

Automatically Discovering How Misogyny is Framed on Social Media

Rakshitha Rao Ailneni and Sanda M. Harabagiu

Human Language Technology Research Institute, The University of Texas at Dallas
{rxa220074, sanda}@utdallas.edu

Abstract

Misogyny, which is widespread on social media, can be identified not only by recognizing its many forms but also by discovering how misogyny is framed. This paper considers the automatic discovery of misogyny problems and their frames through the novel Dis-MP&F method, which enables the generation of a data-driven, rich Taxonomy of Misogyny (ToM), offering new insights in the complexity of the way misogyny is expressed. Furthermore, the reported Dis-MP&F method, informed by the ToM, is capable of producing very promising results on a misogyny benchmark dataset.

1 Introduction

Abusive language used on social media is pervasive, becoming a wide-spread phenomenon that has serious consequences for its victims. Misogyny, defined as hatred of, aversion to, or prejudice against women, is no exception. It is important to identify misogyny in Social Media Postings (SMPs), not only for protecting women, but also for eventually generating counter narratives. However, the identification of misogyny is complicated by the fact that (1) there are many forms of misogyny, each addressing another *Misogyny Problem* (MP), e.g. woman shaming, woman stereotyping, etc; and (2) each misogyny problem can be *framed* in multiple ways.

Figure 1 illustrates two different misogynistic SMPs, addressing the same MP, namely *women stereotyping*. Women stereotyping is defined by the application of broad, oversimplified, and often negative generalizations about women based on their gender. An examination of the content of the SMP shown in Figure 1(a) reveals that it evokes the first Frame of Misogyny (FoM) illustrated in Figure 1, namely FoM₁, while the SMP from Figure 1(b) evokes the second FoM, namely FoM₂. While the automatic detection of MPs is important for understanding the different forms of misogyny that are



Figure 1: Example of two misogynistic Social Media Postings (SMPs) that address the same Misogyny Problem (MP), but evoke different Frames of Misogyny (FoMs).

spreading on social media, the detection of the way misogyny is framed is also important, as it informs counter narratives addressing misogyny, cf (Chung et al., 2019; Guest et al., 2021).

Unlike previous work (Anzovino et al., 2018; Parikh et al., 2019; Guest et al., 2021; Zeinert et al., 2021) that detects misogyny in social media by focusing on the recognition of the MPs addressed in SMPs, we are also interested in the discovery of the FoMs that articulate the cause of the MP. To discover FoMs like those shown in Figure 1, we

considered the widely cited definition of framings, Entman (1993a), which notes that “to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described.” Based on this definition, the MPs represent the aspects of misogyny that receive a causal interpretation in the articulation of FoCs.

Recently, (Weinzierl and Harabagiu, 2024) have introduced a method for automatically discovering and articulating frames of communication evoked in SMPs. However, this method, based on Chain-of-Thought (CoT) prompting (Wei et al., 2023) of Large Language Models (LLMs) assumes that we are aware of all the MPs. However, an inspection of the previous work aiming misogyny detection indicates that each method has considered a different set of MPs, and some have even annotated taxonomies of MPs, e.g. (Zeinert et al., 2021). Consequently, there is no consistent or complete set of MPs. In this paper we present a method that Discovers Misogyny Problems and the way misogyny is Framed on social media by (1) identifying automatically all the MPs addressed in a dataset of SMPs; (2) generating a taxonomy of MPs; and (3) articulating all the FoMs addressing each MP, as evoked in SMPs. We call this method **Dis-MP&F**, allowing us to make the following contributions:

<1> We introduce the first automatic method that identifies all the Misogyny Problems (MPs) that are discussed in a collection of SMPs, by prompting a Large Multimedia Model (LMM). The few-shot framework used by the prompting methods of Dis-MP&F eliminates the need of expensive annotations of various forms of misogyny on collections of SMPs.

<2> We describe the first method of automatically organizing MPs into taxonomies, which resulted in the largest misogyny taxonomy yet considered.

<3> We present first method that is able to discover and articulate Frames of Misogyny (FoMs) without being aware of all the MPs discussed in a collection of SMPs. Furthermore, this method produces new state-of-the-art (SOTA) results on a benchmark dataset. We make available all discovered MPs, the Taxonomy of Misogyny (ToM) and the FoMs that were discovered, as well as all prompts that we used on GitHub¹.

¹<https://github.com/rak55/Auto-Taxonomy-of-Misogyny>

| | |
|---|-----------------------------------|
| SMP = Social Media Posting | MP = Misogyny Problem |
| FoM = Frame of Misogyny | ToM = Taxonomy of Misogyny |
| CoT = Chain of Thought | LLM = Large Language Model |
| Dis-MP&F = Discovery of Misogyny Problems & Frames | |
| RTMP = Reference Taxonomy of Misogyny Problems | |

| |
|--|
| 5 different CoT (Chain of Thought) prompting methods: |
| 1) CoT^{FoM} = CoT prompting for discovering FoMs |
| 2) CoT^{MP} = CoT prompting for correcting the discovered MPs |
| 3) CoT^A = CoT prompting for discovering analogous MPs |
| 4) CoT^H = CoT prompting for uncovering the hierarchies of MPs |
| 5) CoT^{Ret} = CoT prompting for discovering relations between FoMs |

Figure 2: Acronyms used in the paper.

2 Related Work

All previous work aiming the identification of misogyny in SMPs relied on various annotations of MPs. For example, (Anzovino et al., 2018) has cast misogyny across five MPs: (1) *Discredit*, (2) *Harassment & Threats of Violence*, (3) *Derailing*, (4) *Stereotype & Objectification*, and (5) *Dominance*. Given these annotations, their method has considered a set of features that combined N-grams, bag-of-POS tags, word embeddings as well as features characterizing the SMPs, e.g. the length of the tweet, number of adjectives used in the tweets, etc. These features informed several classifiers using (a) Support Vector Machines (SVM); (b) Random Forest (RF); (c) Naive Bayes (NB); and (d) and a Multi-layer Perceptron Neural Network (MPNN). The best performance was obtained when using the classifier implementing the SVM, with an accuracy of 79.95% and an F1-score of 38.25%.

Another multi-label approach for the discovery of MPs was taken in (Parikh et al., 2019) where SMPs from the “Everyday Sexism Project” were annotated with as many as 23 different MPs, e.g. body shaming, gaslighting, mansplaining, victim blaming. A sequential neural architecture that uses LSTMs and CNNs was experimented with to identify MPs, obtaining accuracy as high as 63.6% and an F1-score of 75.3%.

A taxonomy for misogyny labeling applied to a corpus of primarily English Reddit posts was presented in (Guest et al., 2021). In this taxonomy, there are four overarching MPs: (i) *Misogynistic Pejoratives*, (ii) *Descriptions of Misogynistic Treatment*, (iii) *Acts of Misogynistic Derogation* and (iv) *Gendered Personal attacks against women*. Methods based on BERT (Devlin et al., 2019) were able to identify MPs on this dataset with accuracy of 93% and F1-score of 43%.

The SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification (MAMI) (Fersini

| MP | No. of SMPs |
|------------------------|-------------|
| <i>Stereotyping</i> | 2810 |
| <i>Objectification</i> | 2202 |
| <i>Shaming</i> | 1274 |
| <i>Violence</i> | 953 |

Table 1: Distribution of MPs in the MAMI training dataset.

et al., 2022) introduced a benchmark dataset for misogyny detection, detailed in Section 3. During SemEval-2022, 41 teams participated in the challenge of automatically identifying the MPs in the MAMI dataset. The best F1-score for MP identification was 73.1%. Some of the most successful methods have implemented neural models grounded on CLIP encoding (Radford et al., 2021) of the images present in the SMPs combined with text encoding generated with LSTM-based sequential models.

In contrast, in this paper we address the problem of discovering all the MPs that can be *inferred* from a dataset of SMPs. For this reason, we use several promptings of an LMM and of an LLM to discover not only the MPs addressed throughout SMPs, but also the FoMs that are evoked. In addition, we generate automatically the largest taxonomy of MPs, which we derived from the MAMI dataset.

3 The Dataset

The MAMI dataset, released for the SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification (MAMI) (Fersini et al., 2022) contains 11,000 SMPs using memes, divided into a training dataset containing 10,000 memes; and a test set containing 1,000 memes. The SMPs originate from various social media platforms like X / Twitter, Reddit, 9GAG, KnowYourMeme, and Imgur. The SMPs were annotated with four MP labels:

◊1◊ *Shaming*, which is defined as the practice of criticising women who violate expectations of behaviour and appearance regarding issues related to gender typology or related to physical appearance (Zhang et al., 2024).

◊2◊ *Stereotyping*, which is defined as the practice of assigning a fixed, conventional idea or set of characteristics to a woman. The MAMI annotations considered a stereotype as a fixed, conventional idea or set of characteristics assigned to a woman (Eagly and Mladinic, 1989). Table 1, which lists the distribution of SMPs for each of the MPs annotated in the MAMI dataset, shows that the stereotyping MP is the prevalent MP in the MAMI

dataset.

◊3◊ *Objectification*, which is defined as a practice of seeing and/or treating a woman like an object (Szymanski et al., 2011). Degrading women to the status of an object is captured by this MP.

◊4◊ *Violence* which is defined by physical, emotional, psychological, sexual, or economic harm directed at women or girls, driven by gender-based hatred, control, or discrimination (Andreassen, 2020).

A total of 7,220 of the 11,000 SMPs were annotated with MPs. A closer inspection of these annotations indicated that 2893 of the SMPs were annotated with only one MP, while 1342 SMPs were annotated with two different MPs. A total of 475 SMPs were annotated with three different MPs and only 475 SMPs were annotated with all four MPs considered in this dataset.

The inter-annotator agreement of the labeling produced in MAMI, as reported in (Fersini et al., 2022), was computed using the the Fleiss K -coefficient (Fleiss, 1971), resulting in a value of 0.3373. This very low value of the K -coefficient motivated us to consider a misogyny discovery method that does not have to rely on these annotations, but that can instead use the reasoning capabilities of LMMs to discover the MPs as well as the FoMs.

4 The Method

The DiS-MP&F method operates in three phases:

- ◊ In **Phase 1**, it discovers and articulates FoMs by using Chain-of-Thought (CoT) prompting (Wei et al., 2023) of a Large Multimedia Model (LMM).
- ◊ In **Phase 2**, the MPs that were revealed in the rationales generated by the LMM in Phase 1 are identified and verified whether they reflect a woman’s perspective. Furthermore, MPs which are analogous are also recognized and filtered out. But most importantly, MPs are organized into an ample Taxonomy of Misogyny (ToM).
- ◊ In **Phase 3**, all FoMs that address the same MP are further organized, by automatically recognizing the relations that connect them.

It should be noted that in all three phases, multiple forms of prompting LMMs and LLMs are used: one prompting method is used in Phase 1, three prompting methods are used in Phase 2 and one prompting method is used in Phase 3.

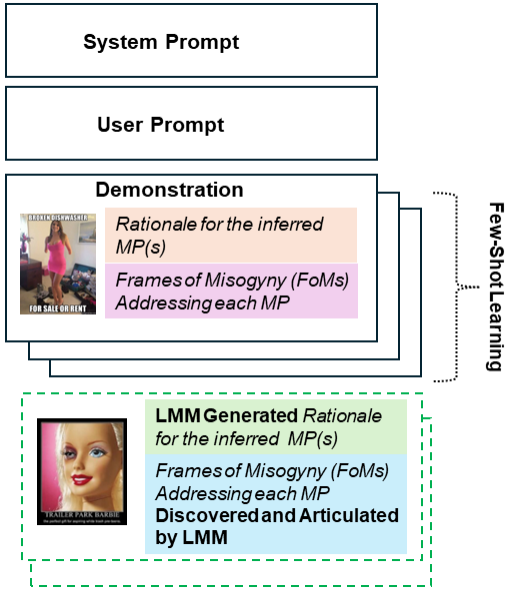


Figure 3: CoT^{FoM} Prompting for Phase 1.

4.1 Phase 1 of Dis-MP&F

The only existing method capable to discover and articulate frames (Weinzierl and Harabagiu, 2024) operates only on the text of SMPs. To discover FoMs from SMPs containing memes, we decided to also use Chain-of-Thought (CoT) prompting of an LMM in a few-shot setting, illustrated in Figure 3. This CoT^{FoM} prompting facilitates the identification of FoMs by breaking down the rationale into distinct problems, guiding the model through a step-by-step analysis that mirrors human reasoning. The System Prompt, detailed in Appendix A, defines FoMs as well as misogyny in general. The User Prompt, also detailed in Appendix A, instructs the LMM to (a) decide based on its reasoning if the meme contained in the SMP evokes any FoM; and (b) to articulate the FoM. Since multiple FoMs can be evoked by a single meme, the LMM is explicitly directed to generate all applicable FoMs in the User Prompt.

The few-shot learning used in CoT^{FoM} consists of five demonstrations which show the LMM (1) a meme; (2) the rationales for the MPs inferred from the meme; and (3) the FoM that were articulated to give a causal interpretation to each of the MPs. The selection of the SMPs used in the demonstrations of CoT^{FoM} is completely different than in the CoT prompting of (Weinzierl and Harabagiu, 2024). This is due to the fact that we did not rely on curriculum learning, but instead we selected the demonstrations based on the observation that 58% of the SMPs from the MAMI dataset were



Rationale:

- 1) The problem of stereotyping arises as the meme compares women to a dishwasher, reinforcing the idea that they belong to the kitchen.
- 2) The problem of violence arises as the meme normalizes and downplays the severity of physical abuse within the context of household roles.

Frames of Misogyny:

- 1) Women are primarily responsible for domestic labor.
- 2) Violence is an acceptable and easy solution when a woman fails to perform her domestic duties.

Figure 4: Demonstration example.

annotated with either one or two MPs. This motivated our decision to select for demonstrations only SMPs that were annotated only with one or two MPs that were predominant in the dataset. As detailed in Appendix B, among the SMPs that received only one annotation, those that were labeled with the MP of *Stereotyping*, *Objectification* and *Shaming* were predominant. Therefore, from each of these subsets of SMPs we selected randomly an SMP, for which we created a demonstration. In Appendix B we also show the the distribution of SMPs that received two MP annotations. It can be noted that the SMPs that were annotated both with *Objectification* and *Stereotyping*, as well as those that were annotated both with *Stereotyping* and *Violence* are predominant. Therefore we created two additional demonstrations by randomly selecting one SMP from each predominant sub-class of SMPs with two MP annotations.

Figure 4 illustrates one of the demonstrations that we provided to the LMM in the CoT^{FoM} prompting. Because the meme used in the demonstration was annotated with $MP_1 = \textit{Stereotyping}$ and $MP_2 = \textit{Violence}$, we illustrate also the rationales for the MPs. In addition, we provide two FoMs, where FoM_1 addresses the causal interpretation of MP_1 , while FoM_2 addresses the causal interpretation of MP_2 .

A second notable difference of CoT^{FoM} from the CoT prompting of (Weinzierl and Harabagiu, 2024) stems from the fact that we did not employ any kind of active-learning. When we experimented with two different LMMs, namely GPT-4o

(OpenAI, 2024) and LLaVA-1.5 (Liu et al., 2023), we noticed that GPT-4o performed the best and it did so without any human intervention. When the first 100 unseen memes were presented, it produced rationales for the MPs as well as it articulated FoMs that did not need any editing.

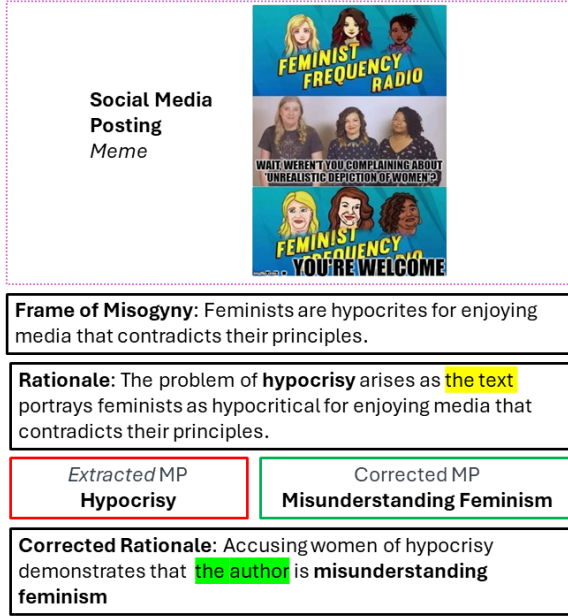


Figure 5: Example of corrected Misogyny Problem.

4.2 Phase 2 of Dis-MP&F

Phase 2 of the Dis-MP&F consists of four steps:

Step 1: Extraction and validation of MPs We extract the MPs addressed by the LMM in its rationales for the 9192 FoMs discovered in Phase 1 and have assembled an initial list $L_{MP}^0 = [MP_1, MP_2, \dots, MP_{N_0}]$, of unique MPs, where $N_0 = 224$. The validation of any MP_i from L_{MP}^0 was performed by trying to find if those MPs have been discussed in academic papers focusing on misogyny. For this purpose, we queried the Semantic Scholar search engine (Lo et al., 2020) with a query $Q = [MP_i \text{ AND } \textit{misogyny}]$. If among the retrieved papers there were papers having the title that combines the two terms of Q , we assumed that MP_i is validated. Furthermore, the inspection of the content of such retrieved papers allowed us to gather citations from literature that define the MP. Only 99 MPs from L_{MP}^0 were validated. These MPs were assembled in the list L_{MP}^V .

Step 2: Correction of non-valid MPs All MPs that have not been validated have however been used in the rationales generated in Phase 1. A closer inspection of the articulated FoMs and

their non-valid MPs generated in Phase 1 showed that sometimes the MPs needed to be corrected by replacing them with some MP from L_{MP}^V . For instance, as illustrated in Figure 5, when extracting the MP from the rationale of the FoM, we obtained the non-valid MP of 'hypocrisy'. This MP illustrates the point-view of the meme's author, who created a misogynistic SMP. The author sarcastically asks in the text superimposed on the posted meme *Wait, didn't you complain about unrealistic depictions of women?*. The correct MP for this meme captures a woman's view-point, i.e. 'misunderstanding feminism', because of accusing women of hypocrisy for daring to enjoy certain media. Figure 5 also shows the corrected rationale, highlighting the role of the author while enabling the extraction of the correct MP.

We further addressed the validity of the MPs by engaging three female graduate students that did not participate in the design of the method. These students provided interpretations and evaluations of the MPs from the female point of view. They were tasked with evaluating the validity of the MPs extracted in Step 1. For non-valid MPs, they were instructed to select a replacement from the list L_{MP}^V . Detailed instructions provided to the annotators are included in Appendix D. The inter-annotator agreement for the correctness of the MPs, computed with Fleiss' Kappa (Fleiss, 1971) was 0.708. Then we randomly selected small number ($=3$) of FoMs addressing a corrected MPs upon which all three students unanimously agreed. The rationales for the corrected MPs were used as demonstrations to guide the correction of all non-valid MPs. To achieve this, a tailored Chain-of-Thought (CoT) prompt, denoted as CoT^{CMP} and detailed in Appendix C, was employed. When presented with any FoMs discovered in Phase 1 using the CoT^{CMP} prompt, addressing an MP_x , the LLM generates a correction of the MP_x if needed, replacing it with some $MP_y \in L_{MP}^V$, also providing a rationale for the correction.

Step 3: Detection of analogous MPs The identification of analogous MPs, that may be expressed as paraphrases or near-paraphrases is performed through CoT^A prompting, detailed in Appendix C. The CoT^A prompt presents only MPs to the LLM, which generates for each valid MP all its possible paraphrases. For each set of paraphrased MPs, S_{MP}^A , we select the MP addressed by the largest number of FoMs to substitute all MPs in S_{MP}^A . In this way we obtained a consolidated list of MP,

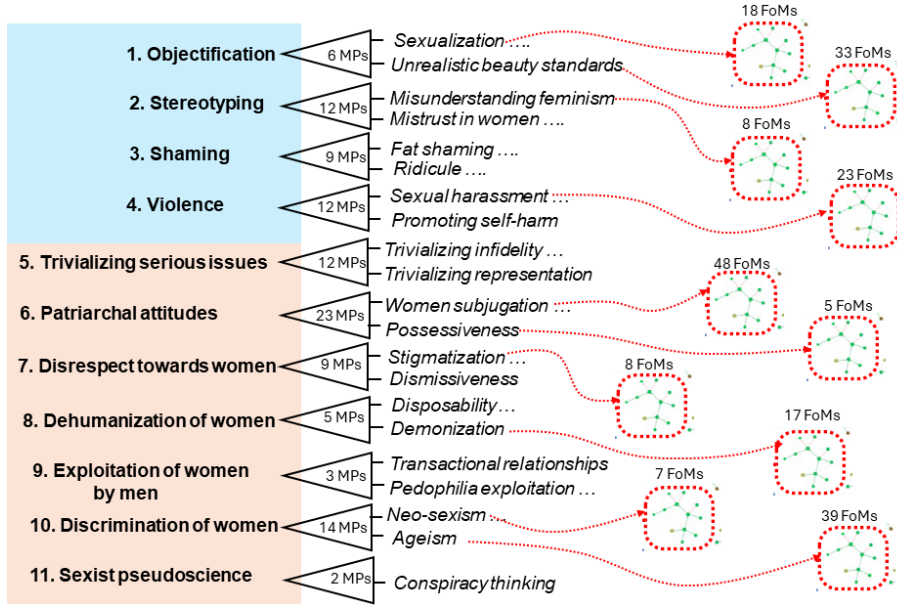


Figure 6: The Taxonomy of Misogyny.

denoted as L_{MP}^C . This list contains 99 different MPs. To note that each MP from L_{MP}^C is linked to all the FoMs (discovered in Phase 1) that it, or its analogous MPs, address.

Step 4: Generation of a Hierarchy of MPs

(shown in Figure 6). The L_{MP}^C list (of 99 MPs) is much larger than the list of 4 MPs originally annotated in the MAMI dataset. Not only that L_{MP}^C contains the 4 MPs originally annotated, but it also contains many additional MPs. Therefore, the MPs from L_{MP}^C were organized conceptually in the Taxonomy of Misogyny Problems (ToM) which is fully detailed in Appendix G and sketched in Figure 6. Note that the ToM has eleven different hierarchies, four of which are rooted in the MP originally considered in the MAMI dataset annotations. Figure 6 also shows the number of MPs distributed across all hierarchies, and it illustrates a couple of examples of MPs from each hierarchy. Additionally, each MP from the ToM is linked to the FoMs that address it. These FoMs are structured as graphs, because of the relations between FoMs discovered in Phase 3 of the Dis-MP&F method. Figure 6 also lists the number of FoMs connected to MPs from various hierarchies of the ToM.

To generate the ToM, the LLM was prompted with CoT^H in a few-shot setting. We have selected 3 examples of MP pairs (MP_x, MP_y) where MP_y specializes MP_x , e.g. $MP_x = \text{Stereotyping}$ and $MP_y = \text{Reductionism}$, providing them along with demonstrations in the CoT^H prompt, which is de-

tailed in Appendix C.

4.3 Phase 3 of Dis-MP&F

As in (Weinzierl and Harabagiu, 2024), we were interested to discover the binary relations that span FoMs. However, only relations between FoMs addressing the same MP were sought. Because FoMs provide a causal interpretation of MPs, we considered that when two FoMs address the same MP using different language, but with consistent underlying causes, they should share a *Paraphrase* relation. When a pair of FoMs address the same MP, but they express opposing causes of the MP, they should share a *Contradiction* relation. Examples of these relations between FoMs are shown in Figure 7.

The automatic detection of relations between FoMs was implemented by another few-shot CoT prompting of an LLM, namely CoT^{rel} . First, two examples of paraphrased pairs of FoMs and two examples of contradiction pairs were presented to the LLM as demonstrations, along with their rationales. We also provided two examples when no relations between a pair of FoMs can be established. Then, for any FoM_a , the top-k most similar FoMs addressing the same MP were selected in S_{FoM}^a , based on the distance provided by Sentence-BERT (SBERT) (Reimers and Gurevych, 2019). Secondly, for each pair $[FoM_a, FoM_b]$, where $FoM_b \in S_{FoM}^a$, the LLM was prompted to decide whether any relation exists between the pair of FoMs, and if they do, to predict the kind

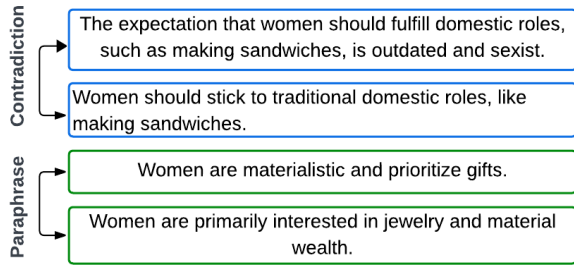


Figure 7: Example of relations identified by GPT-4o between the FoCs.

of relation. The details of the CoT^{rel} prompt are provided in Appendix E.

5 Experimental Results

In our experiments we employed GPT-4o as the LMM prompted with CoT^{FoM} , for discovering and articulating FoMs in Phase 1 of Dis-MP&F. We also used GPT-4o as the LLM prompted in Phase 2 of Dis-MP&F (a) with the CoT^{C_MP} prompt to correct the extracted MPs; (b) with the CoT^A prompting, which consolidates the MPs into L_{MP}^C ; and (c) with CoT^H prompt to discover all hierarchical relations between MPs, informing the ToM. GPT-4o was again used in Phase 3 of Dis-MP&F with the CoT^{rel} prompt, which led to the discovery of relations between FoMs. All prompting experiments were performed through the OpenAI API, utilizing the same hyperparameters: 512 max generated tokens, a temperature of 1.0, and 0.7 top-p sampling.

Quantitative results: Phase 1 of Dis-MP&F resulted in the articulation of 9,192 FoMs. In Phase 2 of Dis-MP&F 224 MPs were extracted from these FoMs. 48 of them were corrected due to CoT^{C_MP} prompt and 77 were found to be analogous, resulting in a list of 99 unique MPs. These MPs were organized in the ToM due to the discovery of 88 problem-subproblem relations facilitated by the CoT^H prompt. The ToM has a maximum depth of 3, but it also contains 11 unique tops of the hierarchy, therefore greatly enhancing the understanding of misogyny through its problems, from the original four MPs considered in the MAMI dataset. Table 2 details the number of FoMs available for each sub-hierarchy from the ToM produced by Dis-MP&F.

The FoMs are spanned by 15,504 paraphrase relations and 55 contradiction relations, discovered when using CoT^{rel} . After eliminating paraphrases, the final set of FoMs was reduced to $N_T = 758$

| Misogyny Hierarchy | No. of FoMs |
|------------------------------|-------------|
| Objectification | 155 |
| Stereotyping | 148 |
| Violence | 55 |
| Shaming | 30 |
| Patriarchal Attitudes | 150 |
| Disrespect towards women | 27 |
| Trivializing serious issues | 70 |
| Discrimination of Women | 82 |
| Dehumanization of Women | 33 |
| Exploitation of Women by Men | 6 |
| Pseudoscience | 2 |

Table 2: Distribution of FoMs across the Taxonomy of Misogyny.

unique FoMs. Therefore the ToM contains 99 MPs across 11 hierarchies as well as 758 FoMs, each connected to some MP from the ToM. Table 2 shows the number of FoMs which are linked to MPs from each hierarchy of the ToM, indicating that the most numerous FoMs address MPs from the hierarchies of *Objectification* and *Patriarchal Attitudes*.

Qualitative results We conducted both an intrinsic and an extrinsic evaluation of the ToM. The intrinsic evaluation concerned the measurement of the quality of the discovered FoMs, of the MPs as well as of the ToM. For this reason we first measured the quality of FoMs encoded in the ToM, which were discovered in Phase 1 of Dis-MP&F, using the metrics introduced in (Weinzierl and Harabagiu, 2024), namely (a) the *soundness* of the rationale provided by the LMM when prompted with CoT^{FoM} ; and (b) the *clarity* of the FoM articulation itself. Three female graduate students were responsible for evaluating both the soundness and clarity of the articulated FoMs, judging N_S as the number of FoMs found to have sound rationales, and N_C indicating the number of FoMs having clear articulations. The quality of reasoning (Z) in identifying FoMs is expressed as $Z = N_S/N_T$, while the quality of articulation (A) is calculated as $A = N_C/N_T$. We obtained $Z = 0.91$ and $A = 0.92$. We used majority voting to arrive at a final decision ensuring consistency in the evaluation process.

The quality of MPs is based on how well we uncovered analogous MPs and later paraphrased MPs performed in Phase 2 of Dis-MP&F. This was evaluated in terms of F1 score. Of the 176 valid MPs generated after Step 2 in Phase 2, 77 MPs were actually analogous. The LLM correctly identified 76 MPs as analogous (True positives) and incorrectly identified 2 MPs as analogous (False positives). Subsequently, there were 97 true neg-

| System | Taxonomy | Task | F1 score |
|-------------------------------|-----------|-------------|--------------|
| Dis-MP&F (Ours) | - | Binary | 0.885 |
| SRCB ((Zhang and Wang, 2022)) | - | Binary | 0.834 |
| Dis-MP&F (Ours) | 4 labels | Multi-label | 0.745 |
| SRCB ((Zhang and Wang, 2022)) | 4 labels | Multi-label | 0.731 |
| Dis-MP&F (Ours) | 11 labels | Multi-label | 0.812 |

Table 3: Results obtained for the Semeval-2022 Task 5 and its extension.

atives and 1 false negative, yielding a precision of 0.974 and a recall of 0.987. Consequently, the resulting F1 score was computed to be 0.980.

The ToM, generated in Phase 2 of Dis-MP&F, was evaluated in two ways, first by considering the quality of the nodes from the taxonomy and then by considering the quality of the hierarchical relations. The quality of the nodes in ToM was measured by *Sibling Distinctiveness* (SD), a metric introduced in (Huang et al., 2020) which we slightly modified. A larger SD means the sibling MPs, sharing a common parent, are truly separate from each other. Since each MP_i from the ToM is associated with a set of FoMs, namely S_{FoM}^i , SD quantifies how distinct these sets of FoMs are from each other. For this reason we first computed C_i , the centroid of each S_{FoM}^i , informed by the embedding representation of FoMs, available through Sentence-BERT (SBERT) (Reimers and Gurevych, 2019). The distance between a pair (S_{FoM}^i, S_{FoM}^j) is computed by the cosine similarity between their centroids, i.e. $dist(S_{FoM}^i, S_{FoM}^j) = C_i \cdot C_j / (\|C_i\| \times \|C_j\|)$. Given S_{FoM}^i , $D_{max}^i = \max_k \{dist(S_{FoM}^i, S_{FoM}^k)\}$, the largest distance S_{FoM}^i has to any other S_{FoM}^k . Then $SD(MP_i) = 1 - D_{max}^i$. We obtained the average SD across all nodes of ToM $SD_{ave} = 0.823$, indicating that in general the MPs are pretty separated from each other.

The second evaluation metric for ToM was the *Rand Index* (RI). The computation of RI required a *Reference Taxonomy of Misogyny Problems* of s (RTMP), generated by a research linguist from the list L_{MP}^C consolidated in the same Phase 2 of Dis-MP&F. RI evaluates the quality of the MHP by using comparisons of the relationships between MPs in both the RHMP and ToM. Let n_{11} represent the number of MP pairs that are siblings in both hierarchies, and n_{00} represent the number of MP pairs that have no common parent in either hierarchy. Similarly, let n_{10} denote the number of MP pairs that are siblings in the ToM but not in RTMP, and n_{01} denote the number of MP pairs that are siblings in the RTMP but not in the ToM. Then $RI = (n_{11} + n_{00}) / (n_{11} + n_{10} + n_{01} + n_{00})$. We

obtained $RI = \mathbf{0.9625}$, indicating great fidelity of the ToM to the RTMP.

The evaluation of the relations between FoMs identified by GPT-4 in phase C, conducted by two research linguists, determined that 92.48% of these relations were accurately identified. Specifically, 98.15% of paraphrase relations were judged to be correct, while 80.30% of contradiction relations were deemed accurate. The annotators exhibited substantial agreement, as reflected by a Cohen’s Kappa (Cohen, 1960) score of 0.636.

The extrinsic evaluation was performed by considering the operation of the Dis-MP&F method on the misogyny detection tasks performed in the Semeval-2022 Task 5 (Fersini et al., 2022). SemEval Task 5 comprises two subtasks: Subtask A is a binary classification task where input memes are categorized as misogynous or non-misogynous, with the systems being evaluated using the Macro-F1 score. Subtask B is a multi-label classification task, where the four original MPs annotated on the MAMI dataset were considered as possible labels. In addition to these subtasks, we extended our evaluation to consider in a multi-label classification using all the hierarchies of the ToM. The results obtained for all these tasks are presented in the Table 3. The Dis-MP&F method outperforms the best performing system from the Semeval-2022 Task 5 on both Task A and B. But more importantly, when we considered the ToM, with all its 11 hierarchies, state-of-the-art misogyny recognition results were obtained for the MAMI dataset, indicating that the more MPs are known, the better they can be identified.

6 Discussion

The automatically derived ToM is providing unprecedented detailed knowledge about how misogyny is discussed on social media, by organizing 99 different MPs, backed by academic literature, into 11 different hierarchies and showcasing in how many ways each MP is framed. In addition to the 99 MPs, we can inspect 758 FoMs. Interestingly, it seems that the hierarchy of *Patriarchal Attitudes*

dominates in terms of the number of MPs it covers, as shown in Figure 6. Table 2 suggests that the hierarchy of *Objectification* dominates in terms of the number of FoMs that address MPs from its hierarchy, immediately followed by the hierarchies of *Patriarchal Attitudes* and *Stereotyping*, providing new insights into the forms of misogyny that dominate on social media.

The Dis-MP&F method was proven to discover misogyny with SOTA results, without needing more than a very few examples in the five different forms of CoT prompting of LMMs/LLMs. This may be due to the impressive knowledge available in GPT-4o, as well as to its reasoning capabilities.

The design and evaluation of the the Dis-MP&F addresses a new form of bias, namely *point-of-view* bias, which was afforded by the specifics of misogyny, in which women are victimized. Interestingly, this bias was manifested through (a) incorrect selection of the Misogyny Problem (MP) by the LLM; and (b) the biased rationale produced through CoT reasoning.

7 Conclusion

This paper presents an automatic method for generating a Taxonomy of Misogyny that encodes not only Misogyny Problems (MPs) but also the way misogyny is framed on social media. Several ways of employing Chain-of-Thought prompting of Large Multimodal Models proved to be successful in creating the taxonomy as well as in automatically discovering misogyny on a benchmark dataset with state-of-the-art results.

8 Ethical Statement

We took extensive measures to respect the privacy and confidentiality of users whose posts were included in SemEval-2022 Task 5 dataset. Given the sensitive nature of this data, we implemented stringent ethical protocols to ensure the research was conducted responsibly. We received approval from the Institutional Review Board at the University of Texas at Dallas for working with this social media dataset. To ensure high-quality evaluation, we employed rigorous standards throughout the annotation process, including the use of inter-annotation agreement metrics to verify consistency. Clear and detailed instructions were provided to the annotators, emphasizing the importance of accurate identification of misogynistic content. All experimental settings, configurations, and procedures

were clearly laid out in this work, the supplemental material, and the linked GitHub repository. We acknowledge the potential harm associated with handling such offensive content, but our work is aimed at advancing understanding of online misogyny and fostering tools to mitigate its impact. This research ultimately serves the public good by contributing to both natural language processing and social justice efforts.

9 Limitations

The Dis-MP&F method introduced in this work processes memes collected from diverse social media platforms, including X/Twitter, Reddit, etc. In this context, memes refer to text superimposed on images, typically intended to be humorous. However, many social media platforms (SMPs) incorporate multiple images, GIFs, videos, and other multimedia content, which our current approach does not yet accommodate. In future work, we aim to extend our method to support more modalities.

A significant limitation of our approach arises from the reliance on Large Language Models (LLMs) or Large Multimodal Models (LMMs), which require a deep infusion of cultural and moral knowledge. While certain reasoning capabilities appear to emerge in LLMs at scale, the extent to which they genuinely exhibit reasoning remains debated (Huang and Chang, 2023). This uncertainty underscores the need to verify the misogyny problems (MPs) generated by LMMs, posing a key limitation in our approach. Furthermore, a reference taxonomy of misogyny problems is essential to evaluate the quality of the taxonomy produced by our method.

Finally, our approach articulates FoMs and generates a taxonomy of misogyny (ToM) based solely on SMPs from the MAMI dataset. In future research, we plan to incorporate additional datasets related to misogyny to explore whether new MPs, overlooked in this study, can be uncovered. Additionally, we aim to expand our methodology to develop taxonomies of problems addressed across a broader range of topics, such as policy issues related to immigration, tobacco use, or same-sex marriage and to discover and articulate the frames that address them. Furthermore, we aim to discover frames evoked in the Media Frames Corpus (Card et al., 2015), for which the problems are known, and eventually explore the generation of additional policy problems and their ontological organization.

References

- Maja Brandt Andreassen. 2020. ‘rapeable’ and ‘unrapeable’ women: the portrayal of sexual violence in internet memes about #metoo. *Journal of Gender Studies*, 30:102 – 113.
- Mary E. Anzovino, Elisabetta Fersini, and Paolo Rosso. 2018. Automatic identification and classification of misogynistic language on twitter. In *International Conference on Applications of Natural Language to Data Bases*.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Annual Meeting of the Association for Computational Linguistics*.
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, 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.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37 – 46.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alice H. Eagly and Antonio Mladinic. 1989. Gender stereotypes and attitudes toward women and men. *Personality and Social Psychology Bulletin*, 15:543 – 558.
- Robert M. Entman. 1993a. Framing: Toward clarification of a fractured paradigm. *Journal of Communications*.
- Robert M. Entman. 1993b. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4):51–58.
- Elisabetta Fersini, Francesca Gasparini, Giulia Rizzi, Aurora Saibene, Berta Chulvi, Paolo Rosso, Alyssa Lees, and Jeffrey Sorensen. 2022. SemEval-2022 task 5: Multimedia automatic misogyny identification. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 533–549, Seattle, United States. Association for Computational Linguistics.
- Joseph L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76:378–382.
- Ella Guest, Bertie Vidgen, Alexandros Mittos, Nishanth Sastry, Gareth Tyson, and Helen Margetts. 2021. An expert annotated dataset for the detection of online misogyny. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1336–1350. Association for Computational Linguistics.
- Jiaxin Huang, Yiqing Xie, Yu Meng, Yunyi Zhang, and Jiawei Han. 2020. Corel: Seed-guided topical taxonomy construction by concept learning and relation transferring. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD ’20*, page 1928–1936, New York, NY, USA. Association for Computing Machinery.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065, Toronto, Canada. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Preprint*, arXiv:2304.08485.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Michael Kinney, and Daniel S. Weld. 2020. S2orc: The semantic scholar open research corpus. In *Annual Meeting of the Association for Computational Linguistics*.
- OpenAI. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.
- Pulkit Parikh, Harika Abburi, Pinkesh Badjatiya, Radhika Krishnan, Niyati Chhaya, Manish Gupta, and Vasudeva Varma. 2019. Multi-label categorization of accounts of sexism using a neural framework. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1642–1652, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *Preprint*, arXiv:1908.10084.
- Dawn M. Szymanski, Lauren Moffitt, and Erika R. Carr. 2011. Sexual objectification of women: Advances to theory and research 1ψ7. *The Counseling Psychologist*, 39:38 – 6.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and

Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models. *Preprint*, arXiv:2201.11903.

Maxwell Weinzierl and Sanda Harabagiu. 2024. Discovering and articulating frames of communication from social media using chain-of-thought reasoning. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1617–1631, St. Julian’s, Malta. Association for Computational Linguistics.

Philine Zeinert, Nanna Inie, and Leon Derczynski. 2021. Annotating online misogyny. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3181–3197, Online. Association for Computational Linguistics.

Jing Zhang and Yujin Wang. 2022. Srcl at semeval-2022 task 5: Pretraining based image to text late sequential fusion system for multimodal misogynous meme identification. In *International Workshop on Semantic Evaluation*.

Qian Zhang, Qinxuan Chen, and Xinrong Gu. 2024. Trolling, cyberstalking, body-shaming, slut-shaming a study on online abuses of social media. *Communications in Humanities Research*, pages 232–241.

A Prompting details for Phase 1 of the Disc-MP&F method

The system and user prompts used to articulate Frames of Misogyny (FoMs), part of the CoT^{FoM} prompting method used in Phase 1 of the Disc-MP&F method are shown in Figure 8a and 8b.

The system prompt determines the style, formality, or tone of the responses, whereas the user prompt drives the actual conversation while prompting a LMM. In the system prompt, we give the definition of frames of communication from (Entman, 1993b) and a brief definition of misogyny.

In the user prompt, we ask the LMM to articulate FoMs from the Social Media Postings that are presented to it, containing memes.

B Distribution of Annotations of Misogyny Problems in the MAMI dataset.

The misogyny benchmark dataset was released for the SemEval-2022168 Task 5: Multimedia Automatic Misogyny Identification (MAMI). In this dataset, Social Media Postings (SMPs) may be annotated with one, two, three, four Misogyny Problems (MPs) or with none of them. The set of SMPs

Chain-of-thought prompting for articulating frames from memes

System prompt

You are an expert linguistic assistant. Frames of communication select particular aspects of an issue and make them salient in communicating a message. They are ubiquitous in social media discourse and can impact how people understand issues and, more importantly, how they form their opinions. It has been argued that "to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote problem definition, causal interpretation, moral evaluation, and/or treatment recommendation. You will be tasked with identifying and articulating misogyny framings on memes. A meme in this context is a combination of image and text, typically meant to be humorous in nature, that is copied and spread rapidly by internet users, often with slight variations. Misogyny is defined as dislike of, contempt for, or ingrained prejudice against women.

(a)

User prompt

You will be tasked with identifying and articulating misogyny framings on the following memes. You should discuss your reasoning first, and then provide a final decision. Each image provided may or may not contain one or more framings, so your first step is (a) Reason about whether the image contains a framing (or more framings), or just states something factual or an experience. If the image contains a framing, the next step is (b) Articulate that framing succinctly. You will perform these steps until the answer to (a) is false, either because there are no framings in the image, or because you have already expressed all the framings.

(b)

Figure 8: Details of the CoT^{FoM} prompt.

that received annotations pertaining to only one MP are considered to be a set SMP_1 . Those that have received annotations pertaining to two different MPs are considered to be a set SMP_2 , whereas those that received annotations pertaining to three different MPs are considered to be a set SMP_3 . The SMPs annotated with all four MPs considered in MAMI are gathered in a set SMP_4 .

Figure 9(a) shows that in the set SMP_1 , those SMPs with labels corresponding to the MP of *Stereotyping*, *Objectification* and *Shaming* were predominant. Figure 9(b), shows the distribution of SMPs in SMP_2 . It can be noted that the SMPs that were annotated both with *Objectification* and *Stereotyping*, as well as those that were annotated both with *Stereotyping* and *Violence* are predominant in the set SMP_3 .

Table 4 lists the number of SMPs in set SMP_1 , SMP_2 , SMP_3 and SMP_4 , showing that it is rea-

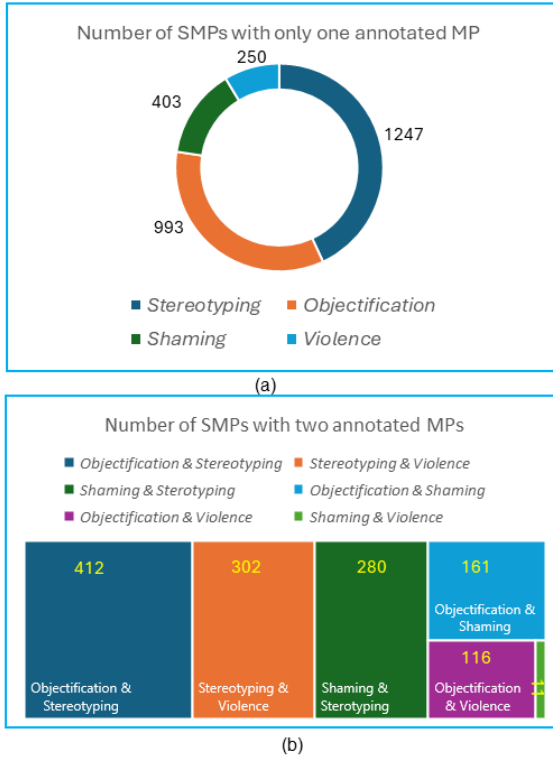


Figure 9: (a) Distribution of MP labels used in SMPs that received only one MP annotation; (b) Distribution of MP labels used in SMPs that received two MP annotations.

sonable to select examples for the CoT^{FoM} prompt only from SMP_1 and SMP_2 .

| Set | No. of SMPs |
|---------|-------------|
| SMP_1 | 2893 |
| SMP_2 | 1288 |
| SMP_3 | 530 |
| SMP_4 | 45 |

Table 4: Distribution of MP labels in SMPs.

C Prompting details for Phase 2 of the Disc-MP&F method

Phase 2 of the Disc-MP&F method uses several prompting methods, namely the CoT^{C_MP} prompt used for correcting non-valid MPs (Step 2); the CoT^A prompt used for identifying analogous MPs (Step 3); and the CoT^H prompt for detecting sub-concept relations among MPs (Step 4). All these prompt use the same System Prompt, which is illustrated in Figure 10, while each of them has a different user prompt, also shown in Figure 10.

A consistent system prompt, encompassing the definition of frames of communication and misogyny, is utilized across all steps. In the CoT^{C_MP}

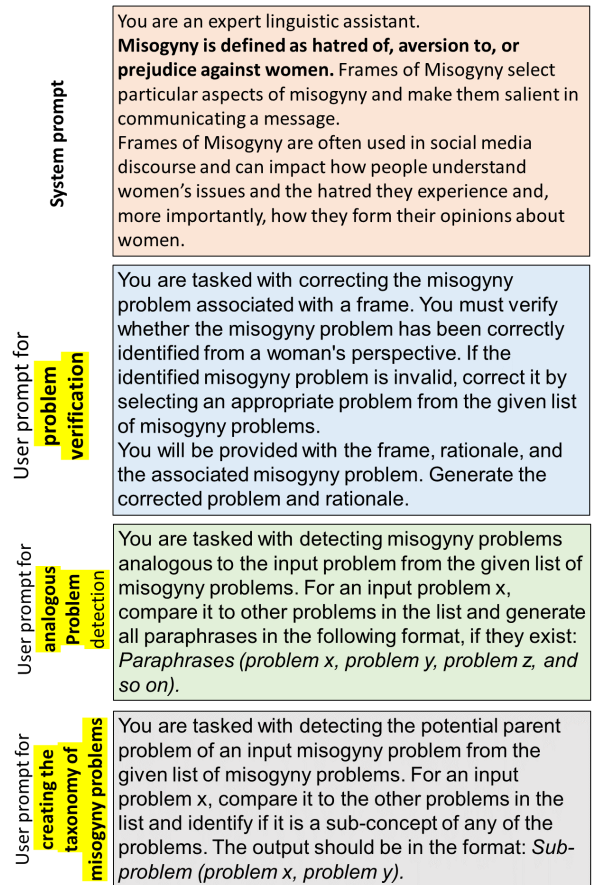


Figure 10: Three different Chain-of-Thought (CoT) prompting methods used in Phase 2 of the Disc-MP&F method.

prompt, the user prompt includes the frame, the corresponding rationale, and the MP extracted from the rationale, supplemented by a list of valid misogyny problems derived in Step 1 of the Phase 2 of the Disc-MP&F method. The Large Language Model (LLM) is tasked with outputting the corrected problem along with the rationale.

In the CoT^A prompt, the user prompt considers as input (1) an MP_x and (2) a list of valid MPs. The LLM generates a list of MPs that are analogous to MP_x .

In CoT^H prompt the list of valid MPs is again supplied to the user prompt alongside some MP_x . The LLM is asked to identify the potential parent MP_y of MP_x .

D Annotator Instructions for Evaluating Misogyny Problems

Along with the instructions outlined in Figure 12, we provide a carefully crafted example with the desired annotations to facilitate a clear understanding for the annotators.

| | |
|----------------------|---|
| System prompt | <p>You are an expert linguistic assistant.</p> <p>Misogyny is defined as hatred of, aversion to, or prejudice against women. Frames of Misogyny select particular aspects of misogyny and make them salient in communicating a message.</p> <p>Frames of Misogyny are often used in social media discourse and can impact how people understand women's issues and the hatred they experience and, more importantly, how they form their opinions about women.</p> |
| User prompt | <p>You will be tasked with identifying relationships between the given misogyny frames. You should discuss your reasoning first, and then provide a final decision. Each framing provided may or may not be involved in a single relationship with one framing from a provided set of similar framings. We will consider two possible relationships:</p> <ol style="list-style-type: none"> 1. Paraphrases(X,Y): X and Y say essentially the same exact thing, with different words or phrasing. If one person agreed with X, they would agree with Y, and vice versa. Frames should share the same cause and the same problem to be considered paraphrases. 2. Contradicts(X,Y): X and Y contradict each other, such that they frame the same exact issue from opposing perspectives. If one person agreed with X, they would disagree with Y, and vice versa. The two frames X and Y should essentially paraphrase each other, sharing the same problem and cause but from opposing perspectives. 3. No relationship: There are no relationships between the new framing and any of the provided framings. Your first step for each framing is <ol style="list-style-type: none"> (a) Reason about if the framing holds one of the above relationships with any of the provided framings. (b) Multiple relationships could be true, but prioritize in the order provided: If a paraphrase relationship holds, it must be provided. If there is no paraphrase, then look for contradicts. Finally, if there is no contradicts relationship, answer no relationship. If a relationship is identified, then state that relationship, using the IDs for each framing. |

Figure 11: CoT^{rel} prompt.

E Prompting details for Phase 3 of the Disc-MP&F method

The prompt employed for discovering paraphrases and contradictions in phase 3 of our methodology is shown in Figure 11. In the user prompt, the LLM is instructed to detect relationships between a given FoM and a set of similar FoMs. We also provide the LLM with definitions of paraphrases and contradictions in the context of FoMs. Additionally, the LLM is prompted to generate rationales alongside the detected relationships, offering justifications for each.

F Examples of Misogyny Problems and Frames of Misogyny

Examples of SMPs corresponding to the misogyny problem (MP) 'patriarchal attitudes' and their associated Frames of Misogyny (FoMs) are shown in Figure 13. Figure 13(a) illustrates an SMP aligned

Objective

You are tasked with evaluating whether the identified misogyny problem associated with a frame is valid. If the problem is deemed incorrect, you must replace it with a more appropriate problem from the provided list and provide a rationale explaining your decision.

Definitions

1. **Meme:** A combination of image and text, typically meant to be humorous in nature, shared widely on the internet. It often reflects societal attitudes and biases.
2. **Frame:** A communication strategy that highlights certain aspects of an issue to promote specific interpretations, moral evaluations, or solutions.
3. **Misogyny Problem:** A specific issue or expression reflecting dislike, contempt, or prejudice against women.

Provided Materials

For each item, you will be given:

1. **The Meme:** Image and text forming the meme.
2. **Generated Frame:** The identified misogyny framing of the meme.
3. **Rationale:** Explanation connecting the meme to the associated misogyny problem.
4. **List of Valid Misogyny Problems:** A predefined list of misogyny problems for reference.

Annotation Tasks

For each item:

1. **Validate the Identified Problem:**
 - Assess whether the misogyny problem mentioned in the rationale is appropriate for the given meme and frame.
 - Answer **Yes** if the problem is valid or **No** if it is not.
2. **Replace the Problem (if needed):**
 - If you answer **No**, select a more appropriate misogyny problem from the provided list.
3. **Provide a Rationale:**
 - Explain why you think the original problem is incorrect and justify your chosen replacement. Your rationale should be clear, concise, and evidence-based.

Guidelines

1. **Be Objective:** Focus solely on the content of the meme, the generated frame, and the rationale.
2. **Use the Valid Problem List:** Only select replacement problems from the provided list.
3. **Rationale is Key:** Provide a clear explanation for your choices to ensure consistency and clarity.
4. **Respect Sensitivities:** This task involves analyzing potentially offensive content. Please approach the task professionally.
5. In case of uncertainty, prioritize the closest match from the valid problem list.

Figure 12: Annotator Instructions for Evaluating Misogyny Problems

with the MP 'patriarchal attitudes,' reinforcing traditional gender roles by suggesting that women should remain silent.



Figure 13: Example of FoMs for MPs that are sub-concepts of *patriarchal attitudes*.

In Figure 13(b), the SMP evokes the MP 'conditional respect,' implying that women are only deserving of respect if they comply with male authority. The SMP in Figure 13(c) normalizes infidelity by asserting that having extramarital relationships with women is socially acceptable. Lastly, the reductionist perspective in Figure 13(d) invokes the MP 'gender essentialism,' as it reduces women to mere biological functions.

Examples of SMPs categorized under the MP 'discrimination of women' along with the associated FoMs are depicted in Figure 14. The MP 'discrimination of women' and its sub-concepts demonstrate the kind of discriminatory practices that coexist with misogyny on social media platforms. A total of 117 SMPs were identified as corresponding to this MP.

In Figure 14(a) the SMP reflects the MP 'ableism' as the text portrays women with disabilities as inadequate. Figure 14(b) corresponds to the MP 'classism', where the SMP normalizes domestic violence, implying that it is an accepted part of life for lower socio-economic demographics. Figure 14(c) exemplifies the MP 'racism', by perpet-

uating stereotypes based on both race and gender. Finally, in Figure 14(d), the MP 'homophobia' is invoked, as the SMP deems same-sex attraction as preposterous, further reinforcing homophobic biases.



Figure 14: Example of FoMs for MPs that are sub-concepts of *discrimination of women*.

Additional examples of FoMs associated with the MPs 'stereotyping' and 'patriarchal attitudes' are shown in Table 5 and 6 respectively.

Table 5 illustrates ten different ways of framing the same Misogyny Problem (MP) of 'stereotyping'. This was one of the MPs originally annotated in the MAMI dataset. Each FoM from Table 5 provides a different causal interpretation of the MP. Moreover it allows us to better understand why this is a misogyny problem. Importantly, it greatly helps us to identify automatically women stereotyping in text, when we are able to uncover these FoMs.

Table 6 illustrates ten different FoMs that are associated with ten different MPs organized in the hierarchy of 'patriarchal attitudes'. None of these MPs were originally annotated in the MAMI dataset, although, as seen in Table 6, they represent Misogyny Problems (MP) or Misogyny Sub-Problems and, more importantly, they represent different forms of misogyny than 'shaming' stereo-

| Index | FoM |
|-------|---|
| 1. | Men are consumers of pornography, while women are producers of it. |
| 2. | Women are more pleasant or rational when they are no longer in a relationship with a man. |
| 3. | A woman's fashion choices are dictated by her relationship and parental status. |
| 4. | Women who put effort into their appearance are shallow and have no right to complain about men. |
| 5. | Women are irrational and use zodiac signs to justify their behavior. |
| 6. | Wives are two-faced and act differently with their husbands compared to others. |
| 7. | Women, especially religious ones, can be easily manipulated into forgiving infidelity. |
| 8. | Women care more about animals than human life. |
| 9. | All women keep a list of potential partners in reserve. |
| 10. | Women are only suited for running the kitchen. |

Table 5: Examples of Frames of Misogyny associated with the Misogyny Problem of 'stereotyping'.

| Index | FoM | MP |
|-------|--|--------------------------|
| 1. | Women who claim to be independent are actually dependent on their partners for money. | Patriarchal attitudes |
| 2. | Women should be exclusive to one man and should not engage with other men. | Possessiveness |
| 3. | Women are not intelligent and can only use their looks to get what they want. | Intellectual degradation |
| 4. | Drug use is acceptable or impressive for men but not for women. | Double standards |
| 5. | Women's bodies are meant to be sexualized and critiqued based on how much skin they show. | Policing women's bodies |
| 6. | Men who are attracted to feminists are weak and inferior. | Toxic masculinity |
| 7. | A woman's value is based on her virginity and lack of tattoos. | Gender essentialism |
| 8. | A woman's self-worth is tied to receiving attention from men, even if that attention is inappropriate. | Male validation |
| 9. | Women should have no say in their marriage decisions. | Forced marriage |
| 10. | Men need to assert dominance and win arguments in relationships. | Patriarchal control |

Table 6: Examples of FoMs and associated sub-concept MPs under 'patriarchal attitudes'.

typing', 'objectification' or 'violence', annotated in the MAMI dataset.

For each MP shown in Table 6, we also illustrate one of the Frames of Misogyny that address it. The Taxonomy of Misogyny (ToM) presented in this paper, contains not only 11 hierarchies of MPs, but for each MP from the ToM, there are many FoMs, as we have seen in Table 5. Therefore, the ToM, organized in hierarchies of MPs, also links each MP to all the FoMs that explain it.

Table 6 shows a small number of FoMs that explain the different forms of misogyny characterizing the patriarchal attitudes that support misogyny. In addition, the ToM provides information of how many SPMs are evoked by the same FoM, linking the ToM to the MAMI benchmark dataset. We also note that sometimes, one Social Media Posting (SMP) may evoke multiple FoMs.

G The Taxonomy of Misogyny

The Taxonomy of Misogyny (ToM) consists of 11 different hierarchies. We illustrate all these hierarchies of Misogyny Problems (MPs) along with their definitions. The levels of each taxonomy is indexed

in the following way: in hierarchy i , a Misogyny Problem (MP) received the index i , which is a number, if it is the root of the hierarchy. Two MPs received the index $i.j$ if they are in the hierarchy i and both of them at the level j . We also list the number of Frames of Misogyny (FoMs) linked to each MP from the ToM.

| Misogyny Problem (No. of FoMs addressing it) | Definition Of MP | Index of MP |
|--|--|-------------|
| 1. Objectification (117) | <i>The action of degrading women to the status of a mere object</i> | 1 |
| 1.1 Sexualization (14) | <i>The act of viewing someone or something in a sexual way, or of emphasizing their sexual characteristics</i> | 2 |
| 1.1.1 Sexual entitlement (4) | <i>Belief that men's own sexual needs are more important than the feelings and rights of women</i> | 3 |
| 1.1.2 Sexual innuendo (5) | <i>Subtly or indirectly suggesting sexual content, ideas, or connotations through language, gestures, or humor</i> | 4 |
| 1.1.3 Age-related sexualization (4) | <i>Inappropriate or harmful imposing of sexual attributes, behaviors, or expectations onto women based on their age</i> | 5 |
| 1.2 Unrealistic beauty standards (11) | <i>Oppressive and often unattainable ideals of physical appearance imposed on women by patriarchal society</i> | 6 |
| 2. Stereotyping (96) | <i>Application of broad, oversimplified, and often negative generalizations about women based on their gender</i> | 7 |
| 2.1 Misrepresentation (7) | <i>The distortion or incorrect portrayal of women and their experiences, often driven by sexist biases and stereotypes</i> | 8 |
| 2.1.1 False equivalence (4) | <i>The misleading comparison of women's experiences, struggles, or behavior to those of men in ways that ignore or downplay the gender-based power imbalances and discrimination women face.</i> | 9 |
| 2.1.2 Deflection (6) | <i>Shifting responsibility or accountability for one's own actions, mistakes, or harmful behavior onto women</i> | 10 |
| 2.1.3 Misunderstanding feminism (5) | <i>The distortion or misinterpretation of feminist principles, often in ways that vilify or dismiss the movement's goals.</i> | 11 |
| 2.2 Mistrust in women (5) | <i>skepticism or doubt directed toward women based on generalized, often negative assumptions and stereotypes about their gender</i> | 12 |
| 2.2.1 Commitment phobia (1) | <i>Fear or avoidance of long-term relationships/ lasting commitments, particularly in romantic or emotional contexts involving women</i> | 13 |
| 2.3 Biased Judgement (10) | <i>Biased evaluation or criticism of women based on sexist attitudes and stereotypes</i> | 14 |
| 2.3.1 False accusations (1) | <i>baseless or exaggerated claims made against women, often to undermine their credibility, harm their reputation, or perpetuate gender-based discrimination</i> | 15 |
| 2.3.2 Scapegoating (2) | <i>Unfairly blaming women for problems, failures, or negative events, even though they are not responsible</i> | 16 |
| 2.3.3 Divisiveness (8) | <i>Deliberate or unintended creation of conflict, hostility, or separation between different groups, particularly by blaming women for societal issues or personal failures</i> | 17 |
| 2.4 / 7.5 Reductionism (3) | <i>Oversimplification of women's identities, experiences, or roles by reducing them to narrow or stereotypical aspects, such as their appearance, reproductive function, or societal roles as mothers or caregivers</i> | 18 |
| 3. Violence (26) | <i>Physical, emotional, psychological, sexual, or economic harm directed at women or girls, driven by gender-based hatred, control, or discrimination</i> | 19 |
| 3.1 Sexual harassment (2) | <i>Unwelcome, inappropriate, or coercive behavior of a sexual nature that creates a hostile or intimidating environment</i> | 20 |

| | | |
|--|---|----|
| 3.1.1 Promoting rape culture (8) | <i>Perpetuation of attitudes, norms, and practices that normalize, trivialize, or excuse sexual violence and harassment</i> | 21 |
| 3.1.2 Trivializing consent (2) | <i>Downplaying or minimizing the importance of obtaining and respecting consent in sexual or interpersonal situations</i> | 22 |
| 3.1.3 Victim blaming (12) | <i>Practice of attributing responsibility or fault to women for the violence, harassment, or discrimination they experience</i> | 23 |
| 3.1.4 Endorsing marital rape (1) | <i>Acceptance, justification, or normalization of non-consensual sexual activity within marriage</i> | 24 |
| 3.1.5 Endorsing necrophilia (2) | <i>Perverse and harmful sexual behavior directed towards deceased females</i> | 25 |
| 3.2. Promoting self-harm (1) | <i>Actions, messages, or behaviors that encourage or pressure women to inflict harm upon themselves, either physically or psychologically</i> | 26 |
| 3.3 Professional misconduct (1) | <i>Improper, unethical, or unlawful behavior by a professional in the course of their work, violating the standards, duties, and ethical guidelines set by their profession</i> | 27 |
| 4. Shaming (7) | <i>Publicly or privately humiliating, degrading, or criticizing women based on their behavior, appearance, or choices</i> | 28 |
| 4.1. / 7.2 Derogatory labeling (9) | <i>Use of negative, demeaning, or contemptuous terms to describe or refer to women</i> | 29 |
| 4.2 / 7.3 Ridicule (10) | <i>Mocking, belittling, or making fun of women in a way that undermines their dignity, worth, or autonomy</i> | 30 |
| 4.3 Fat shaming (4) | <i>Criticizing, mocking, or demeaning women based on their body size or weight.</i> | 31 |
| 5. Trivializing serious issues (33) | <i>Downplaying or minimizing the importance, impact, or severity of problems that warrant attention and concern</i> | 32 |
| 5.1 Trivializing women's issues (7) | <i>Downplaying, dismissing, or minimizing the significance of problems and challenges that predominantly affect women</i> | 33 |
| 5.1.1 / 3.3 Trivializing sexual assault (8) | <i>Downplaying, dismissing, or minimizing the seriousness and impact of sexual violence.</i> | 34 |
| 5.1.2 Minimizing feminist efforts (7) | <i>Dismissing, mocking, or minimizing the importance and goals of the feminist movement, which seeks to achieve gender equality and address systemic issues</i> | 35 |
| 5.1.3 Trivializing women's sexual satisfaction (1) | <i>Dismissing or downplaying a woman's sexual pleasure, needs, or experiences.</i> | 36 |
| 5.2 Trivializing mental health issues (8) | <i>Downplaying, dismissing, or minimizing the importance of mental health issues and concerns</i> | 37 |
| 5.3 Trivializing oppression (5) | <i>Downplaying, dismissing, or minimizing the significance and impact of systemic discrimination, exploitation, and inequality that certain groups face based on factors like race, gender, class, sexuality, or religion</i> | 38 |
| 5.4 Trivializing infidelity (2) | <i>Downplaying, dismissing, or minimizing the significance and impact of a partner's betrayal</i> | 39 |
| 5.5 Trivializing prostitution (1) | <i>Downplaying or minimizing the significance and complexities associated with prostitution, which is the act of exchanging sexual services for money or other goods.</i> | 40 |
| 5.6 Trivializing eating disorders (1) | <i>Downplaying, dismissing, or minimizing the significance and seriousness of eating disorders which are serious mental health conditions characterized by persistent and unhealthy patterns of eating behaviors and attitudes toward food</i> | 41 |
| 5.7 Trivializing addiction (1) | <i>Downplaying, dismissing, or minimizing the seriousness and</i> | 42 |

| | | |
|---|---|----|
| | <i>complexity of addiction, which is a chronic, often relapsing disorder characterized by compulsive substance use or behaviors despite adverse consequences</i> | |
| 5.8 Trivializing the need for representation (1) | Downplaying the importance of including under-represented people in an industry | 43 |
| 6. Patriarchal attitudes (27) | A system of society or government in which men hold the power and women are largely excluded from it. | 44 |
| 6.1 Dismissing women's rights (2) | Disregarding, minimizing, or invalidating the importance and validity of women's rights and issues related to gender equality | 45 |
| 6.1.1 Undermining women's rights movements (3) | Actions or attitudes that diminish, obstruct, or discredit the efforts and goals of movements advocating for gender equality and women's rights. | 46 |
| 6.2 Intellectual degradation (4) | Systematic belittling, undermining, or dismissal of women's intellectual capabilities and contributions based on sexist beliefs | 47 |
| 6.3 Women Subjugation (3) | Bringing women under domination or control, often through oppression or coercion | 48 |
| 6.4 Patriarchal control (3) | Systemic and institutional mechanisms through which patriarchal societies enforce male dominance and maintain gender inequalities | 49 |
| 6.4.1 Possessiveness (4) | Strong desire to control, dominate, or claim ownership over women | 50 |
| 6.4.2 Ownership (2) | Control and entitlement exercised by patriarchal systems over women as though they are possessions or property rather than autonomous beings | 51 |
| 6.3.3 Gatekeeping (5) | Controlling, and usually limiting, general access to women | 52 |
| 6.4.4 Coercion (1) | Use of force, threats, manipulation, or pressure to enforce compliance with patriarchal norms and expectations | 53 |
| 6.4.5 Forced marriage (1) | Marriage in which one or both individuals are coerced into the union without their free and informed consent, often due to patriarchal norms, pressures, or practices | 54 |
| 6.4.6 Policing women's bodies (2) | Practices and mechanisms through which societal, cultural, or institutional forces regulate, control, and scrutinize women's bodies and behaviors | 55 |
| 6.5 Enforced gender norms (40) | Rigid, culturally and socially prescribed roles, behaviors, and expectations assigned to women, which are imposed and maintained by societal institutions, cultural practices, and interpersonal interactions | 56 |
| 6.5.1 Gender essentialism (7) | Belief that gender differences are inherent, natural, and biologically determined, and that these differences define and prescribe specific roles, behaviors, and attributes for men and women | 57 |
| 6.5.2 Double standards (14) | Practice of applying different sets of rules, expectations, or judgments to individuals based on their gender, with men and women often being treated unequally. | 58 |
| 6.6 Conditional respect (3) | Respect or acknowledgment that is granted to individuals based on their adherence to traditional gender roles, behaviors, or expectations set by patriarchal norms | 59 |
| 6.7 Undermining women's capabilities (6) | Dismissing, devaluing, or doubting the skills, talents, and potential of women based on gender stereotypes or biases | 60 |
| 6.8 Demeaning aspirations (4) | Belittling, devaluing, or dismissing the goals and ambitions of women, based on traditional patriarchal norms | 61 |

| | | |
|--|---|----|
| 6.9 Male validation (2) | Seeking approval, recognition, or acceptance from men, often as a measure of worth or success | 62 |
| 6.10 Toxic masculinity (8) | Cultural norms and behaviors associated with traditional male gender roles that promote harmful and restrictive attitudes towards masculinity | 63 |
| 6.11 Promoting infidelity (3) | Accepting unfaithful or disloyal behavior towards a romantic or sexual partner | 64 |
| 6.12 Incest (1) | Sexual activity or relationships between close relatives that are rooted in or perpetuated by patriarchal and sexist beliefs | 65 |
| 6.13/10.14 Wage disparity (5) | Unequal pay that women receive compared to men for performing the same or similar work, which is often rooted in patriarchal systems that prioritize men's labor and devalue women's contributions | 66 |
| 7. Disrespect towards women (6) | Behaviors, attitudes, or actions that devalue, undermine, or disregard women's dignity, autonomy, or worth | 67 |
| 7.1 Cultural insensitivity (5) | Disregard or disrespect for the unique cultural experiences, traditions, and identities of women, especially when those experiences intersect with gender discrimination | 68 |
| 7.4 Stigmatization (4) | Regarding women as worthy of disgrace or great disapproval | 69 |
| 7.4.1/2.3.4 Single motherhood stigmatization (2) | Negative judgments, stereotypes , and societal discrimination faced by women raising children without a male partner | 70 |
| 7.4.2 Menstruation stigma (3) | Negative social and cultural perceptions, beliefs, and practices that surround menstruation | 71 |
| 7.6 Dismissiveness (3) | Having or showing a disdainful attitude toward women regarded as unworthy of serious attention | 72 |
| 7.6.1 Invalidating women's experiences (1) | The denial, dismissal , or undermining of women's personal experiences, feelings, or perceptions | 73 |
| 7.7 Lack of accountability by men (1) | Men's failure to take responsibility for actions or decisions regarding women | 74 |
| 7.8 Disrespecting sex workers (1) | Behaviors, attitudes, or actions that devalue, demean, or undermine individuals who work in the sex industry | 75 |
| 8. Dehumanization of women (26) | Depriving women of human qualities, personality, or dignity | 76 |
| 8.1 Disposability (2) | Perception or treatment of women as easily replaceable or expendable, reducing their value to a narrow, often utilitarian purpose | 77 |
| 8.2 Demonization (3) | Portraying women as inherently evil, dangerous, or morally corrupt | 78 |
| 8.3 Women Valuation (1) | Practice of assigning women a value or status based on perceived traits | 79 |
| 8.4 Fearmongering (1) | Deliberate promotion of fear or anxiety about women's behavior, roles, or societal influence | 80 |
| 9. Exploitation of Women by Men (2) | Unfair and abusive use of women's labor, bodies, or emotions for personal gain, often driven by gender-based power imbalances and societal expectations that devalue or commodify women | 81 |
| 9.1 Transactional relationships (3) | Relationships where interactions between men and women are reduced to exchanges based on material, emotional, or sexual transactions, rather than mutual respect, equality, or emotional connection | 82 |
| 9.2 Pedophilia exploitation (1) | Sexual interest in prepubescent girls or attempts to engage in sexual acts with prepubescent girls | 83 |
| 10. Discrimination of women (4) | Unjust or prejudicial treatment of women | 84 |
| 10.1 Ageism (7) | Prejudice against women based on their age, often rooted in sexist beliefs that devalue women as they grow older | 85 |
| 10.2 Appearance-based discrimination (15) | Unfair treatment or judgment of women based on their physical appearance | 86 |

| | | |
|--------------------------------------|---|----|
| 10.3 Exclusion (10) | <i>Deliberate or systemic marginalization of women from social, professional, political, or cultural opportunities and spaces</i> | 87 |
| 10.4 Reverse sexism (1) | <i>Sexism directed toward men or boys, often as a reaction to the discrimination and inequalities faced by women</i> | 88 |
| 10.5 Workplace discrimination (2) | <i>Unfair or prejudicial treatment of female employees or female job applicants</i> | 89 |
| 10.6 Ableism (9) | <i>Discrimination in favor of able-bodied people</i> | 90 |
| 10.7 Racism (9) | <i>Prejudice or antagonism against a woman on the basis of her membership in a particular racial or ethnic group</i> | 91 |
| 10.8 Constrained gender identity (4) | <i>Limitation or restriction of a woman's ability to express and experience their gender identity freely due to societal norms, expectations, or pressures</i> | 92 |
| 10.9 Intersectional prejudice (1) | <i>Prejudice based on the simultaneous overlap of multiple social identities, like race, gender, class, or sexual orientation, where the discrimination stemming from each identity combines to create a unique and often more severe form of prejudice against a woman</i> | 93 |
| 10.10 Neo sexism (1) | <i>Assertion that genders have already achieved equity, and that discrimination against a gender does not exist</i> | 94 |
| 10.11 Homophobia (3) | <i>Dislike of or prejudice against lesbians</i> | 95 |
| 10.12 Transphobia (6) | <i>Dislike of or strong prejudice against transgender people</i> | 96 |
| 10.13 Classism (7) | <i>Prejudice against or in favor of people belonging to a particular social class.</i> | 97 |
| 11. Sexist pseudoscience (1) | <i>Use of false, misleading, or unscientific claims to justify or perpetuate sexist beliefs and practices</i> | 98 |
| 11.1 Conspiracy thinking (1) | <i>Belief that gender-related inequalities, feminist movements, or efforts to promote women's rights are part of a secret, coordinated plot to undermine men or traditional societal structures</i> | 99 |