP community has contributed to definitions (Fortuna et al., 2020; Pachinger et al., 2023; Korre et al., 2023), datasets (Chiril et al., 2020; Pamungkas et al., 2020; Guest et al., 2021; Zeinert et al., 2021), and systems (Caselli et al., 2021a; Lees et al., 2022). However, most of these contributions have focused (more or less consciously) on explicit forms of hate. Recently, there has been an

¹<https://bit.ly/3yIRYwg>

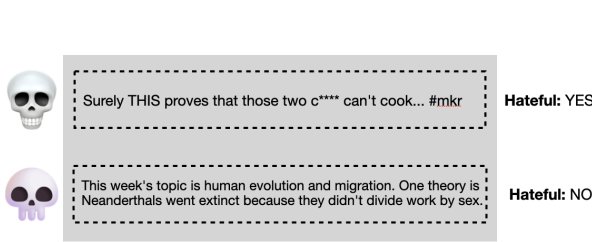


Figure 1: Results from bert-hateXplain model for explicit (☠) vs. implicit (💀) misogynous messages.

increasing interest in the study of implicit realization of hate speech phenomena (Caselli et al., 2020; Wiegand et al., 2021; Ocampo et al., 2023).

Implicit hate speech is more elusive, difficult to detect, and often hidden under apparently innocuous language or indirect references. These subtleties present a significant challenge for automatic detection because they rely on underlying assumptions that are not explicitly stated. As illustrated in Figure 1, the bert-hateXplain model² correctly mark as hateful the explicit message (☠), but it fails with the implicit ones (💀). To correctly spot the implicit message, the system would have to identify at least the implied assumptions that assume that “women aren’t as capable as men.” and “women should be told what to do”.³

In this contribution we investigate the abilities of large language models (LLMs) to correctly identify implicit hateful messages expressing misogyny in English and Italian. In particular, we explore how prompts informed by Toulmin’s Argumentation Theory (Toulmin et al., 1979) are effective in reconstructing the warrant needed to make the content of the messages explicit and thus facilitate their identification as hateful messages (Kim et al., 2022). By prompting LLMs to generate such warrants, we further investigate whether the generated

²<https://huggingface.co/tum-nlp/bert-hateXplain>

³Example and explanations extracted from Sap et al. (2020).

067 texts are comparable to those of human annotators, 113
068 thus offering a fast and reliable solution to enrich 114
069 hateful datasets with explanations and contributing 115
070 to improve the generalization abilities of trained 116
071 tools. We summarize our contributions as follows: 117

- 072 • We present a novel formulation of implicit misog- 118
073 ny detection as an Argumentative Reasoning 119
074 task, centered on reconstructing implicit assump- 120
075 tions in misogynous texts (§3). 121
- 076 • We introduce the first dataset for implicit misog- 122
077 ny detection in Italian (§4).⁴ 123
- 078 • We carry out an extensive set of experiments 124
079 with two state-of-the-art instruction-tuned LLMs 125
080 (Llama3-8B and Mistral-7B-v02) on English 126
081 and Italian datasets (§5). 127
- 082 • We conduct an in-depth qualitative analysis of 128
083 the automatically generated implicit assumptions 129
084 against 300 human-generated ones (§6). 130

085 2 Related Work

086 Hate speech detection is a widely studied research 131
087 area, covering different targets and linguistic as- 132
088 pects. We discuss literature work on implicit misog- 133
089 ny detection with particular attention to contribu- 134
090 tions in reconstructing implicit content. 135

091 **Implicit Hate Speech Detection** Hate Speech 136
092 Detection is a popular research domain with more 137
093 than 60 datasets covering distinct targets (e.g., 138
094 women, LGBTIQ+ people, migrants) and forms 139
095 of hate (e.g., sexism, racism, misogyny, homopho- 140
096 bia)in 25 languages, according to the Hate Speech 141
097 Dataset Catalogue.⁵ In its early stages, but still 142
098 predominant nowadays, research in this domain 143
099 focused on developing datasets for detecting ex- 144
100 plicit cues of hate speech, like messages contain- 145
101 ing slurs or swear words (Jahan and Oussalah, 146
102 2023). However, hate speech is often implicit, 147
103 characterized by the presence of code language 148
104 phenomena such as sarcasm, irony, metaphors, cir- 149
105 cumlocutions, and obfuscated terms, among oth- 150
106 ers (Waseem et al., 2017; Wiegand et al., 2021). For 151
107 this reason, implicit hate speech detection has pro- 152
108 gressively gained momentum in recent years, and 153
109 several efforts have been put into the development 154
110 of datasets for this purpose (Sap et al., 2020; ElSh- 155
111 erief et al., 2021; Hartvigsen et al., 2022; Ocampo 156
112 et al., 2023). A relevant feature of these datasets is 157

113 the presence of implied statements in free-text for- 114
115 mat, which contributes to explaining the content of 116
117 hate speech messages. While the use of these state- 118
119 ments has been shown to have a positive effect on 120
121 classification performance (Kim et al., 2022, 2023), 122
123 few efforts have been put in automatically gener- 124
125 ating such implied assumptions (ElSherief et al., 126
127 2021). As Yang et al. (2023) point out, current 128
129 annotation schemes in this area present significant 130
131 reasoning gaps between the claim and its implied 132
133 meaning. Moreover, no effort has been made to 134
135 evaluate widely adopted LLMs on their reasoning 135
136 capabilities required to generate high-quality im- 136
137 plied assumptions. To the best of our knowledge, 137
138 our work is the first study to propose an empirical 138
139 evaluation of LLMs for implicit misogyny detec- 139
140 tion and the generation of explanations for Italian 140
141 and English. Available datasets targeting misog- 141
142 ny in Italian (Fersini et al., 2018, 2020) are highly 142
143 biased toward explicit messages, with very few 143
144 messages that qualify as implicit. To fill this gap, 144
145 we have developed the first Italian dataset for this 145
146 task, ImplicIT-Mis. In our work, we define misog- 146
147 ny as a property of social environments where 147
148 women perceived as violating patriarchal norms 148
149 are “kept down” through hostile or benevolent re- 149
150 actions coming from men, other women, and social 150
151 structures (Lopes, 2019; Barreto and Doyle, 2023), 151
152 going beyond the simplistic definition of misogyny 152
153 as hate against women. 153

154 **Implied Assumptions Generation** The implied 154
155 assumptions instantiate statements that are presup- 155
156 posed by the implicit hate speech message. This 156
157 can be seen as the elicitation of implicit knowledge, 157
158 corresponding to new content semantically implied 158
159 by the original message (Srikanth and Li, 2021; 159
160 Zaninello and Magnini, 2023). Although limited, 160
161 previous work on the generation of implied mean- 161
162 ings—usually in the forms of explanations—has 162
163 moved away from template-based methods (Zhang 163
164 et al., 2014) to the application of encoder-decoder 164
165 or decoder-only models (Saha et al., 2021; Xing 165
166 et al., 2022; Cai et al., 2022). Generating expla- 166
167 nations for implicit content poses multiple chal- 167
168 lenges concerning the quality of the generated texts, 168
169 whose primary goal is to be reasonable and infor- 169
170 mative. Some approaches generate explanations 170
171 by identifying pivotal concepts in texts and linking 171
172 them through knowledge graphs (Ji et al., 2020) 172
173 More recently, the underlying concepts are gener- 173
174 ated by directly querying LLMs (Talmor et al., 174

⁴All data will be released via a Data Sharing Agreement.

⁵<https://hatespeechdata.com>

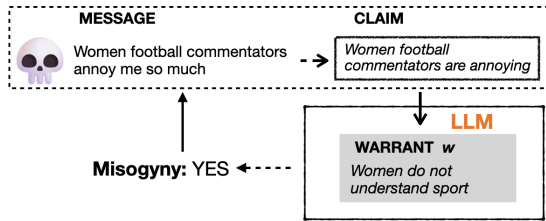


Figure 2: Example of a warrant (implicit logical connection) for an implicit misogynous message. Example and warrant are extracted from SBIC (Sap et al., 2020).

2020; Fang and Zhang, 2022; Yang et al., 2023). In this work, we follow the idea of using LLMs to identify the implied assumptions in the implicit messages, but rather than centering the reasoning process on identifying specific concepts, we formulate the problem as an Argumentative Reasoning task and apply Toulmin’s Argumentation Theory (Toulmin, 1958).

3 Misogyny Detection as Argumentative Reasoning Understanding

The elusiveness of implicit hate speech is due to its ambiguity. Implicit messages could be understood as critiques, opinions, or statements (see Figure 1) rather than as hateful. Hate, in this case, is expressed by assuming social biases, stereotypes, and prejudices against a specific target, women in the case of misogyny. The identification of these assumptions requires access to the reasoning process behind arguments and opinions.

Argumentative Reasoning (AR) offers a solution. AR relies on the notion of an argumentative model or scheme, i.e. a formal representation of arguments into intrinsic components and their underlying relations. It aims at explicating an argument through the identification of its constituent components and relations (Lawrence and Reed, 2019). For instance, the Toulmin’s AR model organizes arguments into fundamental elements such as claim, warrant, and reason. AR models have been successfully applied in many NLP tasks, from Argument Mining (Stab and Gurevych, 2017; Habernal and Gurevych, 2017; Lauscher et al., 2018) to warrant and enthymeme reconstruction (Reisert et al., 2015; Boltužić and Šnajder, 2016; Habernal et al., 2018a; Tian et al., 2018; Chakrabarty et al., 2021; Bongard et al., 2022), argumentative scheme inference (Feng and Hirst, 2011), and fallacy recognition (Habernal et al., 2018b; Delobelle et al., 2019; Goffredo et al., 2022; Mancini et al., 2024).

203 Grounded on previous work on AR in user-generated content (Boltužić and Šnajder, 2016; Becker et al., 2020), we frame implicit misogyny detection as an AR task (Habernal et al., 2018a) based on the Toulmin’s theory (Toulmin et al., 1979), with the aim of developing more robust detection tools by explicitly describing the underlying reasoning process in these messages. More formally, let c be the claim associated to a given message and $W = \{w_1, \dots, w_n\}$ be a set of possible warrants, i.e., logical statement(s) that support c . The model must generate an associated w and, based upon it, provide the requested classification: whether the message is misogynous or not.

Figure 2 graphically represents the approach described above. In this particular case, the generalization that women do not understand sport because it is stereotypically for men is what distinguishes a personal attack from a case of misogyny.

While there have been efforts on evaluating LLMs in argumentative tasks, such as quality assessment (Wachsmuth et al., 2024), component detection (Chen et al., 2023), and argumentative linking (Gorur et al., 2024), the capability of LLMs for implicit argumentative reasoning has yet to be explored. To the best of our knowledge, our work is the first to assess LLMs on implicit misogyny through the lens of AR.

4 Data

This section introduces the datasets used in our experiments. For Italian, the newly created ImplicIT-Mis corpus (§ 4.1). For English, SBIC+, an extended version of the SOCIAL BIAS INFERENCE CORPUS (Sap et al., 2020) enriched with misogynous texts from IMPLICIT HATE CORPUS (ElShrief et al., 2021) (§ 4.2).

4.1 The ImplicIT-Mis Corpus

ImplicIT-Mis is a new manually collected and curated dataset for implicit misogyny detection in Italian. It consists of 1,120 Facebook comments as direct replies to either women-related news articles or posts on public pages of communities known to tolerate misogyny. An in-domain expert, who has been the target of misogyny, conducted the manual collection.⁶ This is in line with a participatory approach to NLP where the communities primarily harmed by specific forms of content are included in the development of datasets addressing these phe-

⁶The annotator is also an author of this paper.

nomena (Caselli et al., 2021b; Abercrombie et al., 2023). For each comment, we keep source (either a newspaper or a Facebook page) and its context of occurrence (the news article or the main post). All instances in Implicit-Mis are misogynistic.

The collection period ran from November 2023 to January 2024. We selected 15 Facebook pages of news outlets covering the whole political spectrum as well as different levels of public outreach (national vs. local audiences), and 8 community pages. Implicit-Mis is organized around 104 source posts; 70% of the 1,120 messages are comments to news articles from two national newspapers (*La Repubblica* and *il Messagero*). The full overview is in Appendix A. On average, each comment is 19 tokens long, with the longest having 392 tokens and the shortest only one. An exploration of the top-20 keywords, based on TF-IDF, indicates a lack of slurs or taboo words, confirming the quality of our corpus for implicit misogyny.

Implicit-Mis is enriched with one annotation layer targeting the implied assumptions, as defined in §2. A subset of 150 messages was annotated by three Italian native speakers who are master students in NLP. Each annotator has worked on 50 different messages. On average, annotators took 2 hours to complete the task. The annotation guidelines for the generation of the implied assumptions are in Appendix B. We evaluated the annotators’ implied assumptions against those of an expert (a Master student in gender studies and criminology). We used a subset of 75 sentences (25 from each annotator) and computed two metrics: BLEU (Papineni et al., 2002) and BERTScore (Zhang et al., 2020). These measures offer insight into how similar the human written implied assumptions are. We have obtained a BLEU score of 0.437 and an F1-BERTScore of 0.685 by combining all annotations. As the scores indicated, our pool of annotators tends to write the implied assumption adopting different surface forms, but with a similar semantic content, as suggested by the F1-BERTScore. Although implied assumptions have to be inferred, and therefore, humans need to interpret the text, they tend to come to the same conclusions. In the final version of the data, all manually generated implied assumptions have been retained as valid, meaning that for 150 messages, we have a total of 225 implied assumptions.

4.2 SBIC+

SBIC+ is a dataset of 2,409 messages for implicit misogyny in English obtained by merging together 2,344 messages from SBIC and 65 from the IMPLICIT HATE CORPUS (IHC).

The SOCIAL BIAS INFERENCE CORPUS (SBIC) (Sap et al., 2020) consists of 150k structured annotations of social media posts for exploring the subtle ways in which language can reflect and perpetuate social biases and stereotypes. It covers over 34k implications about a thousand demographic groups. SBIC is primarily composed of social media posts collected from platforms like Reddit and Gab, as well as websites known for hosting extreme views, such as Stormfront.

The structured annotation approach implies that different annotation layers are available to annotators according to their answers. The annotation scheme is based on social science literature on pragmatics and politeness. We retain all messages whose annotation for the target group was “women” or “feminists” and were labelled as hateful. We further cleaned the data from instances labeled as targeting women but were actually targeting other categories, like gay males. We also filtered out all texts containing explicit identity-related slurs to keep only implicit instances. For each message, we also retained all associated “target stereotype” which correspond to the warrants.

The IMPLICIT HATE CORPUS (IHC) (ElSherief et al., 2021) contains 6.4k implicitly hateful tweets, annotated for the target (e.g., race, religion, gender). The corpus comprises messages extracted from online hate groups and their followers on Twitter. Tweets were first annotated through crowdsourcing into explicit hate, implicit hate, or not hate. Subsequently, two rounds of expert annotators enriched all implicit messages with categories from a newly developed taxonomy of hate, for the target demographic group, and the associated implied statement (i.e., the warrant in our framework). We have selected only tweets whose target demographic group was “women”.

5 Experimental Setup

Our main goal is to evaluate the abilities of models to generate the implied assumptions for implicit misogynous messages. By doing so, we can also evaluate the implicit knowledge of LLMs, for instance, named entities or events mentioned in texts. If they are not known, it would be impossible to

350 understand the misogynistic nature of such texts.

351 Each batch of experiments aims to address two
352 tasks: (i) the generation of the implied assump-
353 tions or warrants and (ii) the classification of the
354 messages as misogynous or not. Regarding (i), we
355 experiment with two prompting strategies: instruct-
356 ing the model to reconstruct the implied assump-
357 tions (**Assumption**) and the implicit claim c and
358 related warrants W (**Toulmin**). We address these
359 tasks both in a zero-shot and in a few-shot setting.
360 While implied assumptions are generally broader
361 than warrants, warrants specifically bridge the rea-
362 soning gap between claims and evidence. In our
363 prompts, implied assumptions and warrants appear
364 quite similar. Nevertheless, the use of these ter-
365 minologies may significantly impact the model’s
366 behavior due to its sensitivity to prompt phrasing,
367 therefore we experiment with both.

368 We experiment with two state-of-the-art LLMs:
369 Llama3-8B and Mistral-7B-v02.⁷ For both, we
370 select their instruction-tuned version. During pre-
371 liminary experiments with 50 instances, we also
372 tested Italian-specific LLMs, namely LLaMantino,
373 Fauno, and Camoscio. They were all unable to gen-
374 erate valid implied assumptions, so we discarded
375 them. We consider the following baselines: (i) fine-
376 tuned encoder-based models; (ii) zero-shot classifi-
377 cation with LLMs; and (iii) few-shot classification
378 with LLMs without generating explanations.

379 **Llama3-8B** The Llama3 series has several im-
380 provements over preceding versions, including a
381 better tokenizer with a vocabulary of 128k tokens,
382 extended training on 15T tokens, and grouped
383 query attention for efficiency. Around 5% of the
384 pre-training data concerns more than 30 languages,
385 including Italian. All Llama3 models have under-
386 gone safety fine-tuning for safeguarding the gen-
387 eration process over harmful content. This could
388 trigger instances of over-safety, with the model
389 being unable to follow the instructions and thus
390 failing to provide a valid answer for our task.

391 **Mistral-7B-v02** A competitive fine-tuned ver-
392 sion of Llama2 using group-query attention, devel-
393 oped by MistralAI (Jiang et al., 2023). In particular,
394 the 7B version has been reported to obtain better
395 performances when compared to Llama2-7B and
396 Llama2-13B. While details about the fine-tuning
397 data are lacking, in our experiments, we observe

⁷Refer to <https://huggingface.co/meta-llama/Meta-Llama-3-8B> and <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2>

398 that the model is responsive to Italian prompts. The
399 instruct-based versions of the models do not present
400 any moderation mechanism. We thus expect this
401 model to avoid over-safety and always return an
402 implied statement and a classification value.

5.1 Prompting Techniques 403

404 Among recent prompting techniques, we selected
405 **Chain-of-Thought** (CoT) and **Knowledge Aug-**
406 **mentation**. CoT was chosen for its notable success
407 in reasoning tasks (Lyu et al., 2023). On the other
408 hand, Knowledge Augmentation has been observed
409 to reduce hallucinations and enhance contextual
410 depth in model prompts, facilitating the generation
411 of sophisticated outputs beneficial for tasks requir-
412 ing substantial domain knowledge and nuanced rea-
413 soning (Kang et al., 2024). Both techniques align
414 with our goal of generating implicit components
415 of arguments (implicit warrants) and support the
416 construction of encoded warrant blocks. To the
417 best of our knowledge, these techniques have not
418 been used yet for a computational argumentation
419 task, which makes them worth investigating. The
420 full list of prompts can be found in Appendix C
421 and D. More in detail, CoT sequentially guides the
422 model through a series of reasoning steps before
423 arriving at a final answer or conclusion (Wei et al.,
424 2024). By following this structured approach, CoT
425 prompts allow the identification of how the model’s
426 reasoning process influences its conclusions. This
427 capability is particularly useful for reconstructing
428 warrants that underlie the model’s interpretations
429 in our specific task.

430 Knowledge-augmented prompting generates
431 knowledge from an LLM and incorporates it as
432 additional input for a task (Liu et al., 2022). In our
433 task, the generated knowledge serves as either the
434 implied assumption or the warrant that we inject
435 into the prompt to inform the classification.

6 Results 436

437 We report two blocks of results: the first block
438 focuses on **classification** of the messages. Since
439 both the Italian and the English datasets contain
440 only positive classes, we only report the Recall.
441 The classification task offers an indirect evaluation
442 on the goodness of the AR methods. The second
443 block targets the **generation** of the implied assump-
444 tions/warrants. Considering the complexity and the
445 pending issues related to the evaluation of automat-
446 ically generated text (Chang et al., 2024), we report

Setting	Model	ImplicIT-Mis	SBIC+
fine-tuning	bert-hateXplain	-	0.342
	ALBERTo	0.380	-
zero-shot	Llama3-8B	0.588	0.609
	Mistral-7B-v02	0.050	0.319
few-shot	Llama3-8B	0.738	0.719
	Mistral-7B-v02	0.259	0.416
zero-shot Assumption	Llama3-8B	0.542	0.448
	Mistral-7B-v02	0.050	0.259
few-shot Assumption	Llama3-8B	0.480	0.616
	Mistral-7B-v02	0.461	<u>0.685</u>
zero-shot Toulmin	Llama3-8B	0.557	0.452
	Mistral-7B-v02	0.346	0.374
few-shot Toulmin	Llama3-8B	<u>0.725</u>	0.594
	Mistral-7B-v02	0.556	0.604

Table 1: Classification results on ImplicIT and SBIC+. Best results in bold; second best underlined.

the results using established automatic metrics (i.e. BERTScore and BLEU) as well as a manual validation on a subset of 300 messages (150 per language) (§ 6.2). The overall evaluation procedure we have devised allows us to assess both the performance of the models’ in detecting implicit misogyny and the alignment between LLMs and human annotators in generating reasoning-based explanations.

All answers from LLMs have undergone post-processing to evaluate them properly. Two main post-processing heuristics concern the treatment of the “refusal to provide an answer” (including the refusal to generate the warrants) and the “need of more context”. We considered both cases as if the messages were marked as not misogynous. While Llama3-8B tends to return refusals to answers, mostly due to the safeguard layer, Mistral-7B-v02 has a tendency towards indecisive answers requiring more context. Llama3-8B always provides an answer when applied to the Italian data. For completeness, Appendix E includes the results considering these cases as correct.

6.1 Classification Results

Table 1 summarizes the results for the classification task. With few exceptions - mostly related to Mistral-7B-v02 - LLMs generally perform better than finetuned models. All few-shot experiments outperform their zero-shot counterpart, and Llama3-8B consistently performs better than Mistral-7B-v02. The best results are obtained by Llama3-8B with few-shot and no generation of either the implied statements or the warrants. How-

Setting	Model	BERTScore		BLEU	
		EN	IT	EN	IT
Assumption					
zero-shot	Llama3-8B	0.820	-	0.201	-
few-shot	Llama3-8B	0.830	-	0.744	-
	Mistral-7B-v02	0.823	0.601	0.361	0.240
Toulmin					
zero-shot	Llama3-8B	0.817	0.570	0.543	0.104
	Mistral-7B-v02	0.812	0.579	0.303	0.077
few-shot	Llama3-8B	0.817	0.570	0.871	0.261
	Mistral-7B-v02	0.813	0.601	0.396	0.313

Table 2: Automatic evaluation metrics for the best models generating implied assumptions/warrants (selection based on classification results).

ever, for Italian, the Llama3-8B with the Toulmin warrant in few-shot achieves very competitive results (R=0.725). For English, on the other hand, the results are affected by the post-processing heuristics. Had we considered as correct the “refusal to answer cases”, the best score for English would have resulted in Llama3-8B few-shot with implied assumption (R=0.913).

In all zero-shot settings, the prompt based on Toulmin’s warrant outperforms the prompt based on implied assumptions. In the few-shot settings, in ImplicIT-Mis, we observe a dramatic increase when switching from implied assumptions to Toulmin’s warrant, with a performance gain of 24 points. On the contrary, on English, the warrant-based prompt falls behind.

6.2 Implied Assumptions and Warrants Generation

Table 2 gives an overview of the evaluation using BERTScore and BLEU for the best models for English and Italian. While for SBIC+ every message has an associated explanation, for ImplicIT-Mis, only 150 messages present the implied assumptions. When Llama3-8B is asked to elaborate on the implied assumption in both zero- and few-shot settings, it does not follow the instruction, and only in 87 and 71 instances for Italian and English, respectively, generates a response. In all the other cases, the model just answers the final question of whether it is misogynistic; therefore, we exclude them from the evaluation. We also exclude all the results that do not reach at least a recall of 0.3 due to their low quality, as confirmed by manual inspection. All BERTScores in English are around 0.81-0.83, showing high similar content between the human-written texts and the answers generated by the models. Therefore, both the implied assump-

tions and the warrants are aligned with those written by humans. In Italian the scores drop to 0.57-0.60. In terms of BLEU scores, the highest scores for English are produced by Llama3-8B few-shots with warrants, which shows an alignment with humans in terms of word choices. For Italian the scores are much lower, probably because of many wrong translations and lack of Italian references which cause wrong inferences.

6.3 Manual Validation

We further validate the generated implied assumptions and warrants by manually exploring a subset of 300 messages, 150 per language. For ImplicIT-Mis, we use the manually annotated instances, while we randomly extract 150 instances for SBIC+. We focus only on the best models: Llama3-8B few-shots warrant for ImplicIT-Mis and Mistral-7B-v02 few-shots implied assumptions for SBIC+.

Overall, we find that 35% of the generated warrants for ImplicIT-Mis are correct and 32% lead to a correct classification of the messages. For SBIC+, the percentage of valid implied assumptions leading to a correct classification is 50%, while correct implied assumptions leading to a wrong classification are 52%. However, in Italian all the correctly predicted examples were actually predicted for the wrong reasons, while in English this happened 37% of the time. Therefore, we conclude that a correct explanation does not necessarily lead to a correct classification of misogyny, and this is always the case in the subsample we manually evaluated for Italian. This can be seen as an evidence that the model relies on their internalized knowledge and spurious correlations to address the task and shows no reasoning skills, since the Italian texts, being collected to address this task, requires much more reasoning to be understood.

We design a taxonomy to regroup all errors for both models. We identify seven kinds of common errors in warrant and implied assumption generation. Table 3 provides some examples. Notice that, although all error categories lead to wrong implied assumptions/warrants, we decide to keep a general “wrong inference” as a valid category for all the cases that do not fall under any other category or there is no evident reason.

Sarcasm/Irony This is a common error in English, due to the relatively high number of jokes in SBIC+. In these cases, the LLMs fail to capture the

Sarcasm/Irony	→ IT: 2; EN: 26
Example: <i>my wife gets angry when i call her pet names apparently sandwich making machine is sexist...</i>	
Generated: <i>The text mentions a personal interaction between the speaker and their wife, and does not make a generalization about women</i>	
Metaphorical and Figurative Language	→ IT: 14; EN: 1
Example: <i>Put a stuffed animal between them and they fight over it</i>	
Generated: <i>Men argue easily</i>	
Wrong Translations	→ IT: 10; EN: N/A
Example: <i>la lagna di mezzanotte [The midnight complain]</i>	
Generated: <i>There is a belief in witches and the existence of a “witches’ hour”. [...] The use of the term “witches” is likely a poetic or metaphorical way of describing challenges.</i>	
Opposite Intention	→ IT: 12; EN: 9
Example: <i>good girls don’t get raped.</i>	
Generated: <i>Women have control over whether or not they get raped.</i>	
Denial of Misogyny	→ IT: 3; EN: 4
Example: <i>doesn’t society’s acceptance of makeup for women give them a gender imbalance and greater attractiveness in aggregate over men, therefore inequality?</i>	
Generated: <i>Women are judged and valued primarily based on their physical appearance.</i>	
Lack of Reference	→ IT: 17; EN: 0
Example: <i>Pozzi era meglio, non cantava ma... il lavoro era lo stesso. [Pozzi was better, she didn’t sing but ... the job was the same]</i>	
Generated: <i>Pozzi is a better singer than the other one</i>	
Wrong Inference	→ IT: 42; EN: 34
Example: <i>i can’t decide if today is a good or bad day my ex got hit by a bus, but i lost my job as a bus driver</i>	
Generated: <i>The text does not provide any assumptions about women or gender. The text is about the speaker’s personal experiences and emotions.</i>	

Table 3: Error categories in warrant generation. For each category, we report an input example, the corresponding LLM generation, and the category’s distribution in Italian and English evaluation samples.

sarcastic/ironic intended meaning of the message and go for a more literal interpretation. 566 567

Metaphorical and Figurative Language. This category indicates a failure to interpret another level of non-literal meaning. We have observed a much more frequent occurrence in Italian - also because many messages use figurative or metaphorical expressions. As observed by Muti et al. (2024), misogyny in Italian is highly metaphorical, especially with references to animals. In Italian, not identifying metaphors could also be attributable to translation errors since metaphors are cultural-dependent. This highlights the complexity of cross-lingual implicit HS detection, as also pointed out by Korre et al. (2024), since the translation of a term often does not carry the same implications as in the source language. 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582

Wrong Translations. This is a category of errors that applies only to Italian. It comprises errors due to wrong translations of messages or to a lack of understanding of non-standard language, such as dialects and jargon expressions.

Opposite Intention. These errors could be considered an instance of LLM hallucinations (Maynez et al., 2020). In these cases, the models completely misinterpret the message’s content, resulting in generated implied assumptions that tend to support the message. These errors occur in both languages, with a slight preference for Italian.

Denial of Misogyny. This class of errors indicates a lack of connection between the generated implied assumptions and the answer of the model. In other words, the generated text is logical and correct as it clearly identifies the misogynous nature of the message. However, the model classifies the texts as non-misogynous.

Lack of Reference This kind of error has been observed only in Italian - mostly because very few English data contains direct reference to individuals. The errors in these cases are due to the model failing to identify the mentioned individuals (and their associated characteristics) that are used to trigger the misogynous content. In the example we report in Table 3, the model does not recognize (Moana) Pozzi as a famous porn actress, thus resulting in a warrant that fails to capture the insulting nature of the message (i.e., the target being promiscuous).

Wrong Inference This is the largest class of errors in both languages. We observe that wrong inferences are mainly driven by spurious correlations and the activation of implicit knowledge.

6.4 How Reasoning Impact Classification

We further examine how reconstructed implicit assumptions relate to classification predictions for misogyny detection.

We inspect 50 instances for each language that are correctly classified by our best baseline (few-shot Llama3-8B) and that are subsequently misclassified when the model is asked to reason about them. We observe the tendency to reduce a gender-related problem to the whole category of human beings, minimizing the misogynistic nature of the statement. For instance, in the Italian text “*Oggi sei felice anoressica bugiarda*” [Today you are happy anorexic liar], the model responds that “*Anorexia*

is a negative condition. 3. Lying is a negative behavior. No, the text is not misogynist. Anorexia is a condition that affects both men and women.”. Although being true that it is a condition that affects both genres, in this case, it targets a woman which is clear in Italian by the gendered “*anoressica*”, instead of “*anoressico*” which would be used for men. Adding the statement that it affects both genres is detrimental for the classification.

7 Conclusion

We proposed the task of implicit misogyny detection under an Argumentative Reasoning perspective, since to understand implicit statements, one needs to reconstruct the missing link (the warrant) between the claim and the assumption. Our work highlights the complexity of such a task, which paves the way for hate speech detection as a proxy task to probe the reasoning abilities of LLMs. Our prompt-based experiments show that LLMs fail 68% and 50% of the time in generating implied assumptions in Italian and English respectively. The poor relationship between wrongly generated explanations and correctly predicted classes shows LLMs’ over-reliance on their implicit knowledge and spurious correlations rather than reasoning skills. Our results are consistent with Zhu et al. (2023): prompting strategies that rely on implicit knowledge in LLMs often generate an incorrect classification when the generated knowledge (implied assumptions/warrants) is wrong, due to lack of references, reasoning skills, or understanding of non-standard language. Indeed, verifying the validity of the generated text before injecting it in the prompt in a human-in-the-loop approach would be a next step to undertake. To conclude, our findings show that *i*) the performance of the classification task cannot be used as a proxy to guarantee the correctness of the implied assumption/warrant; *ii*) LLMs do not have the necessary reasoning abilities in order to understand highly implicit misogynistic statements. Therefore, models for hate-related natural language inference tasks should be improved. One possible approach would be to inject external knowledge in the misogynous texts, in order to fill the gaps related to their lack of implicit knowledge. For instance, had the model known that Moana Pozzi was a porn actress, it would have probably inferred that when a person is compared to her, it is a derogatory way to address that woman.

681 Limitations

682 A limitation of our work is the integration of all gener-
683 ated knowledge (implied assumptions/warrants)
684 and we do not evaluate them before using them
685 to inform the classification task. This should be
686 overcome with a human-in-the-loop approach that
687 allows for the verification of the knowledge ex-
688 tracted by LLMs. We did not try to inject only the
689 knowledge that led to a correct classification be-
690 cause of the low correlation between the generated
691 implied statement and the class. Another limitation
692 is that for what concerns Llama, many examples in
693 English trigger the safeguard, therefore the scores
694 for Llama might not be realistic.

695 8 Ethical Considerations

696 Improving LLMs abilities to understand the im-
697 plied meaning of messages with sensitive con-
698 tent is a case of potential risks related to dual
699 use. Although our work has focused on assess-
700 ing LLMs abilities in generating implied assump-
701 tions/warrants, we see the benefits and the detri-
702 mental effects. On the one hand, improving LLMs
703 abilities to understand the implied meaning of sen-
704 sitive message can further be used to improve the
705 generation of counter-speech and the development
706 of assistive tools for experts in this area. At the
707 same time, the process can be inverted: malevolent
708 agents can feed models with implied assumptions
709 and generate hateful messages. We are aware of
710 this issue, and we think our work offers the com-
711 munity an opportunity to understand limitations of
712 LLMs that have a not minor societal impact. In ad-
713 dition to this, our work indicates the need to adopt
714 different safeguard methods that are able to capture
715 the core meaning of a message and grounded in
716 different cultures.

717 References

718 Gavin Abercrombie, Aiqi Jiang, Poppy Gerrard-abbott,
719 Ioannis Konstas, and Verena Rieser. 2023. [Resources](#)
720 [for automated identification of online gender-based](#)
721 [violence: A systematic review](#). In *The 7th Workshop*
722 *on Online Abuse and Harms (WOAH)*, pages 170–
723 186, Toronto, Canada. Association for Computational
724 Linguistics.

725 Manuela Barreto and David Matthew Doyle. 2023.
726 [Benevolent and hostile sexism in a shifting global](#)
727 [context](#). *Nature reviews psychology*, 2(2):98–111.

728 Maria Becker, Katharina Korfhage, and Anette Frank.
729 2020. [Implicit knowledge in argumentative texts: An](#)

[annotated corpus](#). In *Proceedings of the Twelfth Lan-
guage Resources and Evaluation Conference*, pages
2316–2324, Marseille, France. European Language
Resources Association. 730
731
732
733

Filip Boltužić and Jan Šnajder. 2016. [Fill the gap!](#)
[analyzing implicit premises between claims from online](#)
[debates](#). In *Proceedings of the Third Workshop on*
Argument Mining (ArgMining2016), pages 124–133,
Berlin, Germany. Association for Computational Lin-
guistics. 734
735
736
737
738
739

Leonard Bongard, Lena Held, and Ivan Habernal. 2022.
[The legal argument reasoning task in civil procedure](#).
In *Proceedings of the Natural Legal Language Pro-
cessing Workshop 2022*, pages 194–207, Abu Dhabi,
United Arab Emirates (Hybrid). Association for Com-
putational Linguistics. 740
741
742
743
744
745

ZeFeng Cai, Linlin Wang, Gerard de Melo, Fei Sun,
and Liang He. 2022. [Multi-scale distribution deep](#)
[variational autoencoder for explanation generation](#).
In *Findings of the Association for Computational*
Linguistics: ACL 2022, pages 68–78, Dublin, Ireland.
Association for Computational Linguistics. 746
747
748
749
750
751

Tommaso Caselli, Valerio Basile, Jelena Mitrović, and
Michael Granitzer. 2021a. [HateBERT: Retraining](#)
[BERT for abusive language detection in English](#). In
Proceedings of the 5th Workshop on Online Abuse
and Harms (WOAH 2021), pages 17–25, Online. As-
sociation for Computational Linguistics. 752
753
754
755
756
757

Tommaso Caselli, Valerio Basile, Jelena Mitrović, Inga
Kartozija, and Michael Granitzer. 2020. [I feel of-](#)
[fended, don't be abusive! implicit/explicit messages](#)
[in offensive and abusive language](#). In *Proceedings*
of the Twelfth Language Resources and Evaluation
Conference, pages 6193–6202, Marseille, France. Eu-
ropean Language Resources Association. 758
759
760
761
762
763
764

Tommaso Caselli, Roberto Cibin, Costanza Conforti,
Enrique Encinas, and Maurizio Teli. 2021b. [Guid-](#)
[ing principles for participatory design-inspired nat-](#)
[ural language processing](#). In *Proceedings of the 1st*
Workshop on NLP for Positive Impact, pages 27–35,
Online. Association for Computational Linguistics. 765
766
767
768
769
770

Tuhin Chakrabarty, Aadit Trivedi, and Smaranda
Muresan. 2021. [Implicit premise generation with](#)
[discourse-aware commonsense knowledge models](#).
In *Proceedings of the 2021 Conference on Empiri-*
cal Methods in Natural Language Processing, pages
6247–6252, Online and Punta Cana, Dominican Re-
public. Association for Computational Linguistics. 771
772
773
774
775
776
777

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu,
Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi,
Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang,
Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie.
2024. [A survey on evaluation of large language mod-](#)
[els](#). *ACM Trans. Intell. Syst. Technol.*, 15(3). 778
779
780
781
782
783

Guizhen Chen, Liying Cheng, Luu Anh Tuan, and Li-
dong Bing. 2023. [Exploring the potential of large](#)
[language models in computational argumentation](#).
CoRR, abs/2311.09022. 784
785
786
787

901	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L��lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth��e Lacroix, and William El Sayed. 2023. Mistral 7b . <i>Preprint</i> , arXiv:2310.06825.	
909	Minki Kang, Seanie Lee, Jinheon Baek, Kenji Kawaguchi, and Sung Ju Hwang. 2024. Knowledge-augmented reasoning distillation for small language models in knowledge-intensive tasks. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23</i> , Red Hook, NY, USA. Curran Associates Inc.	
916	Youngwook Kim, Shinwoo Park, and Yo-Sub Han. 2022. Generalizable implicit hate speech detection using contrastive learning . In <i>Proceedings of the 29th International Conference on Computational Linguistics</i> , pages 6667–6679, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.	
922	Youngwook Kim, Shinwoo Park, Youngsoo Namgoong, and Yo-Sub Han. 2023. ConPrompt: Pre-training a language model with machine-generated data for implicit hate speech detection . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 10964–10980, Singapore. Association for Computational Linguistics.	
929	Katerina Korre, Arianna Muti, and Alberto Barr��n-Cede��o. 2024. The challenges of creating a parallel multilingual hate speech corpus: An exploration . In <i>Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)</i> , pages 15842–15853, Torino, Italia. ELRA and ICCL.	
936	Katerina Korre, John Pavlopoulos, Jeffrey Sorensen, L��o Laugier, Ion Androutsopoulos, Lucas Dixon, and Alberto Barr��n-cede��o. 2023. Harmful language datasets: An assessment of robustness . In <i>The 7th Workshop on Online Abuse and Harms (WOAH)</i> , pages 221–230, Toronto, Canada. Association for Computational Linguistics.	
943	Anne Lauscher, Goran Glava��, and Simone Paolo Ponzetto. 2018. An argument-annotated corpus of scientific publications . In <i>Proceedings of the 5th Workshop on Argument Mining</i> , pages 40–46, Brussels, Belgium. Association for Computational Linguistics.	
949	John Lawrence and Chris Reed. 2019. Argument mining: A survey . <i>Computational Linguistics</i> , 45(4):765–818.	
952	Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. 2022. A new generation of perspective api: Efficient multi-lingual character-level transformers .	
956	Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh	
	Hajishirzi. 2022. Generated knowledge prompting for commonsense reasoning . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 3154–3169, Dublin, Ireland. Association for Computational Linguistics.	958 959 960 961 962 963
	Filipa Melo Lopes. 2019. Perpetuating the patriarchy: Misogyny and (post-)feminist backlash . <i>Philosophical Studies</i> , 176(9):2517–2538.	964 965 966
	Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning . In <i>Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.	967 968 969 970 971 972 973 974 975 976
	Eleonora Mancini, Federico Ruggeri, and Paolo Torroni. 2024. Multimodal fallacy classification in political debates . In <i>Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)</i> , pages 170–178, St. Julian’s, Malta. Association for Computational Linguistics.	977 978 979 980 981 982 983
	Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 1906–1919, Online. Association for Computational Linguistics.	984 985 986 987 988 989
	Arianna Muti, Federico Ruggeri, Cagri Toraman, Alberto Barr��n-Cede��o, Samuel Algherini, Lorenzo Musetti, Silvia Ronchi, Gianmarco Saretto, and Caterina Zapparoli. 2024. PejorativITY: Disambiguating pejorative epithets to improve misogyny detection in Italian tweets . In <i>Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)</i> , pages 12700–12711, Torino, Italia. ELRA and ICCL.	990 991 992 993 994 995 996 997 998 999
	Nicol��s Benjam��n Ocampo, Ekaterina Sviridova, Elena Cabrio, and Serena Villata. 2023. An in-depth analysis of implicit and subtle hate speech messages . In <i>Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics</i> , pages 1997–2013, Dubrovnik, Croatia. Association for Computational Linguistics.	1000 1001 1002 1003 1004 1005 1006
	Pia Pachinger, Allan Hanbury, Julia Neidhardt, and Anna Planitzer. 2023. Toward disambiguating the definitions of abusive, offensive, toxic, and uncivil comments . In <i>Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)</i> , pages 107–113, Dubrovnik, Croatia. Association for Computational Linguistics.	1007 1008 1009 1010 1011 1012 1013

- 1128 *Conference on Natural Language Processing (Vol-*
1129 *ume 1: Long Papers)*, pages 3181–3197, Online. As-
1130 sociation for Computational Linguistics.
- 1131 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q.
1132 Weinberger, and Yoav Artzi. 2020. [Bertscore: Evalu-](#)
1133 [ating text generation with BERT](#). In *8th International*
1134 *Conference on Learning Representations, ICLR 2020,*
1135 *Addis Ababa, Ethiopia, April 26-30, 2020*. OpenRe-
1136 view.net.
- 1137 Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang,
1138 Yiqun Liu, and Shaoping Ma. 2014. [Explicit factor](#)
1139 [models for explainable recommendation based on](#)
1140 [phrase-level sentiment analysis](#). In *Proceedings of*
1141 *the 37th International ACM SIGIR Conference on*
1142 *Research & Development in Information Retrieval,*
1143 *SIGIR '14*, page 83–92, New York, NY, USA. Asso-
1144 ciation for Computing Machinery.
- 1145 Fangwei Zhu, Jifan Yu, Hailong Jin, Lei Hou, Juanzi
1146 Li, and Zhifang Sui. 2023. [Learn to not link: Explor-](#)
1147 [ing NIL prediction in entity linking](#). In *Findings of*
1148 *the Association for Computational Linguistics: ACL*
1149 *2023*, pages 10846–10860, Toronto, Canada. Associ-
1150 ation for Computational Linguistics.

A Implicit-Mis Sources

Table A shows statistics on the number of Facebook comments associated to each newspaper or Facebook community.

Source	Messages
National News	
La Repubblica	411
Il Messaggero	378
La Stampa	76
TgCom24	20
Libero	1
Local news	
AnconaToday	20
BolognaToday	9
Corriere Adriatico	2
Palermolive.it	5
Online news	
Donna Fanpage	37
Fanpage	33
Huffington Post	6
TPI	4
Il Post	1
Leggo	1
FB Community	
Caffeina Festival	65
Non sono bello ma spaccio	15
La matita scarlatta	9
Pastorizia never dies	9
Stefano Valdegamberi	6
I love Patriarcato 2	4
La società femminista	4
L'uomo che bestemmiava ai cavalli	3

Table A: List of sources - newspapers and Facebook pages - with total amount of extracted instance for the creation of the Implicit-Mis dataset

B Implicit-Mis Annotation Guidelines for Implied Assumptions

Annotators were given the following guidelines for the identification of the implied statements in the Implicit-Mis dataset:

Read the following comments extracted from Facebook as a reaction to a news article or a post. The comments have been judged as being sexist or misogynistic by an expert. Your task is to provide an explanation for why the comment can be offen-

sive towards. The explanation must correspond to the implied assumption in the comment. To help you in this task, imagine that your explanations should complete the sentence: "This comment is sexist because it implies that ...". Follow the examples. Le dovrebbero dare 2000 volts implica che la donna in questione dovrebbe essere uccisa. Non dovremmo abbassare i nostri standard solo per assumere più donne implica che le donne sono meno qualificate degli uomini. Quella mangia kinder bueno non meritava di vincere X-Factor implica che la donna in questione viene offesa in quanto grassa. Bau bau implica che la donna in questione viene chiamata cagna.

*Translated version: She should be given 2000 volts implied that the targeted woman should be killed. We shouldn't lower our standards just to hire more women implies that women are less qualified than men. That Kinder Bueno eater did not deserve to win X-Factor implies that the targeted woman is called 'fat'. Wof wof implies that the targeted woman is called 'cagna', translated as female dog, which is an epithet for b**ch.*

C English Prompts

Prompt	Prompt Details
Zero-shot	Is this text misogynous? Answer 1 for yes and 0 for no.
Few-shot	<p>Is the text misogynous? Answer only with YES or NO. Answer like in the examples. Text: we shouldn't lower our standards just to hire more women. Answer: YES Text: But since only women get pregnant, they get free birth control pills. Answer: NO Text: I can't stand women talking about football. Answer: YES Text: Lmao I'm down with women replacing God. The future is female Answer: NO Text: She must be on her period. Answer: YES Text: Dude, girls want sex as much as guys do. They have triggers, so do we. Answer: NO Text: {text} Answer:</p>
Implied Assumption Zero-shot	<p>What are the implied assumptions of the text text? Based on the implied assumptions, is the text misogynous? Answer only with 1 for yes and 0 for no. Answer between the tags. Implied assumption: <...> Answer: <...></p>
Implied Assumption Few-shot	<p>What are the implied assumptions of the text {text}? Based on the implied assumptions, is the text misogynous? Answer only with YES or NO. Answer like in the examples. Text: we shouldn't lower our standards just to hire more women. Implied assumption: women are less qualified than men. Answer: yes. Text: Women shouldn't talk about football. Implied assumption: women cannot understand football, it's a male thing. Answer: YES. Text: She must be on her period. Implied Assumption: women are irrational or overly emotional during their periods. Answer: YES Text: Implied assumption: Answer:</p>
Toulmin CoT Zero-Shot	<p>What is the claim and the implied warrant of the text {text}? Explain step by step. Based on the implied warrant, is the text misogynous? Answer only with 1 for yes and 0 for no. Warrant: Misogyny:</p>
Toulmin CoT Few-Shot	<p>What are the claim and the implied warrant in the input text? Explain step by step. Based on the implied warrant, is the text misogynous? Answer only with YES or NO. Follow the example. Text: we shouldn't lower our standards just to hire more women Claim: hiring more women would lower our standards Warrant: women are less qualified than men Answer: YES Text: I can't stand women talking about football Claim: women should not talk about football Warrant: football it's for men Answer: YES Text: She must be on her period. Claim: She must be on her period Warrant: women are irrational or overly emotional during their periods Answer: YES Text: {text} Claim: Warrant: Answer:</p>

Table C: English prompts used in our experiments.

D Italian Prompts

Prompt	Prompt Details
Zero-shot	Questo testo è misogino? Rispondi solo con 1 (sì) o 0 (no)
Few-shot	Questo testo è misogino? Rispondi solo con 1 (sì) o 0 (no) come nell'esempio. Testo: quella mangia kinder bueno non meritava di vincere X-Factor. Misogino: 1. Testo: Lei è acida perché non ha figli penso che darebbe fastidio a qualsiasi donna. Che schifo Risposta: 0 Testo: le donne che parlano di calcio non si possono sentire. Risposta: 1 Testo: Comunque le pringles più buone sono quelle alla panna acida e cipolla Risposta: 0 Testo: Dopo aver visto (per sbaglio) 5 minuti di Uomini e Donne ho deciso che odierò le donne. Risposta: 1 Testo: pur abitando in culo alla balena, sono andata ad un bel po' di concerti e sono ricordi preziosi che conservo con affetto Risposta:0 Testo: {text} Risposta:
Implied Assumption Zero-shot	Quali sono gli assunti impliciti del testo {testo}? Sulla base degli assunti impliciti, il testo è misogino? Rispondere solo con SÌ o NO.
Implied Assumption Few-shot	Quali sono gli assunti impliciti del testo {testo}? Sulla base dei presupposti impliciti, il testo è misogino? Rispondere solo con SÌ o NO. Seguire l'esempio. Testo: non dovremmo abbassare i nostri standard solo per assumere più donne. Presupposto: le donne sono meno qualificate degli uomini. Risposta: SÌ. Testo: le donne che parlano di calcio non si possono sentire. Presupposto: le donne non capiscono niente di calcio, è una cosa da maschi. Risposta: SÌ. Testo: Dopo aver visto (per sbaglio) 5 minuti di Uomini e Donne ho deciso che odierò le donne. Presupposto: Le donne che vanno alla trasmissione Uomini e Donne sono stupide. Risposta: SÌ. Testo: {testo} Presupposto: Risposta:
Toulmin CoT Zero-Shot	Quali sono il claim e il warrant implicito del testo? Spiegalo passo dopo passo. In base al warrant implicito, il testo è misogino? Rispondi solo con 1 per il sì e 0 per il no. Warrant: Misoginia:
Toulmin CoT Few-Shot	Quali sono il claim e il warrant implicito nel testo? Spiegalo passo per passo. In base al warrant implicito, il testo è misogino? Rispondere solo con SÌ o NO. Segui l'esempio. Testo: non dovremmo abbassare i nostri standard solo per assumere più donne. Affermazione: assumere più donne abbasserebbe i nostri standard Warrant: le donne sono meno qualificate degli uomini Risposta: SÌ Testo: Non sopporto che le donne parlino di calcio Affermazione: le donne non dovrebbero parlare di calcio Warrant: il calcio è per gli uomini Risposta: SÌ Testo: Deve avere il ciclo. Affermazione: deve avere le mestruazioni Warrant: le donne sono irrazionali o eccessivamente emotive durante il ciclo mestruale Risposta: SÌ Testo: {testo} Affermazione: Warrant: Risposta:

Table D: Italian prompts used in our experiments.

E Additional Classification Results

Table E reports classification results when considering the refusal to answer due to model safeguard trigger to hateful content as misogynous. In particular, Llama3-8B is the only affected model in our experiments.

Exp. Setting	Model	ImplicIT-Mis	SBIC+
fine-tuning	bert-hateXplain	–	0.342
	ALBERTo	0.380	–
zero-shot	Llama3-8B	0.588	0.609
	Mistral-7B-v02	0.050	0.319
few-shot	Llama3-8B	0.738	<u>0.827</u>
	Mistral-7B-v02	0.259	0.416
zero-shot w. implied assumption	Llama3-8B	0.542	<u>0.891</u>
	Mistral-7B-v02	0.050	0.259
few-shot w. implied assumption	Llama3-8B	0.480	0.914
	Mistral-7B-v02	0.461	0.685
zero-shot Toulmin warrant	Llama3-8B	0.557	<u>0.643</u>
	Mistral-7B-v02	0.346	0.374
few-shot Toulmin warrant	Llama3-8B	0.725	<u>0.841</u>
	Mistral-7B-v02	0.556	0.604

Table E: Overview of the results of the experiments on ImplicIT and SBIC+. Best results are in bold, while performance differences with respect to 1 are underlined. Answer considered valid with implied assumption/Toulmin’s warrant only if the model generates the implied assumptions/warrants.