

Manospheres: exploring an Italian incel community through the lens of NLP and Frame Semantics

Sara Gemelli^a and Gosse Minnema^b

University of Pavia^a; University of Bergamo^a; University of Groningen^b

sara.gemelli@unibg.it; gosseminnema@gmail.com

Abstract

We introduce a large corpus of comments extracted from an Italian online *incel* ('involuntary celibate') forum, a community of men who build a collective identity and anti-feminist ideology centered around their inability to find a sexual or romantic partner and who frequently use explicitly misogynistic language. Our corpus consists of 2.4K comments that have been manually collected, analyzed and annotated with topic labels, and a further 32K threads (700K comments) that have been automatically scraped and automatically annotated with FrameNet. We show how large-scale frame semantic analysis can shed a light on what is discussed in the community, and introduce incel topic classification as a new NLP task and benchmark.

Keywords: Incels, Typical frames

1. Introduction

Among the many communities that prosper on the Internet, the so-called *manosphere* seems to be one of the most obscure yet active of the last years. More precisely, the manosphere is better described as a *network* of forums, groups, and blogs that are heterogeneous with respect to their approaches to society, but united by the commitment to expose the oppression that men supposedly suffer at the hands of women (Ging, 2019). Within this landscape, the incel community is probably the best-known. The term is a portmanteau for *involuntary celibate*, and it is used by the users to describe themselves; the men who populate incel forums, in fact, argue that their inability to find a partner is due to their appearance or their personality not meeting the standards imposed by the society in general and by women in particular.

The language used within these groups is often explicitly derogatory towards women, who are seen both as active oppressors and object of desire, whereas the users of the forums conceive themselves as passive victims who are doomed to suffer their condition of *inceldom*. Their linguistic repertoire is characterised by what Jane (2014) calls *e-bile*, a term that encapsulates a variety of offensive expressions typical of online communication, often misogynistic and homophobic; it is also rich in neologisms and acronyms, that reflect the users' linguistic creativity and often disguise disparaging expressions. These communities are often conceived as hermetic environments, exceptional and somehow remote, but in fact the very nature of the Web makes it easy for their theories to disperse in the mainstream areas of the online world, with language as the vehicle. Their discursive choices and linguistic creativity contribute to the strengthening of their sense of belonging and

the creation of an in-group (the users of the forum and incels in general, depending on the discussion) and of an out-group (women, and sometimes people that are considered 'not ugly') very rigidly delimited. The misogynistic ideologies perpetuated within these echo chambers are strictly linked to the patriarchal rules that permeate society, and reflect the sense of entitlement (Manne, 2017) that men who live by those rules often feel towards women; these sexist positions become more explicit in the forums through the users' discourse, working as a 'social glue' and justifying their sense of misery and failure, recurring motifs within the community.

We believe that developing computational linguistic analyses of incel discourse is urgently necessary: a better understanding of how the manosphere works is vital for addressing 'modern' forms of misogyny, and computational tools are needed for quantitatively analyzing the enormous amounts of online content that are produced every day. Moreover, incel discourse sometimes leads to extremist violence;¹ this increases the need for real-time automatic monitoring tools.

So far, there has been only a small amount of NLP research about inceldom, all of it very recent and mostly investigating Anglophone contexts (see §2). In this paper, we build on this previous work by analyzing incel language through the lens of frame semantics (Fillmore, 1985) as a tool for systematically analyzing *what incels say* about themselves and others. In doing so, we take a non-anglocentric perspective, analyzing an Italian incel forum: *Il Forum dei Brutti* ('The Forum of the

¹E.g., see the well-known 2018 attack in Toronto; <https://www.theguardian.com/world/2022/jun/13/toronto-van-murders-court-victim-2018-attack> (archived version at <https://archive.is/QurUY>).

Ugly”).²

Contributions Our main contributions are:

- A new dataset comprising over 700K comments from the forum, 2.4K of which have been manually annotated with topic labels;³
- A novel frame semantics-based analysis of the corpus;
- A benchmark of preliminary machine learning experiments for predicting topic labels.

2. Background

In recent years, there has been a growing interest in exploring the manosphere and its complexity from various perspectives. Scholars investigating this network primarily examine anglophone groups or communities and come from various disciplines, including gender studies, sociology, psychology, communication sciences, economics, and linguistics. These analyses cover several important aspects, such as the representation of masculinity within different communities (Schmitz and Kazayak, 2016; Van Valkenburgh, 2021), the presence of misogyny within these groups (Farrell et al., 2019), discourses and perceptions of violence within the communities (Bryan and Warren, 2023; Lounela and Murphy, 2023), and the manosphere’s relationship with the neoliberal economic model (Banet-Weiser and Bratich, 2019).

In the linguistic field, discourse within the manosphere, and especially incel communities, has been often studied with corpus-based Critical Discourse Analysis and Cognitive Linguistics approaches (Heritage and Koller, 2020; Maxwell et al., 2020; Tranchese and Sugiura, 2021), but also by means of Computational Linguistics techniques. In particular, Jaki et al. (2019) worked on the description of the language used in incel forums in order to facilitate the identification of gendered hate speech; Jelodar and Frank (2021) analyse comments from an incel forum from a semantic perspective, through tasks such as topic modeling and opinion mining; Yoder et al. (2023) investigated the terminology used by users to construct their own identity in incel forum.

²This is the same forum that Gajo et al. (2023) used; since their paper had not yet been published at the time that we conducted our study, we collected our dataset independently of their work. However, it is likely that there is considerable overlap between our and their dataset.

³Our scraping and analysis code as well as annotations on a stand-off basis are available on <https://gitlab.com/sociofillmore/manosphrames>. Due to privacy and copyright concerns, we cannot publicly release the full corpus, but it is available upon request for interested researchers.

Despite the limited literature in the Italian context, researchers’ interest in this type of online phenomenon has been increasing over the past few years. The sociologists Farci and Righetti (2019) have focused on analyzing communities of men’s rights activists and their emergence as a response to feminism, particularly online. The International Journal of Gender Studies *About Gender* dedicated an issue (vol. 10, No. 19, 2021: *Doing masculinities online: defining and studying the manosphere*) to exploring the manosphere in both Italian and international contexts. In this monographic issue, in particular, Dordoni and Magaraggia (2021) explore the representations of gender identities and masculinity within Italian Red Pill and incel communities; De Gasperis (2021) aims to analyze interactions within the Italian incel forum Il Forum dei Brutti (also subject to analysis in this contribution) to explore the intertwining of gender identity representations and the Italian literary imagination, particularly by considering threads where users compare themselves to the poet Giacomo Leopardi. The linguists Nodari and Fiorentini (2023) propose a first description of the language used within various communities of the Italian manosphere, analysing data extracted from blogs and webpages. In particular, the authors observe how narratives within the Italian manosphere resemble those in the Anglophone context, where individuals are rigidly classified based on gender definitions. Additionally, they highlight that users primarily discuss gender relations in terms of economic transactions and employ a language that echoes scientific discourse in presenting quantitative data and para-scientific evidence in support of their ideology. Lastly, from a cross-lingual perspective, Gajo et al. (2023) collected and analysed data from English- and Italian-language incel forums. They performed tasks of automatic identification of hate speech, with a specific focus on misogyny and racism, and attempted to forecast the extent to which the posts would trigger more hateful content.

We are not aware of any previous working using frame semantics to analyze manosphere discourse.

3. Datasets

In order to examine the communication dynamics within the Italian forum from multiple perspectives, we collected two different datasets using two different approaches. It is important to note that we only accessed the parts of the forum that are public. While there are private sections of the forum that are only accessible to users who sign up, we believe that the threads that anyone can freely access are representative enough of the discourse

within incel communities; in addition, we expect these threads to have a stronger impact on the expansion of these ideologies and linguistic expressions outside these spaces. In section 3.1, we discuss the collection of the first corpus and the process of manual annotation that was performed on it; section 3.2 presents the second corpus, which was scraped automatically.

3.1. Manually Collected Corpus

The first dataset includes threads and posts that were collected manually between November 2020 and October 2021. We accessed the forum monthly and extracted random threads and posts that were actively being commented on at the time of access. This resulted in a total of 60 threads, 2,406 posts and approximately 94,000 words.

Manual Annotation One human annotator classified comments in the dataset by topic, in order to better understand the most common subjects of discussion in the forum and to obtain different sub-corpora based on topics. Assigning only one topic to each comment was often too reductive, given that many comments were complex and quite long. Therefore, two topics were assigned to each comment when necessary. The set of topics used in the annotation was based on both literature regarding incel communities (Dordoni and Magaraggia, 2021; Heritage and Koller, 2020; Tranchese and Sugiura, 2021) and the experience frequenting the forum during the year in which the data were collected. The corpus was annotated through a process of *open coding* (Khandkar, 2009). The annotator read the corpus in order to identify a first set of recurring topics; subsequently, each comment was annotated by topic, assigning one or two labels according to the topics discussed. During the classification, every new comment was compared to the previously annotated ones, in order to determine if the label was fitting and if the set of labels was sufficiently representative; if not, new labels were added to the set, or the criteria to assign the labels were adjusted. This resulted in a definite set of labels that was suitable to annotate the whole corpus.

Annotation Set We define the following topic labels: D (*donne*, ‘women’), AF (*aspetto fisico*, ‘physical aspect’), IN (*incels*), SS (*sé stessi*, ‘themselves’ – the users), U (*uomini* ‘men’), AL (*altri utenti*, ‘other users’), FDB (*Forum dei Brutti*), SOC (*società*, ‘society’), RP (*Red Pill*, a theory shared within the manosphere), M (*mondo esterno*, ‘outside world’), POL (*politica*, ‘politics’), IRR (*irrilevante*, ‘irrelevant’). See Table 1 for the descriptions

of the labels and examples⁴.

Inter-Annotator Agreement For validating the annotations, a second annotator annotated a subset of the corpus (157 comments⁵). Up to two topic labels could be assigned per comment; in order to compute set similarity between the annotations for each comment, we experimented with both Jaccard Distance and the MASI Distance (Passonneau, 2006). We obtained a Cohen’s Kappa score of $\kappa = 0.59$ using Jaccard Distance, and a more conservative $\kappa = 0.52$ using MASI Distance.

3.2. Scraped Corpus

In addition to the manually collected corpus, we also automatically scraped the entire history of threads on the *Una vita da brutto* (‘a life as a ugly person’) subforum, containing 32,560 threads posted between April 2010 and May 2023, containing 706,086 posts in total (number of posts/thread ranging from 1–2,165; median 13.0) amounting to 31.8 million words (between 1–58,815 words per thread; median 478 words).

4. Automatic Annotation

We performed a small-scale set of machine-learning experiments aiming at automatically predicting topic labels for forum comments in a multi-label setting. We tried two different approaches: on one hand, we trained a linear SVM model using as input either text-based features (raw unigram counts and TF-IDF weighted unigram counts) or frame-based features (a count vector of automatically tagged FrameNet frames); on the other hand, we used ChatGPT⁶ with several zero-shot and few-shot prompts: a prompt including only the label definitions, a prompt including one hand-picked prototypical example for each label, and a prompt including both examples and definitions.

The annotated dataset was split into 70% training samples ($n=1684$), 20% test samples ($n=481$) and 10% development samples ($n=241$). For the ChatGPT experiments, we used the *gpt-3.5-turbo* model from the OpenAI API⁷, with the default generation settings and the default system parameter.

⁴Some words in the examples provided in the table were slightly altered, in order to preserve the anonymity of the users.

⁵For choosing this subset, we randomly selected threads from the corpus until the total number of selected comments exceeded 150, or about 6% of the total corpus.

⁶<https://openai.com/blog/chatgpt>

⁷The experiments were done in May 2023.

| Code | Description | Example | | Typical frames (FFICF) | Freq. |
|------|---|--|---|---|-------|
| D | Comments about women (in general, about specific women, or about relationships) | <i>La maggior parte delle donne, a causa del ciclo mestruale, sono di umore instabile e capricciose.</i> | “The majority of women, because of their menstrual cycle, are unstable and moody” | PEOPLE (1.0), PERSONAL_REL. (0.94), DESIRING (0.79) | 863 |
| AL | Comments about other forum users | <i>Sembra che tu ti senta giudicato o attaccato anche quando non hai fatto niente di male.</i> | “It looks like you feel judged and attacked even if you haven’t done anything bad.” | STATEMENT (1.0), AWARENESS (0.99), DESIRABILITY (0.96) | 697 |
| AF | Comments about physical appearance | <i>Per avere gli addominali bisogna mangiare troppo poco, per come la vedo io.</i> | “In order to have nice abs one has to eat too little, in my opinion.” | BODY_PARTS (1.0), AESTHETICS (0.69), BODY_DESCRIPTION_HOLISTIC (0.62) | 622 |
| SS | Comments by users talking about themselves | <i>Ma cosa state dicendo? Io prima d’ora non avevo neanche mai avvicinato una ragazza.</i> | “What are you talking about? Until now I had never even approached a girl” | AWARENESS (1.0), CALENDRIC_UNIT (0.97), EMOTION_DIRECTED (0.90) | 433 |
| IN | Comments talking about incels (as individuals or as a community) | <i>Tutto questo è assurdo, noi al giorno d’oggi siamo perseguitati come gli ebrei... esagero, ma avete capito.</i> | “All this is crazy, we are persecuted just like the Jews... I am overstating, but you get what I mean.” | STATEMENT (1.0), PEOPLE (0.87), INCREMENT (0.78) | 219 |
| M | Comments about external people | | | | 212 |
| IRR | Irrelevant / not possible to assign a label | | | | 188 |
| FDB | Comments about the forum itself | | | | 153 |
| SOC | Comments about society and social issues | | | | 83 |
| U | Comments about the position of men in society | | | | 71 |
| RP | Comments related to the redpill theory | | | | 22 |
| POL | Comments related to political issues and ideologies | | | | 13 |

Table 1: Annotated topic labels in the manually collected forum, with examples and typical frames (see §5) for the top-5 most frequent topics. N.B.: a comment can have up to two topic labels.

Table 2 lists the results of our experiments. Interestingly, the best overall model are the two word count-based SVM models. ChatGPT performs substantially worse, the best-performing setup being the zero-shot one. While our results are from a single run of the model, the fact that scores are very consistent between the development and test sets suggests that the model’s performance is stable across different runs.

5. Frames

We build on recent work that applies Fillmorean frame semantics to societal issues (Minnema et al., 2022a,b,c). Fillmorean frames (Fillmore, 1985,

2006), catalogued in lexical databases such as Berkeley FrameNet (Baker et al., 1998), are pieces of conceptual information, grounded in human experience and cognition, that pick out a particular event, situation or object in the world around us. By looking at which frames are used in a text, we gain information about *what* is said about the world, but also about *how* and *from whose perspective* it is said. For example, in English, “buying” and “selling” pick out the same type of real-world event, but do so from a different event participant’s point of view. In the context of socio-politically loaded events, studying why one frame is used over another can be used as a tool for Critical Discourse Analysis (CDA): for example, when

| Input representation | Model | dev | | | test | | |
|--|--------------------------|------|------|------|------|------|------|
| | | P | R | F1 | P | R | F1 |
| Frame count vectors | Linear SVM ($C = 4$) | 0.40 | 0.35 | 0.37 | 0.36 | 0.34 | 0.35 |
| Bag-of-words (raw) vectors | Linear SVM ($C = 0.5$) | 0.63 | 0.50 | 0.56 | 0.58 | 0.54 | 0.51 |
| Bag-of-words (tf-idf) vectors | Linear SVM ($C = 2$) | 0.72 | 0.51 | 0.59 | 0.68 | 0.45 | 0.54 |
| Zero-shot prompt (only definitions) | ChatGPT | 0.42 | 0.41 | 0.41 | 0.43 | 0.43 | 0.43 |
| Few-shot prompt (only examples) | ChatGPT | 0.22 | 0.14 | 0.17 | 0.21 | 0.14 | 0.17 |
| Few-shot prompt (examples + definitions) | ChatGPT | 0.42 | 0.32 | 0.37 | 0.45 | 0.34 | 0.39 |

Table 2: Classification experiment results (P, R, and F1 are micro-averaged across topic labels)

discussing traffic incidents, the headlines “cyclist *dies* in traffic” or “driver hits and *kills* a cyclist” could both be factually correct ways of describing the same event, but convey different perspectives on the event and could imply different ideological positions (in this case about the place of cars and pedestrians in urban planning) (Minnema et al., 2022c). This type of variation in framing has also been linked to differences in event perception, e.g. with respect to who is to blame (Minnema et al., 2022a).

In the present study, we are interested in studying how users of incel forums conceptualize the world, especially when relating to gender relations. In this section, we perform two types of analysis as preliminary steps to better understanding the conceptual world of the community: (i) analyzing which frame types are most representative for the different topics discussed in the corpus; (ii) analyzing which semantic frames are used to talk about men versus women.

For both analyses, we use LOME (Xia et al., 2021) to automatically annotate our automatically-scraped corpus with FrameNet frames. LOME has been trained only on the English-language Berkeley FrameNet, but, since it has a multilingual encoder model (XLM-R, Conneau et al. 2020), can be applied to other languages in a zero-shot cross-lingual transfer setting. To our knowledge, LOME is the only recent model to have been tested on an Italian FrameNet benchmark; in Minnema et al. (2022c), we showed that it has acceptable overall performance for Italian on a standard benchmark. We also evaluated LOME’s predictions on a dataset of Italian news articles about gender-based violence and showed that applying the (original) version of LOME that was trained on English data and tested directly (zero-shot) to Italian outperformed a (new) version of LOME that was trained on both English and Italian data.

The use of automated FrameNet-based analysis has several limitations. First of all, frame semantic parsers make errors and performance may vary across different types of texts. Due to the high complexity and cost of manually annotating high-quality evaluation data, we were unable to system-

atically test LOME’s performance on our corpus for this study. However, based on a preliminary check of the automatic annotations, we found that there is at least one serious domain adaptation issue, namely that words relating to (consensual) sex are frequently mistagged. In particular, the verb *scopare* (“fuck”) was frequently mistagged as evoking violence-related frames (KILLING, RAPE); we therefore decided to exclude all instances of *scopare* from our frame analysis. We speculate that these errors may originate from the nature of Berkeley FrameNet; the original FrameNet corpus contains many descriptions of violence (e.g., in the context of geopolitics) and frames corresponding to this, but few descriptions of and frames relating to (consensual) sex. Apart from affecting parsing performance, the semantic coverage of Berkeley FrameNet also forms a limitation for our analysis by itself: we are likely to miss out on many important aspects of incel discourse due to lacking frames. For example, while there are frames related to emotion in general (e.g., EMOTION_DIRECTED) there are no frames specifically for capturing expressions of hate, or for analyzing (misogynistic) emotion descriptions such as “unstable” and “moody” (see the first example in Table 1), which are frequently found in the corpus. In the future, it could be interesting to look at expanding FrameNet’s coverage for our specific domain, as has also been proposed, for example, for the domain of gender-based violence (Dutra et al., 2023).

5.1. Typical Frames

We adopt the notion of *typical frames* proposed in Vossen et al. (2020) and Remijnse et al. (2021): a set of FrameNet frames that is most representative for a particular subcorpus within a larger corpus, and that can be automatically detected using *FF-ICF*, a modified version of the *TF-IDF* metric:⁸

$$FF-ICF_i = \frac{t_i}{f_i} \times \log \frac{m}{\sum_j^n t_j}$$

⁸Specifically, it is a derivative of C-TFIDF (Grootendorst, 2022).

where i is a subcorpus, t_i is the frequency of frame t in i , f_i is the total number of frame instances in i , and m is the total number of documents across all subcorpora (Remijnse et al., 2021, p. 233). This results in each frame being assigned a score in $[0, 1]$ for each subcorpus, with the highest-ranked frames in each subcorpus being most informative for distinguishing between subcorpora.

Table 1 shows the highest-ranking frames for each of the most frequent topic labels in the manually annotated corpus. For example, for topic D (“women”), we find PEOPLE (frequently triggered by words such as *ragazza* “girl”), PERSONAL_RELATIONSHIP (triggered by words such as *amica* “[female] friend”, *fidanzamento* “engagement”), or DESIRING (triggered by words such as *volere* “to want”). In this case, the frames correspond quite closely to aspects of the manually definition of the topic. However, we also find more specific information: for example, when looking at different instances of DESIRING, we find that forum users frequently talk about desires (romantic or otherwise) both from their own (male) perspective (e.g. *non la voglio*, “I don’t want her”) but also from the perspective of women (e.g. *loro vogliono il brivido di far eccitare i maschi e sentirsi desiderate*, “they [women] want the thrill of getting boys turned on and they want to feel desired”). Similarly, for other topics we also find frames closely corresponding to the topic definition — e.g. for topic AF (“physical appearance”) we find BODY_PARTS, triggered in phrases such as *poteva farsi un trapianto di capelli*, “[he] could get a hair transplant” or *bei lineamenti e bonus occhi, 6*, “nice features and a bonus [for her] eyes, [she gets a] 6” — but also less expected, but still informative frames: e.g., for topic SS (“talking about themselves”), we find CALENDRIC_UNIT, triggered by words like *ieri* “yesterday”, *stamattina* “this morning”, which is often an indicator of stories about the users’ personal lives, e.g. *prima di ieri non avevo neanche baciato* “until yesterday, I had never even kissed”. In addition, the presence of the frame EMOTION_DIRECTED as one of the most frequent in comments about themselves, evoked by words like *tristezza* “sadness”, *imbarazzo* “embarrassment”, *ansia* “anxiety”, suggests that the users frequently present their experiences adopting emotional narratives, choosing to share their feelings (often negative) about themselves or their life with the rest of the community.

5.2. Gender and Semantic Roles

In the typical frame analysis, we used semantic frames essentially as a way to group together related lexical units: words that express the same concept. Here, we go a step further and exploit the ability of semantic frames to relate concepts to *frame elements*: semantic roles that express

the prototypical participants of an event or situation type (e.g. in “Chiara sold a book to Tommaso”, “Chiara” fills the Seller role of the COMMERCE_SELL frame whereas “a book” fills the Goods role and “Tommaso” fills the Buyer role). By analyzing the contents of semantic role spans, we can get an insight into what is said about event participants: in which frames does a given participant appear as a role filler, and which roles does that participant fill? Here, we are interested in men vs. women: what kind of conceptual information do forum users typically convey about members of each gender? We are particularly interested in *agentive frames*: semantic frames that describe the main participant as actively doing something. In the literature about language and gender, two (seemingly) conflicting patterns have been observed relating agentivity: on the one hand, there seems to be a general tendency in several languages that in active sentences, men are more often expressed as syntactic subjects than as objects, while women are more often expressed as syntactic objects (Kotek et al., 2021; da Cunha and Abeillé, 2022). On the other hand, linguists and feminist scholars have identified patterns of language use where actions by men are described in a *de-agentivized* way. For example, according to Penelope (1990), expressions without an explicit agent such as “it is widely understood that...” or “the only reason for ordering a war ...” are frequently found in contexts in which the only plausible implicit agent is a man or a group of men, and omitting the agent in such cases can contribute to presenting men’s experiences as ‘universal’ or to present male actions as inevitable and to absolve the agents of responsibility. This latter pattern has also been observed in a number of recent empirical (cognitive, corpus-based and computational) studies on the reporting of gender-based violence, where journalists and other writers frequently use agent-removing constructions, which has been shown to decrease the level of blame that readers attribute to the perpetrator (Henley et al., 1995; Bohner, 2002; Pinelli and Zanchi, 2021; Meluzzi et al., 2021; Minnema et al., 2022a). In the light of this, it is interesting to observe agentivity patterns in the language use of the incel community.

We implement our analysis as follows: first, we perform a keyword search in all automatically detected semantic roles in the 32K threads scraped from the forum for words referring to men or women, respectively.⁹ Next, we identify agen-

⁹We used the following keywords: *donna, donne, ragazza, ragazze, lei, np, co* (“woman”, “women”, “np”, “co”, “girl”, “girls”, “she/her”); *np* and *co* are abbreviations of pejorative terms frequently used in the corpus to mean “woman”) for women, and *uomo, uomini, ragazzo, ragazzi, lui* (“man”, “men”, “boy”, “boys”, “he/him”).




Figure 1: Roles mentioning women




Figure 2: Roles mentioning men

itive frames by checking for each frame if it (indirectly) inherits from either TRANSITIVE_ACTION or INTENTIONALLY_ACT, and identifying which of the frame’s roles expresses an agent.¹⁰ Note that both

¹⁰FrameNet has a rich and complex graph structure of relating frames and frame elements to each other. Here, we only use the *inheritance* frame-to-frame relation. We allow for both direct and indi-




Figure 3: Agentive roles mentioning women

of these steps have limited recall: our keyword search does not include all possible words for referring to men and women, and also ignores all cases where a participant is expressed using a proper name or anaphorically as a pro-drop subject. Moreover, the FrameNet hierarchy is incomplete, and not all frames that are semantically agentive can be detected as such using the inheritance hierarchy.

Figures 1 and 2 show the top-15 most frequent semantic roles across all frames that match one of the keywords. The first observation is that women are mentioned much more frequently than men. For both genders, “man/men” or “woman/women” are the most commonly matched keywords (we find 605K matches in total for the female keywords, of which *donna/donne* accounts for 317K; for men, *uomo/uomini* account for 126K out of 257K total matches). In both cases, the by far most frequent frame-role combination is PEOPLE.Person, which is expected as any and all mentions of “man” or “woman” (and variants/synonyms of those words) trigger this frame. The next-most frequent frames are more interesting: for example, we find 13.5K instances of women matching STATEMENT.Message role, i.e., being mentioned as the content of some-

rect inheritance; e.g., KILLING inherits directly from TRANSITIVE_ACTION (where the KILLING.Killer role is mapped to TRANSITIVE_ACTION.Agent; on the other hand, COOKING_CREATION inherits from INTENTIONALLY_CREATE (mapping the Cook role to Creator), which in turn inherits from INTENTIONALLY_ACT (mapping Creator to Agent).




Figure 4: Agentive roles mentioning men

thing that someone says (e.g., in *ma almeno non potete dire [che nessuna ragazza starebbe mai con voi]* “but at least you cannot **say** [that no girl will ever be with you]”¹¹). Similarly, we find 9.5K mentions of women matching AWARENESS.Content, i.e. being mentioned in the content of a (stated) knowledge or belief (e.g. *io non so [se le donne cercano un uomo per riprodursi]* “I don’t **know** [if women are looking for a man to reproduce]”). Interestingly, these same frame/role pairs are also in the top-5 for mentions of men.

Moving to agentive frames (Figures 3 and 4), INTENTIONALLY_ACT.Agent is the most frequent role for both genders; the majority of these instances correspond to the subject of the verb *fare* “do/make/act” (e.g. *io non ho mai visto [ragazze] fare così*), “I have never seen [women] **acting** like this”). In the rest of the top-5, for women we find CHOOSING.Cognizer (e.g. *Tanto è la donna che sceglie* “It’s the woman who chooses, anyway”), GIVING.Donor, and PERCEPTION_ACTIVE.Perceiver_agentive (e.g., “seeing”, “hearing”). By contrast, for men, we find RAPE.Perpetrator, ATTEMPT.Agent, SELF_MOTION.Mover (“going”), and CHOOSING.Donor.

5.3. Diachronic Analysis

Since our corpus spans more than a decade’s worth of posts (April 2009–May 2023), we were

¹¹Frame trigger highlighted in boldface, semantic role instance between square brackets.




Figure 5: Comments by year

also able to start exploring the question of how incel discourses changes over time. Figure 5 shows the number of comments scraped from *Il Forum dei Brutti*. The annual number of comments steadily increases after 2012, rising to over 127,000 posts in 2021. This trend ended in 2022, which saw substantially fewer comments; our data for 2023 are incomplete. Note that a decrease in comments in our corpus does not necessarily imply a decrease in overall activity: certain sections of the forum are private, and we are unable to monitor the activity trends there. Figures 6 and 7 show the top-7 of most frequent role mentions (collectively accounting for at least 30% of the total number of matching role mentions in each year) that match women-related and men-related keywords, respectively. The clearest visible pattern is the stability over time of the most frequent frames: the top frames are mostly the same ones for each year. However, an interesting development is the decline in frequency of AESTHETICS.Entity (e.g., sentences like *Ho ritenuto da sempre [le donne indiane] le più belle del mondo*, “I have always considered [Indian women] the most **beautiful** in the world”), for both women and men: relating to women, the relative frequency of this frame-role pair peaked at 2.5% in 2011, and then entered a steep decline, falling to 1.4% in 2016 and 0.9% in 2021; relating to men, there is a similar pattern, but with a peak in 2013 (3.1%), falling to 1.0% in 2021. While it is hard to draw strong conclusions from this, it could be an indicator that the importance of discussing physical attractiveness is declining or increasingly expressed in a different way over time.




Figure 6: Roles mentioning women by year




Figure 7: Roles mentioning men by year

6. Conclusion

This paper introduced a large corpus of comments extracted from *Il Forum dei Brutti*, an Italian online incel community. Our corpus consists of 2.4K comments that have been manually collected, analyzed and annotated with topic labels, and a further 32K that have been automatically scraped and automatically annotated with FrameNet annotations. We also provided a benchmark with basic machine learning experiments for automatically predicting topic labels. Our experiments yielded mixed results: while simple SVM-based approaches work surprisingly well, ChatGPT performed surprisingly poorly. Finally, we performed an automatic FrameNet-based analysis of the contents of the corpus. In the first step of our analysis, we showed the usefulness of typical frame detection (Vossen et al., 2020; Remijnse et al., 2021) for analyzing topic-based subcorpora. In the second step of the analysis, we showed that forum users talk twice as frequently (explicitly) about women than about men — this is true both across all frames and when only including frame instances where a man or woman is described in an agentive way — and we found interesting parallels and differences in the patterns of semantic roles in which men and women appear most commonly. We also took a first peek at how incel discourse changes over time; clearly, there is a lot of room for expansion here. In particular, Baele et al. (2023) observed a general increase of violent ex-

tremist discourse in incel forums in recent years; it could be interesting to investigate this through the lens of violence-related frames (e.g. KILLING, CAUSE_HARM, etc.).

7. Acknowledgements

Author G.M.’s contribution was funded by the Dutch National Science organisation (NWO) through the project *Framing situations in the Dutch language*, VC.GW17.083/6215 (see www.dutchframenet.nl). We would like to thank the anonymous reviewers for their extensive comments and insights.

8. Limitations

Some of the peculiarities of the language used by the users of the forum can produce various challenges in computational tasks like automatic frame analysis and annotation. In particular, we came across three issues that we tried solve, but that partially still represent areas for improvement: the first two have to do with the way the users refer to women and to themselves, whereas the last one concerns lexical opacity typical of these type of communities.

To identify semantic roles mentioning women we used a list of keywords that seemed to cover the majority of cases (cfr. note 2), but in some comments the users refer to women by only using the plural personal pronoun *loro* (“they/them”), or even just by relying on verbal morphology, expressing the verb in the third-person plural form, as Italian is a pro-drop language (e.g. *pens-ano*, “[they] think”; *dic-ono*, “[they] say”). These strategies contribute to the idea that women - seen as a homogeneous whole - are the out-group with respect to the users of the forum. At the same time, the keywords used for the men were not always sufficient, as they often talk about themselves and the community in general using the personal pronoun *noi* (“we/us”), sometimes the adverb *qui* (“here” [in the forum]), or, similarly to the case of women, they simply use the first-person singular form of the verb (e.g. *sembr-iamo*, “[we] seem”), marking themselves as in-group. Lastly, the presence of neologisms and acronyms constitutes an obstacle for the automatic exploration of the corpus. If not all the words are understood by the models, it is harder to obtain a correct classification of the comments; similarly, analyzing semantic frames and roles is more difficult if there are opaque terms, e.g. *CO* for *cessa obesa* (“ugly fat [woman]”), *zerbinare* (lit. “doormatting”, the act of submitting completely to a woman in hopes of being noticed and, ultimately, loved by her).

9. Bibliographical References

- Stephane Baele, Lewys Brace, and Debbie Ging. 2023. [A diachronic cross-platforms analysis of violent extremist language in the incel online ecosystem](#). *Terrorism and Political Violence*.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational Linguistics - Volume 1, COLING '98*, pages 86–90.
- Sarah Banet-Weiser and Jack Bratich. 2019. From pick-up artists to incels: Con (fidence) games, networked misogyny, and the failure of neoliberalism. *International Journal of Communication*, 13:1.
- Gerd Bohner. 2002. Writing about rape: Use of the passive voice and other distancing features as an expression of perceived responsibility of the victim. *British Journal of Social Psychology*, 40:515–529.
- Elliott Bryan and Shane Warren. 2023. The manosphere under the microscope: a critical discourse analysis of the news media reporting of rowan baxter’s murder of hannah clarke and her family. *Feminist Media Studies*, pages 1–18.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Yanis da Cunha and Anne Abeillé. 2022. [Objectifying women? a syntactic bias in French and English corpora](#). In *Proceedings of the First Workshop on Language Technology and Resources for a Fair, Inclusive, and Safe Society within the 13th Language Resources and Evaluation Conference*, pages 8–16, Marseille, France. European Language Resources Association.
- Arianna De Gasperis. 2021. “giacomino uno di noi”. letteratura italiana e pratiche di maschilità nel forum dei brutti. *AG About Gender-International Journal of Gender Studies*, 10(19).
- Annalisa Dordoni and Sveva Magaraggia. 2021. Modelli di mascolinità nei gruppi online incel e red pill: Narrazione vittimistica di sé, deumanizzazione e violenza contro le donne. *AG About Gender-International Journal of Gender Studies*, 10(19).
- Livia Dutra, Arthur Lorenzi, Lorena Larré, Federico Belcavello, Ely Matos, Amanda Pestana, Kenneth Brown, Mariana Gonçalves, Victor Herbst, Sofia Reinach, Renato Teixeira, Pedro de Paula, Alessandra Pellini, Cibele Sequeira, Ester Sabino, Fábio Leal, Mônica Conde, Regina Grespan, and Tiago Torrent. 2023. [Building a frame-semantic model of the healthcare domain: Towards the identification of gender-based violence in public health data](#). In *Jornada de Descrição do Português: Anais do XIV Simposio Brasileiro de Tecnologia da Informação e da Linguagem Humana*.
- Manolo Farci and Nicola Righetti. 2019. Italian men’s rights activism and online backlash against feminism. *Rassegna Italiana di Sociologia*, 60(4):765–781.
- Tracie Farrell, Miriam Fernandez, Jakub Novotny, and Harith Alani. 2019. Exploring misogyny across the manosphere in reddit. In *Proceedings of the 10th ACM conference on web science*, pages 87–96.
- Charles J. Fillmore. 1985. Frames and the semantics of understanding. *Quaderni di semantica*, 6(2):222–254.
- Charles J. Fillmore. 2006. Frame semantics. In D. Geeraerts, editor, *Cognitive Linguistics: Basic Readings*, pages 373–400. De Gruyter Mouton, Berlin, Boston. Originally published in 1982.
- Paolo Gajo, Arianna Muti, Katerina Korre, Silvia Bernardini, and Alberto Barrón-Cedeño. 2023. [On the identification and forecasting of hate speech in inceldom](#). In *Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing*, pages 373–384, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Debbie Ging. 2019. Alphas, betas, and incels: Theorizing the masculinities of the manosphere. *Men and masculinities*, 22(4):638–657.
- Maarten Grootendorst. 2022. [Bertopic: Neural topic modeling with a class-based tf-idf procedure](#).
- Nancy M. Henley, Michelle Miller, and Jo Anne Beazley. 1995. [Syntax, semantics, and sexual violence: Agency and passive voice](#). *Journal of Language and Social Psychology*, 14(1-2):60–84.
- Frazer Heritage and Veronika Koller. 2020. Incels, in-groups, and ideologies: The representation of gendered social actors in a sexuality-based online community. *Journal of Language and Sexuality*, 9(2):152–178.

- Sylvia Jaki, Tom De Smedt, Maja Gwózdź, Rudresh Panchal, Alexander Rossa, and Guy De Pauw. 2019. Online hatred of women<? br?> in the incels. me forum: Linguistic analysis and automatic detection. *Journal of Language Aggression and Conflict*, 7(2):240–268.
- Emma A. Jane. 2014. “your a ugly, whorish, slut” understanding e-bile. *Feminist Media Studies*, 14(4):531–546.
- Hamed Jelodar and Richard Frank. 2021. Semantic knowledge discovery and discussion mining of incel online community: Topic modeling. *arXiv preprint arXiv:2104.09586*.
- Shahedul Huq Khandkar. 2009. Open coding. *University of Calgary*, 23:2009.
- Hadas Kotek, Rokker Dockum, Sarah Babinksi, and Christopher Geissler. 2021. [Gender bias and stereotypes in linguistic example sentences](#). *Language*, 97(4):653–677.
- Emilia Lounela and Shane Murphy. 2023. Incel violence and victimhood: Negotiating inceldom in online discussions of the plymouth shooting. *Terrorism and Political Violence*, pages 1–22.
- Kate Manne. 2017. *Down girl: The logic of misogyny*. Oxford University Press.
- December Maxwell, Sarah R Robinson, Jessica R Williams, and Craig Keaton. 2020. “a short story of a lonely guy”: A qualitative thematic analysis of involuntary celibacy using reddit. *Sexuality & Culture*, 24(6):1852–1874.
- Chiara Meluzzi, Erica Pinelli, Elena Valvason, and Chiara Zanchi. 2021. Responsibility attribution in gender-based domestic violence: A study bridging corpus-assisted discourse analysis and readers’ perception. *Journal of pragmatics*, 185:73–92.
- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022a. [Dead or murdered? Predicting responsibility perception in femicide news reports](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1078–1090. Association for Computational Linguistics.
- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022b. [SocioFillmore: A tool for discovering perspectives](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 240–250, Dublin, Ireland. Association for Computational Linguistics.
- Gosse Minnema, Gaetana Ruggiero, Marion Bartl, Sara Gemelli, Tommaso Caselli, Chiara Zanchi, Marco te Brömmelstroet, and Malvina Nissim. 2022c. [Responsibility framing under the magnifying lens of NLP](#). *Computational Linguistics in the Netherlands*, 12:207–233.
- Rosalba Nodari and Ilaria Fiorentini. 2023. Gli alpha e i beta. la lingua della uomosfera nelle comunità online italiane. In V. Caruso and M. Maffia, editors, *Vecchie e nuove forme di comunicazione diseguale: canali, strutture e modelli*, pages 307–321. Officinaventuno.
- Rebecca Passonneau. 2006. [Measuring agreement on set-valued items \(MASI\) for semantic and pragmatic annotation](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Julia Penelope. 1990. *Speaking freely: Unlearning the lies of the fathers’ tongues*. The Athene Series. Permagmon Press, Elmsford, New York.
- Erica Pinelli and Chiara Zanchi. 2021. Gender-based violence in italian local newspapers: How argument structure constructions can diminish a perpetrator’s responsibility. *Discourse Processes between Reason and Emotion: A Post-disciplinary Perspective*, page 117.
- Levi Remijnse, Marten Postma, and Piek Vossen. 2021. [Variation in framing as a function of temporal reporting distance](#). In *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*, pages 228–238, Groningen, The Netherlands (online). Association for Computational Linguistics.
- Rachel M Schmitz and Emily Kazyak. 2016. Masculinities in cyberspace: An analysis of portrayals of manhood in men’s rights activist websites. *Social Sciences*, 5(2):18.
- Alessia Tranchese and Lisa Sugiura. 2021. “i don’t hate all women, just those stuck-up bitches”: How incels and mainstream pornography speak the same extreme language of misogyny. *Violence against women*, 27(14):2709–2734.
- Shawn P Van Valkenburgh. 2021. Digesting the red pill: Masculinity and neoliberalism in the manosphere. *Men and masculinities*, 24(1):84–103.
- Piek Vossen, Filip Ilievski, Marten Postma, Antske Fokkens, Gosse Minnema, and Levi Remijnse.

2020. [Large-scale cross-lingual language resources for referencing and framing](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3162–3171, Marseille, France. European Language Resources Association.

Patrick Xia, Guanghui Qin, Siddharth Vashishtha, Yunmo Chen, Tongfei Chen, Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, and Benjamin Van Durme. 2021. [LOME: Large ontology multilingual extraction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 149–159, Online. Association for Computational Linguistics.

Michael Yoder, Chloe Perry, David Brown, Kathleen Carley, and Meredith Pruden. 2023. [Identity construction in a misogynist incels forum](#). In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 1–13, Toronto, Canada. Association for Computational Linguistics.

10. Language Resource References

Baker, Collin F. and Fillmore, Charles J. and Lowe, John B. 1998. *FrameNet*. Berkeley International Computer Science Institute, version 1.7. PID <https://framenet.icsi.berkeley.edu/>.