

# BiaSWE: An Expert Annotated Dataset for Misogyny Detection in Swedish

Kätriin Kukk<sup>1,2</sup>, Danila Petrelli<sup>1</sup>, Judit Casademont<sup>1</sup>,  
Eric J. W. Orłowski<sup>3</sup>, Michał Dzieliński<sup>4</sup>, Maria Jacobson<sup>5</sup>

<sup>1</sup>AI Sweden

<sup>2</sup>Linköping University

<sup>3</sup>AI Singapore

<sup>4</sup>Stockholm University

<sup>5</sup>Anti-Discrimination Agency West Sweden

katriin.kukk@liu.se, danila.petrelli@ai.se,  
juditcasademont@gmail.com, ericorłowski@aisingapore.org,  
michal.dzielinski@sbs.su.se, maria.jacobson@adbvast.se

## Abstract

In this study, we introduce the process for creating BiaSWE, an expert-annotated dataset tailored for misogyny detection in the Swedish language. To address the cultural and linguistic specificity of misogyny in Swedish, we collaborated with experts from the social sciences and humanities. Our interdisciplinary team developed a rigorous annotation process, incorporating both domain knowledge and language expertise, to capture the nuances of misogyny in a Swedish context. This methodology ensures that the dataset is not only culturally relevant but also aligned with broader efforts in bias detection for low-resource languages. The dataset, along with the annotation guidelines, is publicly available for further research.

## 1 Introduction

Large Language Models (LLMs) have experienced immense growth over the past years due to being capable of solving diverse tasks that previously required a separate model for each specific task (De Angelis et al., 2023). Despite their apparent benefits, it is known that the characteristics of the dataset used to train a language model play a fundamental role in determining the model’s behavior (Geburu et al., 2021). LLMs are typically trained on large amounts of data from the Internet and thus inevitably reflect the opinions and biases of its users. For example, a 2018 survey showed

that about 85% of English Wikipedia contributors identified as male (Oldach, 2022). As LLMs’ behavior “reflects the Collective Intelligence of Western society”, LLMs can perpetuate and even amplify biases and stereotypes of social minorities (Kotek et al., 2023). The widespread presence of misogyny online is illustrated by a study from 2020 where 65% of women reported knowing another woman that had been the target of online violence (The Economist Intelligence Unit, 2020).

The way to avoid harmful machine learning models is to ensure that the datasets used for training are responsibly curated, involving diverse stakeholders (Delgado et al., 2021). However, dataset creation alone is not sufficient, and additional approaches, such as alignment, play a role in guiding model outputs towards human values. In the context of bias detection, misogyny varies by language and culture (Zeinert et al., 2021). Therefore, we consider creating expert-annotated, language-specific datasets crucial for detecting biases, helping to identify areas where models may risk perpetuating harmful stereotypes or undesirable attitudes.

To address these challenges, we make two key contributions<sup>1</sup>:

1. We present BiaSWE, a small annotated dataset for misogyny detection in Swedish, annotated for hate speech, misogyny, misogyny type categories and severity.
2. We share the creation process of the BiaSWE

<sup>1</sup>Link to the dataset and annotation guidelines:  
<https://huggingface.co/datasets/AI-Sweden-Models/BiaSWE>

dataset and our annotation guidelines. By doing this, we show how an existing experiment can be adapted to the Swedish needs and cultural context.

## 2 Related Work

In recent years, work has been done in the field of dataset creation for bias and hate speech in general, paying great attention to data coming from online sources, especially social media, such as Twitter, Facebook, Reddit, or blogs. This kind of work has been carried out in a multitude of languages, across several cultural contexts, and tends to cover various forms of sexism as it presents in written language. This is the case of Chiril et al. (2020), who present a corpus for detecting sexism in French tweets. Another example is the work of Zeinert et al. (2021), who sample their Bajer dataset from Twitter, Facebook and Reddit posts in Danish. However, research is also carried out with the target of more subtle, less explicit misogyny in mind; this is the case of the Biasly dataset by Sheppard et al. (2024), who gathered their data from scripts from North American movies, in English.

The most common method for data collection among the different existing datasets is using keywords (Chiril et al., 2020; Zeinert et al., 2021; Sheppard et al., 2024). The degree of detail or the number of keywords varies from words that do not necessarily imply misogyny (e.g. “she”) to ambiguous keywords, to keywords that are very highly related to misogyny and sexism (e.g. “#MeToo”).

Most of the existing misogyny detection datasets provide a taxonomy for different categories of misogyny in addition to the binary classification. Many regard the addition of a multi-label classification layer as necessary, given that “binary detection [...] disregards the diversity of sexist content, and fails to provide clear explanations for why something is sexist” (Kirk et al., 2023). There is no clear consensus regarding the types of misogyny to classify the sentences into, or even on the optimal level of detail regarding the categories.

To the best of our knowledge, this work is the first attempt to create resources for misogyny detection for the Swedish language.

## 3 Method

This section provides an overview of our data preparation process, introduces the team of expert

annotators and details the annotation workflow.

### 3.1 Data

As a data source for our dataset, we used the Swedish website Flashback, one of the largest Swedish internet forums since the 1990s. Known for its focus on freedom of speech, the forum hosts discussions on controversial subjects, and its anonymity often leads to misuse (Norlund and Stenbom, 2021).

To ensure the presence of enough misogynistic examples in the final dataset, we decided to use keyword search. Taking into account the cultural and linguistic closeness of Danish and Swedish, our initial list of keywords was based on the work of Zeinert et al. (2021) that used both keywords and hashtags in Danish. We excluded all hashtags but “#MeToo” because of their rarity on Flashback. The keywords were first translated from Danish into Swedish with the help of a Danish speaker. Thereafter, we presented the resulting Swedish keywords to our team of expert annotators, who suggested removing some of the keywords and adding others. Our list consists of 118 key terms including words and phrases in Swedish (e.g., “kvinna”) as well as some terms and names in popular English slang (e.g., “Chad”). The full list of keywords is available in the annotation guidelines (see section 1).

Based on these keywords, we gathered 450 data points, each to be annotated by two or more annotators. We did not want several annotators to have the exact same set of data points, so we used a rotation system that distributed them. Each expert was assigned 210 data points.

### 3.2 Annotation

The team of annotators included researchers and experts from the humanities and social sciences, as well as civil society actors. Four of our experts identify as women and the other three as men. Our experts volunteered to participate in the project amongst a bigger pool of experts in humanities, social sciences, and civil society representatives that have been introduced to the basics of LLMs and AI within a broader interdisciplinary project at AI Sweden<sup>2</sup>. From this point on, we refer to them as ‘experts’, acknowledging their role in both

---

<sup>2</sup>Link to the project page:  
<https://www.ai.se/en/project/interdisciplinary-expert-pool-nlu>

annotation and providing critical insights into the interdisciplinary process.

To facilitate the annotation process, we prepared annotation guidelines (see section 1) by taking inspiration from the work by Sheppard et al. (2024). Their guidelines included a definition of misogyny that was modified together with the team of experts to obtain the following final definition<sup>3</sup>:

Hatred of, dislike of, contempt for, ingrained prejudice, control of or oppression against women as well as going against the idea of feminism. It is a form of sexism and can be either intentional or unintentional. Misogyny can contain different types of opinions and values such as seeing women as inferior, asserting men’s sexual entitlement, objectifying women, accepting violence, celebrating traditional gender roles but also the need to ”protect” women as well as thinking that equality and feminism have gone too far. One of the ways misogyny can be expressed is through language and, in this project, we focus on misogynistic language portrayed in text. Misogyny can be perpetrated by people regardless of their gender.

We also present a taxonomy of five misogyny categories constructed by combining the twelve categories by Sheppard et al. (2024) and the six categories by Zeinert et al. (2021) and modifying these with the help of the team of experts. The guidelines also give detailed instructions on how to carry out the annotation task and provide examples. For annotation, we used an open-source platform called Label Studio. After everyone had annotated up to 50 examples, we held a workshop to discuss examples where the experts had disagreed.

We divided the annotation of each data point into four small tasks, each more fine-grained than the previous. Once a negative answer was given by the expert, the annotation process ended and the downstream tasks did not need to be completed.

**Hate speech** Our experts were asked to perform binary classification of whether the post they were reading contained hate speech or not. This gave the experts the chance to mark any type of hate speech or hateful behavior.

<sup>3</sup>Link to the complete guidelines, as well as the definition in Swedish, can be found on Hugging Face (see Footnote 1).




Figure 1: Is this post misogynistic?

**Misogyny** The second task was the binary classification of misogyny based on the instructions and the definition of misogyny provided in the annotation guidelines.

**Category** Once a post was classified as misogynistic, our experts were requested to choose a category label. Our taxonomy of misogyny consists of the following categories: stereotype, erasure and minimization, violence against women, sexualisation and objectification, anti-feminism and denial of sexualisation. The experts could only choose one category and could not choose a subcategory outside of the ones presented.

**Severity** Experts hold that misogyny exists in a spectrum and it depends on individual perception. To portray this, we asked them to give a score ranging from 1 to 10, where 1 is the least misogynistic. Although one would assume that, for example, a post portraying violence would have a high score, we did not give them any specific guidelines they had to follow to assign these scores and asked them to trust their own judgement.

## 4 Results

This section provides an overview of the annotation results.

**Hate speech** Each of the 450 posts was annotated by two to four experts and in almost two-thirds of the cases all experts agreed on whether hate speech was present. In the 334 cases where all experts agreed, slightly more posts were annotated as containing hate speech compared to not containing hate speech but the difference was marginal. However, there was a large difference




Figure 2: Which category of misogyny does this example belong to?

between the number of posts that were annotated as hate speech by all experts (172) and posts that were considered to be hate speech by at least one expert (288).

**Misogyny** Figure 1 shows the annotation results for the misogyny classification task. A negative label in the previous task is considered to be a negative label in the misogyny classification task as well, in which case, all 450 posts were again annotated by two to four experts. In slightly more than two-thirds of the examples, all experts agreed on the label. However, in this case, there was a larger imbalance between the two possible labels: 207 posts were considered to be non-misogynistic and 139 misogynistic by all experts. The number of posts considered to be misogynistic by at least one expert was however larger at 243 posts.

**Category** There was more disagreement in choosing the category of misogyny. There were between zero and four category annotations for each of the 450 posts in the dataset. Out of the 185 posts in the final dataset with more than one category annotation, in 93 cases all experts chose the same category. Figure 2 gives an overview of all 546 category annotations in the dataset, comparing the number of times each of the five possible categories was chosen.

**Severity** The last annotation task asked the experts to estimate the severity of the misogyny in the post. Similarly to the previous task, 185 posts

had at least two severity annotations and a closer analysis of those revealed that although the experts seldom selected the same rating, in 91% of cases the difference between the minimum and the maximum rating was not larger than 3.

## 5 Discussion and Conclusion

This project’s primary contribution lies in its interdisciplinary approach to misogyny detection in Swedish rather than the dataset itself, which remains small. Collaborating with experts from diverse fields, we developed an annotation process that captures the complexity of misogyny as it manifests in Swedish online discourse. This experiment provides a valuable framework for future studies focused on bias detection in under-resourced languages.

In the context of misogyny detection, defining what constitutes misogynistic language is inherently challenging. Attempting to capture a wider range of potentially harmful expressions risks being too broad, while using a stricter approach might fail to recognize subtler forms of misogyny. The challenge lies in determining who defines misogyny, as cultural, linguistic and societal factors have an influence over the definition, making it a complex decision.

The feedback from the experts highlighted the need for clearer operational definitions and stronger contextual support. Misogyny detection, particularly in a complex environment like Flashback, would benefit from additional discussion on cultural nuances and interdisciplinary perspectives. Additionally, better alignment between academic rigor and practical applicability is crucial to ensuring that interdisciplinary projects like this one fully realize their potential. We also took into account the experts’ perspective in section 6.

In conclusion, while the dataset is limited, the interdisciplinary approach and methodology offer a valuable starting point for future research. Refining the annotation process and expanding the dataset could further improve the effectiveness of misogyny detection tools, especially for lower-resourced languages like Swedish.

## 6 Limitations and Future Work

This project has several limitations that provide avenues for future work.

**Dataset Size and Diversity** The current dataset, while robustly annotated, is relatively small, lim-

iting its capacity to capture the full spectrum of misogynistic expressions within the Swedish on-line discourse. The limited number of examples might not adequately represent less overt forms of misogyny, which are increasingly prevalent and harmful. Future work should focus on expanding the dataset to include a larger variety of sources.

**Keyword Selection Bias** The reliance on pre-defined keywords to scrape forum posts inherently introduces selection bias, primarily focusing on explicit forms of misogyny. This method may overlook subtle or emergent forms of misogynistic language that do not necessarily conform to expected patterns. Future iterations of this project should aim to refine the keyword selection process. Additionally, incorporating machine learning techniques to identify potential posts could reduce bias introduced by keyword dependency.

**Decontextualisation** A key challenge was annotating decontextualised posts, which made it difficult to detect subtle misogyny. Without context, the experts had to rely on isolated phrases, often missing nuances that could clarify intent or severity. Providing more context in future datasets would enhance accuracy.

**Consensus Building** Disagreements among the experts highlighted the subjective nature of misogyny detection and the challenges in classifying complex human behaviors and attitudes. While we utilized workshops to align annotator perspectives, a more systematic approach to handling disagreement could enhance the consistency and reliability of annotations. Future work could include developing detailed guidelines based on the initial rounds of annotation to standardize responses and improve inter-annotator reliability. Implementing an adjudication process where the experts discuss and resolve disagreements before finalizing annotations could also be beneficial.

## Acknowledgments

This work is a result of the “Interdisciplinary Expert Pool for NLU” project funded by Vinnova (Sweden’s innovation agency) under grant 2022-02870.

## Experts Involved

- Annika Raapke, Researcher at Uppsala University, Department of History
- Eric Orłowski, Sociocultural Anthropologist, Research Fellow (AI Governance), AI

Singapore

- Michał Dzieliński, Assistant Professor at Stockholm Business School, International Finance
- Maria Jacobson, Anti-Discrimination Agency West Sweden
- Astrid Carsbrin, Swedish Women’s Lobby
- Cia Bohlin, The Swedish Internet Foundation
- Richard Brattlund, The Swedish Internet Foundation

## Special Thanks

Special thanks to:

- Francisca Hoyer at AI Sweden for making the Interdisciplinary Expert Pool possible from the start
- Magnus Sahlgren at AI Sweden for guidance
- Allison Cohen at MILA AI for Humanity for participation and support during the experiment

BiaSWE’s multi-disciplinary engagement process was, in part, inspired by the Biasly project from Mila - Quebec AI Institute.

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

## References

- Patricia Chiril, Véronique Moriceau, Farah Benamara, Alda Mari, Gloria Origgi, and Marlène Coulomb-Gully. 2020. <https://aclanthology.org/2020.lrec-1.175> An annotated corpus for sexism detection in French tweets. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1397–1403, Marseille, France. European Language Resources Association.
- Luigi De Angelis, Francesco Baglivo, Guglielmo Arzilli, Gaetano Pierpaolo Privitera, Paolo Ferragina, Alberto Eugenio Tozzi, and Caterina Rizzo. 2023. <https://doi.org/10.3389/fpubh.2023.1166120> Chatgpt and the rise of large language models: the new ai-driven infodemic threat in public health. *Frontiers in Public Health*, 11.
- Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder participation in ai: Beyond “add diverse stakeholders and stir”. In *Proceedings of the Human-Centered AI Workshop at NeurIPS 2021*.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021.

<https://doi.org/10.1145/3458723> Datasheets for datasets. *Commun. ACM*, 64(12):86–92.

Hannah Kirk, Wenjie Yin, Bertie Vidgen, and Paul Röttger. 2023. <https://doi.org/10.18653/v1/2023.semeval-1.305> SemEval-2023 task 10: Explainable detection of online sexism. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2193–2210, Toronto, Canada. Association for Computational Linguistics.

Hadas Kotek, Rikker Dockum, and David Sun. 2023. <https://doi.org/10.1145/3582269.3615599> Gender bias and stereotypes in large language models. In *Proceedings of The ACM Collective Intelligence Conference, CI '23*, page 12–24, New York, NY, USA. Association for Computing Machinery.

Tobias Norlund and Agnes Stenbom. 2021. <https://aclanthology.org/2021.nodalida-main.38> Building a Swedish open-domain conversational language model. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 357–366, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.

Laura Oldach. 2022. What’s with wikipedia and women? Retrieved April 19, 2024, from <https://www.asbmb.org/asbmb-today/careers/030822/what-s-with-wikipedia-and-women>.

Brooklyn Sheppard, Anna Richter, Allison Cohen, Elizabeth Smith, Tamara Kneese, Carlyne Pelletier, Ioana Baldini, and Yue Dong. 2024. <https://doi.org/10.18653/v1/2024.findings-acl.24> Biasly: An expert-annotated dataset for subtle misogyny detection and mitigation. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 427–452, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

The Economist Intelligence Unit. 2020. Measuring the prevalence of online violence against women. Infographic. Retrieved October 11, 2024 from <https://example.com/online-gender-gap-infographic>.

Philine Zeinert, Nanna Inie, and Leon Derczynski. 2021. <https://doi.org/10.18653/v1/2021.acl-long.247> 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.