

Effects of Forum Participation on Incel Identity Salience: A Computational Analysis

Author:
Leo Peltola

5.5.2024
Turku

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

Bachelor's Thesis

Major: Social Sciences

Author: Leo Peltola

Title: Effects of Forum Participation on Incel Identity Saliency: A Computational Analysis

Supervisor: Senior Lecturer Michael Laakasuo

Page count: 19 pages

Date: 5.5.2024

Keywords: incel, manosphere, social identity, computational social sciences

Abstract

Involuntary celibates, better known as *incels*, are a group of men dissatisfied with their lack of romantic and intimate relationships. In this study, I seek to find out whether participation in a popular online forum *incels.is* strengthens users' identities as incels. Using prototypical behavior as a proxy for identity saliency, I analyzed a sample of one million posts from the forum using computational language analysis methods. Prototypical behavior was operationalized into three metrics: language toxicity, obsession with sex-related topics, and obsession with feminism. This study found a small negative correlation between incel forum participation and behavior prototypicality, suggesting forum participation may contribute to deradicalization.

Table of Contents

1	Introduction	4
1.1	The Involuntary Celibate and His Origin	4
1.2	Social Identity and Prototypicality	6
1.3	Current Study	7
2	Method	9
2.1	Dataset	9
2.2	Procedure	9
2.2.1	Toxicity and Obsession With Sex-Related Topics	9
2.2.2	Obsession With Feminism	10
2.3	Statistical Analysis	11
3	Results	12
4	Discussion	13
4.1	Limitations	14
4.2	Future Studies	14
4.3	Conclusion	15
5	References	16

1 Introduction

Involuntary celibates, more commonly known as *incels*, are a decentralized community of men unsatisfied with their lack of romantic or intimate relationships (Sugiura, 2021, Chapter 1). In recent years, researchers, media, and the general public have paid increasing attention to incels because of their radical views and associations with real-world violence (Solea & Sugiura, 2023; Sugiura, 2021, Chapter 1). This study aims to investigate whether participation in an online incel forum *incels.is* contributes to the strengthening of an individual's identity as an incel.

Drawing on the social identity approach (SIA) and the concept of prototypicality, this study employs computational language analysis methods to examine a sample of one million posts from the popular incel forum *incels.is*. By operationalizing prototypical behavior into three metrics - language toxicity, obsession with sex-related topics, and obsession with feminism - and analyzing their correlation with post rank order, this study seeks to examine the association between forum participation and prototypical behavior.

1.1 The Involuntary Celibate and His Origin

Self-identified incels are part of a larger group called the manosphere (O'Malley et al., 2022). The manosphere is a collection of separate but perhaps ideologically loosely connected communities united by an interest in men's issues and an alleged crisis of masculinity (Horta Ribeiro et al., 2021; Van Valkenburgh, 2021). Some of the more visible communities within the manosphere include *involuntarily celibates* (incels), *pick-up artists* (PUA), *men going their own way* (MGTOW), and *men's rights activists* (MRA) (Lilly, 2016).

As explained by Sugiura (2021, Chapter 2) the manosphere originates from the men's liberation movement (MLM) established in the early 1970s alongside with women's liberation movement. The two sibling movements shared a belief in the harmfulness of patriarchy to both sexes. By the late 1970s, the conservative and moderate wings of MLA transformed into the men's rights movement (MRM) with the major point of disagreement being the importance of sex roles. MRAs started to increasingly move to the web and later transformed into what is now labeled as the manosphere. (Sugiura, 2021, p. 2)

Despite the origins of the manosphere, and consequently incels, being pro-feminist, the movement has now come to be commonly associated with anti-feminist sentiments (Dickel &

Evolvi, 2023; Lindsay, 2022). Much of previous research has repeatedly highlighted misogynistic and sexist tones and values in incel communities (e.g. Dickel & Evolvi, 2023; O'Malley et al., 2022; Price, 2023)

Manosphere communities diverge in their exact origins and preferred platforms (Sugiura, 2021). The social media website Reddit has been popular within the manosphere among other communities. Because of Reddit's research-friendly data policies and openness, the platform has also been a popular source of social media-based text analysis research, including recent incel and manosphere research as well (e.g. Habib et al., 2022; Horta Ribeiro et al., 2021; Stijelja & Mishara, 2023). However, following Reddit changing its hate speech policies and increasingly banning manosphere and other communities, these communities have started transitioning to new, dedicated forums hosted by community members or sympathetic individuals. Previously manosphere communities were required to prohibit radical viewpoints to allow their continued existence on Reddit. The new independent forums lack these kinds of restrictions and have been identified by some researchers as having considerably higher rates of toxic language compared to older manosphere subreddits (Horta Ribeiro et al., 2021).

A prevalent concept within the manosphere, and especially incels, is the pill analogy (Van Valkenburgh, 2021). This concept references the 1999 movie *The Matrix* (Wachowski & Wachowski, 1999) in which Neo, the main character, is offered a choice between a blue and a red pill (Pelzer et al., 2021, p. 5); the latter allowing the protagonist to see life's ugly truths and the former letting him continue his naïve life from before. Within the manosphere, the *blue pill* generally refers to someone who disagrees with the manosphere worldview and/or is not aware of the movement. These people are also referred to as *normies* (normal people). The *red pill*, in contrast to the blue pill, signals agreement with the manosphere worldview. Some social scientists have viewed the so-called red pill ideology as the uniting factor of manosphere communities (Horta Ribeiro et al., 2021, p. 206). Red pill ideology represents the core views inherited by the manosphere from MRM (Sugiura, 2021, Chapter 2). While hard to define, previous research has claimed that red pill ideology is anti-feminist and misogynistic (Dickel & Evolvi, 2023, p. 1394; Ging, 2017, p. 8; Sugiura, 2021, Chapter 1). The *black pill*, mostly in use by incels, has considerable overlap with the red pill. The major difference is that people who have accepted the red pill ("redpillers") generally encourage self-improvement and confidence to improve one's status and mating success while "blackpillers" believe this is not possible. The black pill emphasizes the importance of genetics in looks and generally

claims that incels will never be considered as potential partners by women (Pelzer et al., 2021, p. 5).

Another characteristic of incels and, to a lesser extent, the manosphere is the perceived importance of appearance and attractiveness. *Lookism* is a term used for appearance-based discrimination (Warhurst et al., 2009), which incels perceive themselves as victims of (Halpin, 2022; Jones, 2020). In a lookist hierarchy, the most attractive people occupy the top with individuals of decreasing attractiveness getting the lower spots. For men, this hierarchy is usually divided into three tiers. The most attractive men are called *chads* or *alpha males*. They are perceived as having high access to romantic and intimate relationships and encounters. The middle tier is inhabited by *beta males* or *normies* (“*normies*” is used in multiple contexts; it may also refer to all non-incels or non-redpillers). Incels make up the lowest tier, being unable to enter relationships despite their willingness to. The most attractive women are called *stacies* and the rest *beckies*. A female equivalent of incels is not perceived to exist as the phenomenon is considered unique to men. Incels and the black pill deviate from other manosphere communities by viewing this hierarchy as being caused by genetics and, consequently, immutable (Furl, 2022; Halpin, 2022; Meier & Sharp, 2024).

Ribeiro et al. (2021) show that seemingly higher levels of “toxic language” – left undefined but measured by Perspective API (Jigsaw & Google, n.d.), a machine learning model, and a dictionary-focused approach – are used within manosphere communities compared to the wider internet. Older manosphere communities such as MRA and PUA were identified as less toxic than newer ones. Previous research has also identified incel communities to be particularly high in language toxicity compared to other manosphere communities when measured with Perspective API (Pelzer et al., 2021). This toxicity, defined by the authors as blatantly aggressive and demeaning content, is not only directed towards other groups but also, interestingly, other incels.

1.2 Social Identity and Prototypicality

The social identity approach (SIA) is a theoretical framework that explains how individuals categorize themselves into groups based on shared characteristics (Hogg, 2016, p. 6). This approach has evolved from its original focus on intergroup behavior and status differences between groups to encompass a broader range of human interaction and group categorization processes (Hornsey, 2008). At its core, SIA is made up of two key theories: social identity theory (SIT) and self-categorization theory (SCT).

Social identity theory (SIT) models individuals as seeking membership of high-status groups in order to achieve a positive social identity (Hogg, 2016). Groups are divided into *ingroups* (groups to which an individual belongs) and *outgroups* (groups to which they do not belong). However, this categorization is not always mutually agreed upon, as individuals may disagree on group membership or even the existence of certain social groups (Benwell & Stokoe, 2007).

A concept central to SIA is *salience*, which refers to the perceived membership of a group. When an individual perceives themselves or is perceived by others as part of a group, that group becomes a salient social group. A salient self-categorization can also be referred to as a salient social identity (Turner & Oakes, 1986).

Accentuation is another important phenomenon in SIA. It describes the tendency of individuals to view ingroup members as more similar and outgroup members as more different to themselves. This process can lead to group polarization, where the attitudes and behaviors of group members become more extreme and homogeneous (Hogg, 2016, p. 11).

Self-categorization theory (SCT) introduces the concept of *prototypicality*, which refers to the degree to which an individual embodies the characteristics and norms of a particular social group. Typically, there is wide agreement within a group about the prototypes of both the ingroup and outgroups, serving as cognitive representations of what the group is like and how it differs from other groups (Hogg, 2016). A group is considered to have a prototype, an abstract model of a perfect group member that embodies all the qualities of that group.

Depersonalization, a process of self-stereotyping, occurs when individuals modify their behavior and attitudes to align more closely with a group's prototype, leading to increased prototypicality (Hogg, 2016, pp. 11–12).

In the context of incels, the concepts and processes of social identity may be applied. As incel identity becomes more salient through participation in forum discussions, it should reflect in the prototypicality of their behavior as they undergo depersonalization. Specifically, the more a user identifies as an incel, the more prototypical their behavior would be expected to become.

1.3 Current Study

This study aims to investigate whether participation in the online forum *incels.is* strengthens users' incel identity salience. Prototypical behavior, a key indicator of identity salience as

discussed in previous literature, is used as a proxy for measuring incel identity. The following hypothesis is presented:

Prototypical behavior will be positively associated with a forum post's rank order.

Here we will define the prototypicality of incel community membership with the users' tendency to engage in the use of certain linguistic indicators in their posts: 1) the level of "toxicity" markers in their posts, 2) seeming obsession with sex-related topics, and 3) seeming obsession with feminism.

2 Method

The three metrics presented previously are employed to examine the prototypicality of user behavior. These metrics were selected based on their prototypicality of incels, as discussed in Chapter 1.1, and their feasibility for computational measurement.

Each of the metrics was examined for correlation with post rank order (e.g. *1st post by user*, *37th post by user*). If forum participation results in a stronger incel identity the scores should increase as the user posts more, as in the hypothesis presented earlier.

2.1 Dataset

The dataset comprises the posts and usernames from the incel forum *incels.is* (formerly *incels.me* and *incels.co*). The website was scraped in March 2024 using Python and the BeautifulSoup library (Richardson, 2007). The dataset contains 12,250,240 posts across 525,451 threads, contributed by 24,534 users. The posts span from February 2017 to March 2024.

2.2 Procedure

The dataset was first preprocessed to prepare it for computational language analysis. All quotes, images, and videos were removed to ensure only original text written by forum users remained. Excess whitespace (e.g., tabs) was stripped and stop words (e.g. “*of*”, “*at*”) were removed. The text was lemmatized, reducing words to their base form (e.g. “*universities*” to “*university*”, “*boys*” to “*boy*”). The preprocessing was performed using the Natural Language Toolkit (Bird et al., 2009) and WordNet (Miller, 1995).

The analysis used Python (Van Rossum & Drake, 2009) and Pandas, a data analysis library (McKinney, 2010; The pandas development team, 2024). Because of limited computational resources only a sample of one million posts was analyzed instead of the full dataset.

2.2.1 Toxicity and Obsession With Sex-Related Topics

Both language toxicity and obsession with sex-related topics were computed using Detoxify, a Python library for toxicity classification (Hanu & Unitary team, 2020). The model

“*unbiased*” was used¹. This library was selected due to its open-source code and previous use in research (e.g. Bouchaud et al., 2023; Fang et al., 2024; Snoswell et al., 2023).

Given a text section, Detoxify returns a set of values between zero and one, representing various toxicity metrics. Scores (e.g., *severe_toxicity*, *sexual_explicit*) computed by Detoxify represent the likelihood that a random person would assign the given label to the text. In other words, the scores are not absolute but based on subjective assessment; biases from dataset annotators may be present.

For evaluating language toxicity, as the term is used by Detoxify, the *severe_toxicity* value given by the library was used. *Severe_toxicity* was chosen because it has been used in previous incel language toxicity research (e.g. Horta Ribeiro et al., 2021; Pelzer et al., 2021)². For evaluating obsession with sex-related topics, the *sexual_explicit* score given by Detoxify is used. Table 1 shows examples of Detoxify scores given for various sentences picked from the dataset.

Table 1

Example texts from incels.is and their scores computed with Detoxify

Text	toxicity	severe_toxicity	sexual_explicit
they'll never admit this	0.000523	0.000001	0.000012
I'll kill you	0.974651	0.020522	0.010430
i hate woke tranny shit	0.989405	0.091534	0.084945

Note. *toxicity* is included here as an example of other values given by detoxify, it is not used in the analysis.

2.2.2 Obsession With Feminism

Obsession with feminism was measured by examining the relative frequency of feminist keywords in posts, as suggested by Cohen et al. (2014) and used by Habib et al. (2022). A list of 37 feminism-related keywords was manually curated. The keywords were sourced from relevant peer-reviewed articles related to both feminism and incels, books, and from examination of incel discussions on the topic. Both positive (e.g. “*feminist*”) and negative (e.g. “*feminazi*”) keywords were included.

¹ “*unbiased*” was chosen because it includes the *sexual_explicit* score, which the “*original*” does not.

² These studies use Perspective API (Jigsaw & Google, n.d.) instead of Detoxify but the meaning is comparable due to shared training data.

The obsession scores were calculated with term frequency-inverse document frequencies (TF-IDF). The TF-IDF score, which is calculated individually for every post, considers how often a term is present in the post (term frequency) and how often the term is present in other documents within the dataset (inverse document frequency). This causes terms with high relevancy to a post to get a high score. Terms that are not relevant to a post get a low score. The scores range from 0.0 to 1.0.

2.3 Statistical Analysis

After computing the prototypical behavior metrics correlation with post rank order was examined. The dataset's distribution was analyzed through graphical analysis due to its size, which would make analytical methods (e.g. Shapiro-Wilk) inaccurate. The dataset did not follow a normal distribution and as such non-parametric statistical analysis methods were used.

Spearman correlation coefficients were calculated between each of the metrics and post rank order using Statsmodels, a statistics library for Python (Seabold & Perktold, 2010). The method was chosen because it gives statistically robust results for non-parametric (non-normally distributed) data. Confidence intervals for the coefficients were calculated using Fisher transformations.

The spearman correlation coefficient is a numeric value between -1 and 1 with zero suggesting no relationship between the variables. Negative values indicate a negative correlation (one variable increases when the other decreases) and positive ones a positive correlation (one variable increases when the other does). For example, -0.15 suggests a small negative association and 0.83 a strong positive one.

3 Results

This study found small negative correlations between prototypical behavior metrics and post rank order. As seen in Table 2 the results indicate a small negative association between forum participation and prototypicality. This contradicts the hypothesis that behavior becomes more prototypical as a user posts more.

Table 2

Spearman correlation coefficients of prototypical behavior metrics and post rank order.

Metric	Correlation	P-value	95 % lower	95 % upper
Toxicity	-0.0988	<0.001	-0.1007	-0.0969
Obsession with sex-related topics	-0.0996	<0.001	-0.1015	-0.0977
Obsession with feminism	-0.0354	<0.001	-0.0373	-0.0335

Note. High P-values would indicate that the results are not statistically representative. Low ones, such as seen here, mean the results are statistically significant.

4 Discussion

This study investigated whether participation in the online incel forum *incels.is* strengthened users' identification as incels, using prototypical behavior as a proxy for incel identity salience. The analysis used three metrics to quantify prototypicality: language toxicity, seeming obsession with sex-related topics, and seeming obsession with feminism. The results showed a small negative correlation between these metrics and post rank order, suggesting that forum participation may lead to decreased prototypical behavior and, consequently, weaker incel identity salience. This contradicts with the hypothesis that users would behave more prototypically as they participate more because of the incel identity growing more salient.

These findings may have implications for the dynamics of social identity processes and their role in the formulation of extreme (incel) beliefs. Should participation in *incels.is* lead to decreased incel identity salience, as this study indicates, the salience decrease may be considered a sign of deradicalization because of participation. One possible explanation for this phenomenon could be that interaction with peers and people with similar experiences helps process one's own feelings (Horgan et al., 2013; Shalaby & Agyapong, 2020).

While prototypicality and normative behavior have been proposed as key indicators of social identity salience in the social identity approach (SIA) literature (Hogg, 2016, p. 11), it is possible that this relationship may not hold true in all contexts or the motivation to do so-called "identity work" is decreased (Winkler, 2018). The unique characteristics of an online incel community, such as its decentralized structure and the anonymity of its members, may influence the manifestation of social identity in ways that differ from the predictions of the most popular formulations of SIA (Postmes et al., 2001). Additionally, the dynamics of identity salience and its representation may differ for self-perceived low-status groups such as incels.

The findings of this study contribute to the growing body of research on incel communities and the role of online platforms in the radicalization process. While previous studies have highlighted the high levels of toxicity and extreme views present in incel forums such as *incels.is* (Horta Ribeiro et al., 2021; Pelzer et al., 2021), this analysis raises new questions about intra-forum dynamics. The results suggest that incel forums may not serve as primary drivers of radicalization but may even serve as possible platforms of deradicalization.

4.1 Limitations

The metrics selected for this study - language toxicity, seeming obsession with sex-related topics, and seeming obsession with feminism - were chosen based on their prototypicality to incels and their feasibility for computational measurement. However, it is possible that these metrics fail to capture prototypical behavior within *incels.is* or incels as a whole. The results of this study should be evaluated in context and with the methodological limitations in mind (Ahad et al., 2023).

The fact that only one forum was analyzed also limits the extent of the conclusions that can be drawn from this study. Currently, it is difficult to draw widely convincing results since only a limited section of the incel community is examined. Regardless, this study provides new insight into how prototypicality changes because of participation in one popular incel forum.

The use of computational natural language processing (NLP) methods, such as Detoxify, introduces limitations in the accurate measurement of the chosen metrics. NLP algorithms may struggle to differentiate between nuanced aspects of language, such as tone, intention, and context, which can lead to misclassifications of toxicity or other linguistic features. For example, the presence of sensitive identity labels (e.g., “gay”, “Islam”) or incel-specific terminology may cause non-toxic posts to be incorrectly identified as toxic and the opposite, introducing unwanted noise into the analysis (Dixon et al., 2018). However, since the analysis methods are applied consistently the results of the study remain meaningful.

Additionally, the use of post rank order as an approximator of forum engagement may not accurately reflect the true level of user involvement in the incel community. This study did not examine posting patterns, such as the frequency of posts over time or the distribution of posts across different threads or subforums. Users with the same number of posts may exhibit varying levels of engagement depending on factors such as the period over which their posts were made, their participation in specific discussions, or their interactions with other users. While this study fails to measure the full engagement of a user the results remain meaningful when evaluated in the limited context studied.

4.2 Future Studies

Further research is needed to understand the formation and maintenance of incel identity salience. Whether prototypical behavior is an acceptable proxy for identity salience and the

identification of core normative behaviors remain open questions and may serve as topics of future research. Different methods for measuring and analyzing language in this context may also be explored.

Future studies could employ more sophisticated measures of forum engagement, such as the analysis of posting patterns, network centrality, or user-user interactions, to provide a more nuanced understanding of user activity and its relationship with social identity salience. Additionally, the examination of different forums and online spaces could provide valuable insight into prototypical behavior patterns within online communities.

This study highlights the need for a more nuanced understanding of incel community dynamics and the role of online platforms in shaping incel identity. While the current analysis focused on prototypical behavior as a proxy for identity salience, future research could explore other aspects of incel identity formation and the ways in which online interactions contribute to the maintenance and reinforcement of incel beliefs. This could involve qualitative analyses of forums (Fabrício et al., 2010) or network analyses of the relationships between users and the spread of ideas within forums. Additionally, research using comparable methods to this study could be conducted with a wider dataset consisting of multiple forums and manosphere communities.

4.3 Conclusion

This study provides new insights into the relationship between online forum participation and incel identity salience. The results suggest a small negative association between prototypical behavior, an approximator of identity salience, and post rank order on *incels.is*. This finding contradicts with the original hypothesis derived from social identity theory and brings forth new questions about identity development and the role of dedicated incel forums in radicalization. The door remains open for further research on incel radicalization processes and identity formulation.

5 References

- Ahad, A., Talpur, M. S., & Jumani, A. (2023). Natural Language Processing Challenges and Issues: A Literature Review. *Gazi University Journal of Science*, 36.
<https://doi.org/10.35378/gujs.1032517>
- Benwell, B., & Stokoe, E. (2007). *Discourse and identity* (Reprinted). Edinburgh Univ. Press.
- Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python: Analyzing text with the natural language toolkit*. O'Reilly Media, Inc.
- Bouchaud, P., Chavalarias, D., & Panahi, M. (2023). Crowdsourced audit of Twitter's recommender systems. *Scientific Reports*, 13, 16815. <https://doi.org/10.1038/s41598-023-43980-4>
- Cohen, K., Johansson, F., Kaati, L., & Mork, J. C. (2014). Detecting Linguistic Markers for Radical Violence in Social Media. *Terrorism and Political Violence*, 26(1), 246–256.
<https://doi.org/10.1080/09546553.2014.849948>
- Dickel, V., & Evolvi, G. (2023). “Victims of feminism”: Exploring networked misogyny and #MeToo in the manosphere. *Feminist Media Studies*, 23(4), 1392–1408.
<https://doi.org/10.1080/14680777.2022.2029925>
- Dixon, L., Li, J., Sorensen, J., Thain, N., & Vasserman, L. (2018). Measuring and Mitigating Unintended Bias in Text Classification. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 67–73. <https://doi.org/10.1145/3278721.3278729>
- Fabricao, B., Azevedo, T., Behar, P., & Reategui, E. (2010). Qualitative Analysis of Discussion Forums. *International Journal of Computer Information Systems and Industrial Management Applications*, 1.
- Fang, X., Che, S., Mao, M., Zhang, H., Zhao, M., & Zhao, X. (2024). Bias of AI-generated content: An examination of news produced by large language models. *Scientific Reports*, 14(1), 5224.
<https://doi.org/10.1038/s41598-024-55686-2>
- Furl, K. (2022). Denigrating Women, Venerating “Chad”: Ingroup and Outgroup Evaluations among Male Supremacists on Reddit. *Social Psychology Quarterly*, 85(3), 279–299.
<https://doi.org/10.1177/01902725221090907>

- Ging, D. (2017). Alphas, Betas, and Incels: Theorizing the Masculinities of the Manosphere. *Men and Masculinities*, 22, 1097184X1770640. <https://doi.org/10.1177/1097184X17706401>
- Habib, H., Srinivasan, P., & Nithyanand, R. (2022). Making a Radical Misogynist: How Online Social Engagement with the Manosphere Influences Traits of Radicalization. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2), 1–28. <https://doi.org/10.1145/3555551>
- Halpin, M. (2022). Weaponized Subordination: How Incels Discredit Themselves to Degrade Women. *Gender & Society*, 36(6), 813–837. <https://doi.org/10.1177/08912432221128545>
- Hanu, L., & Unitary team. (2020). *Detoxify*. <https://github.com/unitaryai/detoxify>
- Hogg, M. A. (2016). Social Identity Theory. In S. McKeown, R. Haji, & N. Ferguson (Eds.), *Understanding Peace and Conflict Through Social Identity Theory* (pp. 3–17). Springer International Publishing. https://doi.org/10.1007/978-3-319-29869-6_1
- Horgan, A., McCarthy, G., & Sweeney, J. (2013). An Evaluation of an Online Peer Support Forum for University Students With Depressive Symptoms. *Archives of Psychiatric Nursing*, 27(2), 84–89. <https://doi.org/10.1016/j.apnu.2012.12.005>
- Hornsey, M. J. (2008). Social Identity Theory and Self-categorization Theory: A Historical Review. *Social and Personality Psychology Compass*, 2(1), 204–222. <https://doi.org/10.1111/j.1751-9004.2007.00066.x>
- Horta Ribeiro, M., Blackburn, J., Bradlyn, B., De Cristofaro, E., Stringhini, G., Long, S., Greenberg, S., & Zannettou, S. (2021). The Evolution of the Manosphere across the Web. *Proceedings of the International AAAI Conference on Web and Social Media*, 15, 196–207. <https://doi.org/10.1609/icwsm.v15i1.18053>
- Jigsaw, & Google. (n.d.). *Perspective API*. Retrieved 19 January 2024, from <https://perspectiveapi.com/>
- Jones, A. (2020). Incels and the Manosphere: Tracking Men’s Movements Online. *Electronic Theses and Dissertations, 2020-*. <https://stars.library.ucf.edu/etd2020/65>
- Lilly, M. (2016). *‘The World is Not a Safe Place for Men’: The Representational Politics of the Manosphere* [Thesis, Université d’Ottawa / University of Ottawa]. <https://doi.org/10.20381/ruor-5184>

- Lindsay, A. (2022). Swallowing the black pill: Involuntary celibates' (Incels) anti-feminism within digital society. *International Journal for Crime, Justice and Social Democracy*, 11(1), 210–224. <https://doi.org/10.3316/informit.379622488850885>
- McKinney, W. (2010). Data Structures for Statistical Computing in Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference* (pp. 56–61). <https://doi.org/10.25080/Majora-92bf1922-00a>
- Meier, M. L., & Sharp, K. (2024). Death to Chad and Stacy: Incels and anti-fandom as group identity. *International Journal of Cultural Studies*, 27(3), 349–367. <https://doi.org/10.1177/13678779231220056>
- Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, 38(11), 39–41. <https://doi.org/10.1145/219717.219748>
- O'Malley, R. L., Holt, K., & Holt, T. J. (2022). An Exploration of the Involuntary Celibate (Incel) Subculture Online. *Journal of Interpersonal Violence*, 37(7–8), NP4981–NP5008. <https://doi.org/10.1177/0886260520959625>
- Pelzer, B., Kaati, L., Cohen, K., & Fernquist, J. (2021). Toxic language in online incel communities. *SN Social Sciences*, 1(8), 213. <https://doi.org/10.1007/s43545-021-00220-8>
- Postmes, T., Spears, R., Sakhel, K., & Groot, D. de. (2001). Social Influence in Computer-Mediated Communication: The Effects of Anonymity on Group Behavior. *Personality and Social Psychology Bulletin*, 27(10), 1243–1254. <https://doi.org/10.1177/01461672012710001>
- Price, H. (2023). Anti-feminism as anti-establishment and emancipatory: The gendered metapolitics of Incel. *Feminist Media Studies*, 0(0), 1–17. <https://doi.org/10.1080/14680777.2023.2291314>
- Richardson, L. (2007). Beautiful soup documentation. *April*.
- Seabold, S., & Perktold, J. (2010). statsmodels: Econometric and statistical modeling with python. *9th Python in Science Conference*.
- Shalaby, R., & Agyapong, V. (2020). Peer Support in Mental Health: Literature Review. *JMIR Mental Health*, 7(6), e15572. <https://doi.org/10.2196/15572>

- Snowswell, A. J., Nelson, L., Xue, H., Salim, F. D., Suzor, N., & Burgess, J. (2023). *Measuring Misogyny in Natural Language Generation: Preliminary Results from a Case Study on two Reddit Communities*. <https://doi.org/10.48550/ARXIV.2312.03330>
- Solea, A. I., & Sugiura, L. (2023). Mainstreaming the Blackpill: Understanding the Incel Community on TikTok. *European Journal on Criminal Policy and Research*, 29(3), 311–336. <https://doi.org/10.1007/s10610-023-09559-5>
- Stijelja, S., & Mishara, B. L. (2023). Characteristics of Incel Forum Users: Social Network Analysis and Chronological Posting Patterns. *Studies in Conflict & Terrorism*, 1–21. <https://doi.org/10.1080/1057610X.2023.2208892>
- Sugiura, L. (2021). *The Incel Rebellion: The Rise of the Manosphere and the Virtual War Against Women*. <https://library.oapen.org/handle/20.500.12657/51536>
- The pandas development team. (2024). *pandas-dev/pandas: Pandas (v2.2.2)* [Computer software]. Zenodo. <https://doi.org/10.5281/zenodo.10957263>
- Turner, J. C., & Oakes, P. J. (1986). The significance of the social identity concept for social psychology with reference to individualism, interactionism and social influence. *British Journal of Social Psychology*, 25(3), 237–252. <https://doi.org/10.1111/j.2044-8309.1986.tb00732.x>
- Van Rossum, G., & Drake, F. L. (2009). *Python 3 Reference Manual*. CreateSpace.
- Van Valkenburgh, S. P. (2021). Digesting the Red Pill: Masculinity and Neoliberalism in the Manosphere. *Men and Masculinities*, 24(1), 84–103. <https://doi.org/10.1177/1097184X18816118>
- Wachowski, L., & Wachowski, L. (Directors). (1999). *The Matrix*. Warner Bros.
- Warhurst, C., Van den Broek, D., Hall, R., & Nickson, D. (2009). Lookism: The new frontier of employment discrimination? *Journal of Industrial Relations*, 51(1), 131–136.
- Winkler, I. (2018). Identity Work and Emotions: A Review. *International Journal of Management Reviews*, 20(1), 120–133. <https://doi.org/10.1111/ijmr.12119>