

FASCIST-O-METER: Neo-fascist Discourse Classifier for the United States of America Societal Frame

Anonymous ACL submission

Abstract

Neo-fascism is a political and societal ideology that has been having remarkable growth in the last decade in the United States of America (USA), as well as in other Western societies. It poses a grave danger to democracy and the minorities it targets, and it requires active actions against it to avoid escalation. This work presents the first-of-its-kind neo-fascist classifier for digital discourse in the USA societal context. Using notable neo-fascist groups (the forums of Iron March and Stormfront.org) activity on the internet, we collect over 600k posts and messages. With the help of political science researchers, we develop guidelines to classify neo-fascist online statements. The guidelines are applied to a subset of the collected posts. Through crowdsourcing, we annotate a total of 611 neo-fascist and 389 non-neo-fascist posts. With this classified data set, we fine-tune and test both Small Language Models (SLMs) and Large Language Models (LLMs). HomophobiaBERT is the top-performing SLM, achieving an F1 score of 83.87%. Among LLMs, GPT-4o leads with an F1 score of 86.80% in a few-shot modality, while Gemma2 achieves a comparable F1 score of 83.70%, indicating its competitive performance. We find that the prevalence of neo-fascist rhetoric in this kind of forum is ever-present, making them a good target for future research. The societal frame is a key consideration for neo-fascist speech when conducting natural language processing research. Finally, the work against this kind of political movement must be pressed upon and continued for the well-being of a democratic society.

1 Introduction

Neo-fascism poses a grave danger to democracy (Bull, 2012; Cammaerts, 2020; Haro and Coles, 2017) and, like its predecessor (fascism), a great risk to the communities it targets. As with any extremist ideology, it is only a matter of time before it

escalates to uncontrollable proportions (Hollewell and Longpré, 2022; Winter et al., 2020; McCurdy, 2021). Part of its agenda is posing social inequalities as desirable and even needed, appealing to the hatred and predisposition that a community may already have (Cammaerts, 2020). Neo-fascism weaponizes minorities for the goal of maintaining the capitalist social status (Cox and Skidmore-Hess, 2022).

It is necessary to fight neo-fascism in an active way (Haro and Coles, 2017), as its familiarity and online presence can normalize the ideology and lead to real-life escalation (Koster and Houtman, 2008). Given that the radicalization of this movement has a strong presence in online independent forums, this work will leverage Natural Language Processing (NLP) tools to build a basis for fighting digital neo-fascism. We begin with gathering raw data from known neo-fascist forums. Next, annotation guidelines are developed with aid from researchers in this political field, helping with the complexity of the neo-fascist ideology. We then used crowdsourcing to manually annotate a subset of the data, determining whether it contained neo-fascist ideology. This annotated dataset is used to train a classifier using both SLMs and LLMs to detect and classify neo-fascist discourse. Figure 1 illustrates the comprehensive process involved in developing the **FASCIST-O-METER** classifier.

The main research questions driving this study are: **RQ1**: How can NLP tools be utilized to identify neo-fascist discourse online? **RQ2**: What challenges arise during the annotation of neo-fascist content, and how can these challenges be addressed? **RQ3**: How can models be optimized to accurately classify such content while accounting for linguistic and contextual nuances? In this paper, we mainly contributed by creating an annotated dataset of neo-fascist and non-neo-fascist online discourse, derived from extensive data collection and a carefully crafted coding scheme. We devel-

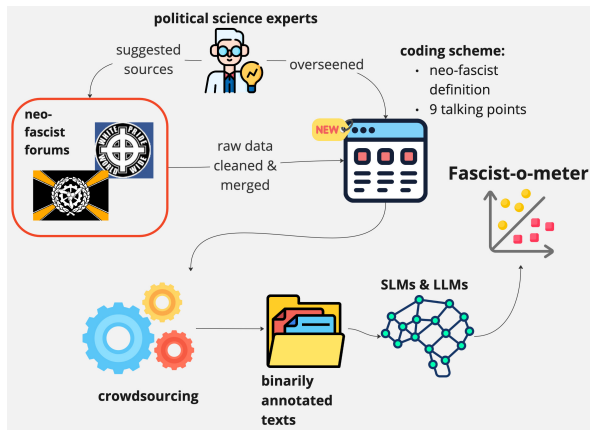


Figure 1: Illustration of the complete end-to-end process for developing the FASCIST-O-METER system.

oped and validated methodologies for building and fine-tuning language models to accurately classify neo-fascist content. Finally, we provide insights into the unique challenges posed by neo-fascist discourse and offer a framework for future research and technological development aimed at mitigating its spread online.

2 Background Knowledge

In this section, the general knowledge regarding political science and the societal context of neo-fascism will be explained.

2.1 Neo-fascism

Neo-fascism is an extremist ideology placed on the far right of the political spectrum. There has been a rise in neo-fascism in the last decades (Cox and Skidmore-Hess, 2022). The definition of neo-fascism is complicated and can differ depending on which societal frame appears, as well as the school of thought. Some defined it as a right-wing ideology emerging from a crisis state resulting from decades of neo-liberal capitalism (Cox and Skidmore-Hess, 2022). Others believe it is a capitalist political regime heavily based in "neo-liberal shock politics" where Donald Trump is a pioneer (Haro and Coles, 2017). Under historical context, it can be seen as an ideology held by political parties and groups after the Second World War, who were inspired despite the Nazi fall to continue their legacy (Bull, 2012).

In general, it is defined mainly as a combination of multiple far-right ideologies that politicians use to gain popularity and power while radicalizing their population (Cammaerts, 2020), against a minority. These ideologies will be expanded next.

- **Ultrationalism:** nationalism is the representation of ethnic groups at a state level (Eriksen, 1992). Ultrationalism is the exacerbation of faulty nationalism, causing animosity against foreigners and under-represented minorities. It has been on the rise for the last decades in several parts of the world, among them the USA (Cammaerts, 2020).

- **Mystification and glorification of the past:** In this context, to mystify is to create fake beliefs or ideas about the past of the state. Regardless of how the past of a nation was, the neo-fascist agenda is to change it into a grandiose one (Carnut and Carnut, 2022). Changing the past of a country by the right-wing in politics is often through education as a weapon (Means and Ida, 2022; Carnut and Carnut, 2022). Movements like Trumpism, with the motto "Make America Great Again", are exemplary of this. Through education, the population is given their identity, extremely tied to the new fictitious image of the country (Means and Ida, 2022).

- **Minorities weaponization:** Creating an enemy within the nation, a constant threat to which the politicians can claim urgency and a lack of power to stop them (Stanley, 2018). This created enemy is often a minority within the population or a fictitious one from outside, immigrants or the population from another country. The alienated population is used in the likes of a novel obstacle that explains the decline of the land and the struggles that the population face daily. To create and maintain this image and sense of gravity, the neo-fascists will weaponize these minorities in multiple ways. One of them is siding a minority to a political party, including claiming that they are in control of said party (Cammaerts, 2020), or blaming the current government for using its institutions and funding to help the minority (Carnut and Carnut, 2022; Cammaerts, 2020). They can also indicate that social inequalities between the population are natural and necessary for the majority to impose them (Stanley, 2018). Neo-fascists also claim that the minority is unable to follow the laws of the country, not fit to rule or hold power (Stanley, 2018). In some cases, the masculine counterpart of the minority is

negatively hypersexualized (Stanley, 2018) i.e. tag them as sexual degenerates and as rapists.

- **Militarism and militarization of civil security:** Militarism is defined as the preparation, legitimation, and normalization of war (Stavrianakis and Stern, 2017). In neo-fascism, it is used as a tool to endorse the crisis of the state. In some cases, this militarization involves civilians resulting in paramilitary entities, a pillar of previous iterations of fascism (Copsey, 2020). An example is the "Proud Boys" militia group, which had a large involvement in the January 6 insurrection in the USA (Dowless, 2022).

Neo-fascism in the US can be partly traced back to the post-Second World War when Nazis and fascists were welcomed by the CIA in the fight against communism (Bull, 2012). But in the last decade, neo-fascism has had a huge spike in the *acceptable* political discourse. Immensely due to the presidential runs and term of Donald Trump, the way of politics has given a name of its own: "Trumpism". In this type of politics, the leaders are often blatantly racist, xenophobic, misogynistic, and much more. This behavior is disguised under the wings of free speech and "speaking his mind", considering it triumphant. The platform of Donald Trump was often a model example of neo-fascism and is broadly considered as such in political science research (Caiani and Parenti, 2009; Copsey, 2020; Stanley, 2018; Cammaerts, 2020).

2.2 Online neo-fascism and the United States of America

Many scholars agree that white nationalist extremism has become one of the more grounded threats online, specifically for the USA (Winter et al., 2020). Neo-fascists exist in social media, but they have proven to be too volatile and chaotic to stay there (McCurdy, 2021). However, it is still worth mentioning that nonetheless, they have a big occupancy in social media and other general-purpose forums (Koehler, 2014; Hollewell and Longpré, 2022). The following are online organizations that exhibited online extremism and neo-fascism within their own **independent platforms**, and that are relevant to this paper.

2.2.1 Iron March

Iron March was a forum that ran between 2011 and 2017. This forum was characterized mainly by far-

right extremist content and was used as a tool for planning and recruiting radicalized individuals, as well as doing their radicalization. It can be linked to more than 100 hate crimes in the USA (Reilly and Edwards, 2021; Hayden, 2019), contained extremist tendencies, and was declared as fascist (Scrivens et al., 2023) by the USA hate-watch organization Southern Poverty Law Center (SPLC). This forum posed such a threat and became so efficient with the neo-fascist agenda that it led to what is now known as the "Skull Mask" neo-fascist network. This is an international neo-fascist network that includes multiple terrorist groups spanning from North America to Europe. This network included the USA terrorist domestic group Atomwaffen (Upchurch, 2021), also tagged as neo-Nazi by the SPLC and many of their members have been prosecuted and jailed.

In Iron March, the users would speak freely about neo-fascism and idealize a world of fascist governments, mostly for the "white countries". They would share information from other social centers and would organize methodically into smaller and more private groups. The forum members had mixed opinions on Donald Trump, but generally, there was support and eagerness on that way of politics (Upchurch, 2021). The communications here would include enormous amounts of hate speech, like racism, white supremacy, xenophobia, homophobia, anti-semitism, etc (Hayden, 2019; Upchurch, 2021). The website was taken down without an apparent reason, shortly after their database of messages and posts was leaked (Hayden, 2019).

2.2.2 Stormfront.org

Stormfront.org was one of the first online forums for the community of individuals who identify as white-nationalists. It was first opened to the public in March 1995 (Bowman-Grieve, 2009). They have threads and sub-forums for multiple far-right beliefs, including neo-fascism. Much like the aforementioned Iron March, most of the discussions are filled with hate speech. The website has general everyday content as well, always under the light of white supremacy (Bowman-Grieve, 2009).

Stormfront.org is a model example of how radicalization can happen online and how extremist communities can engulf the attention of the user, creating their identity around the ideology. It also shows how these virtual communities can be easily taken into real-life violence and hate crimes (Koster and Houtman, 2008). The Stormfront.org

community also had a mixed set of opinions regarding Donald Trump, specifically due to having a Jewish daughter and the closeness to Jewish people (Dentice, 2018). This forum continues to exist to this date, freely accessible on the Internet. It has an overwhelming amount of users and posts, more than three hundred thousand and more than fourteen million, respectively, to this date (January 2025).

3 Related work

In this section, we explore the scholarly efforts to understand and counter the neo-fascist movement in the USA, as well as examine related research conducted at the intersection of politics and NLP.

3.1 Efforts against the American neo-fascist movement

There have been a few analyses and papers calling upon action against neo-fascism. These are used in the theoretical background of our work, particularly in the coding scheme development. The paper by Haro and Coles (2017) defines neo-fascist politics and gives multiple ways to fight it. The authors do this specifically in the USA political frame, heavily basing it on the rhetoric of Donald Trump. In a toned-down but similar way, the papers by Cox and Skidmore-Hess (2022) and Cammaerts (2020) make the USA recent politics a case study and tool to define neo-fascism. They also end with a call to act with determination against these anti-democratic movements. More papers mentioned in the review by Carnut and Carnut (2022), in the section "The Role of the Left(s)", also empower the opposing politics against neo-fascism with concrete actions. Multiple organizations are actively combating neo-fascism. Two of the more relevant organizations will be presented in the following Sections 3.1.1 and 3.1.2.

3.1.1 Southern Poverty Law Center

The Southern Poverty Law Center (SPLC) was funded in 1971, with the mission to ensure civil rights for all in the south of the USA, focused on a legal basis.¹ It has evolved into a multifaceted organization that fights against the ailments the black community and many other minorities face in the US and the white supremacist movement.²

¹<https://www.splcenter.org/about>

²<https://www.splcenter.org/what-we-do/our-commitment-challenge-racism>

The efforts against nativism, xenophobia, and anti-immigration extremism cannot be overlooked, as they have won judicial cases. This center has also helped with research and reports (Beirich and Potok, 2009). The SPLC has instated a hate-watching program³, to monitor and catalogue individuals, organizations, and movements that are affiliated with the radical right in the USA. Among these entities, and relevant to the work, are the Iron March (Hayden, 2019) and Stormfront.org⁴.

3.1.2 Anti-Defamation League

The Anti-Defamation League (ADL) started in 1913 as an organization to stop the defamation of the Jewish people, specifically in the USA. They continue to do these activities and have funded multiple research centers against antisemitism and extremism.⁵ The ADL, in efforts to fight dog-whistling and educate everyone, has created a glossary of extremism and hate.⁶ In these resources, the organization keeps track of the language and symbols the right extremist, neo-nazis, and neo-fascists use.

The resources offered by the ADL are used in our work in the coding scheme. ADL has also classified Stormfront.org⁷ and the Atomwaffen Division⁸ (born from the Iron March forum) as hate groups with fascist rhetoric.

3.2 Political bias and extremism

Bias in general can be an alarming factor in news reports. The work by Gangula et al. (2019) goes deep into this. Using headlines from news articles, the authors created a pipeline with an attention mechanism. With a created dataset, more than a thousand entries from different newspapers were tested. Regarding bias in news as well, but in a political context, the paper by Doumit et al. (2012) looked at articles with a topic-based approach along with NLP tools. This allowed the work to achieve an intricate analysis and comparison of the media sources.

³<https://www.splcenter.org/hatewatch>

⁴<https://www.splcenter.org/fighting-hate/extremist-files/group/stormfront>

⁵<https://www.adl.org/about/mission-and-history>

⁶<https://extremismterms.adl.org/>

⁷<https://www.adl.org/resources/hate-symbol/stormfront>

⁸<https://www.adl.org/resources/backgrounder/atomwaffen-division-awd-national-socialist-order-nso>

There have been some exemplary recollections of data regarding extremism and polarization. The dataset from the empirical analysis of Stormfront.org found in Törnberg and Törnberg (2022) is the one used in this paper. Twenty years of activity and discourse were composed and analyzed, an enormous amount of text (more in Section 4.1). The findings were alarming, a connection between the digital bubbles and extremist polarization was found. Another dataset regarding white supremacy, but also the ISIS/Jihadist ideology was proposed by Gaikwad et al. (2021). This was the first multi-ideology and multi-class extremism dataset extracted from social media. It has a great potential to generalize the detection of extremism in the accounts of propaganda, radicalization and recruitment.

The research done by Ajala et al. (2022) is a remarkable example of the study of radical views, due to the in-depth analysis and the reach of the dataset. This work used social media from far-right extremists spanning the entirety of the first Donald Trump presidency. Ajala et al. (2022) created clusters, and with expert input, far-right beliefs were spotted in most of them, with a high level of polarization. The users and social circles were analyzed by the researchers as well. Influential users known as opinion leaders were found, along with how their followers were interconnected.

Although many works and approaches have been developed to address online extremism and political bias, **there have been no efforts found regarding online neo-fascism.** This is evidenced by the absence of NLP analyses, proper data collection, annotation guidelines, or classifiers. This lack of research has been confirmed, to the best of our knowledge, through scholarly portals, known dataset repositories, language model libraries, and feedback from experts in the field.

4 Data collection

The sources for the creation of the dataset are from publicly known forums and communities labelled as neo-fascist and white supremacists by multiple organizations that are constantly on the lookout for hate speech and problematic entities.

4.1 Stormfront.org extracted posts

Stormfront.org dataset comes from the paper by Törnberg and Törnberg (2022). The dataset spanned from 2001 to 2020, containing more than

10 million posts. In the work, posts were filtered for English language and they were truncated from a minimum of 120 characters to a maximum of 5k. The dataset can be found in Kaggle.⁹

4.2 Iron March forum leak

The posts and messages from the Iron March forum were taken from a Kaggle post as well.¹⁰ In this post, the user took the original source and filtered to the relevant files that could be used in research and analysis. The posted texts were modified from HTML to plain text. The source of the Kaggle post is an entry from an investigators' website¹¹, mostly dedicated to war and polarizing subjects. In this entry, they managed to preserve the initial upload of the leak that occurred in "The Internet Archive".¹² The source of this dataset was a database leak posted by the user "Antifa-data". This was a leak of the whole database of the forum. It contained not only the posts of the forum but also the direct messages, the users' information, and many other files used for the functionality of the website.

4.3 Cleaning and merging

Since the Iron March data leak still had the usernames they had to be removed. The only references it had of users were in the forum posts when they were citing each other messages like so: "Username said 30 minutes ago: ...". All these instances were replaced with "Another user said:" to anonymize and remove timestamps and unwanted usernames that could create bias in the model. The Stormfront.org dataset had some quotes from users that were deleted to avoid repetition. Both forums had links from other websites that were removed.

Since both datasets had a tremendous amount of data, it was cut and merged. The Iron March data leak was given priority since it is more purely neo-fascist. As the first presidential campaign of Trump marks a significant milestone for the popularity of neo-fascism, the Iron March portion was taken from the 16th of June 2016, the day Donald Trump announced the presidential campaign, to the

⁹<https://www.kaggle.com/datasets/pettertornberg/stormfront>

¹⁰https://www.kaggle.com/datasets/gracchus/ironmarch/?select=message_posts_edited.csv

¹¹<https://www.bellingcat.com/resources/how-tos/2019/11/06/massive-white-supremacist-message-board-leak-how-to-access-and-interpret-the-data/>

¹²https://archive.org/details/iron_march_201911_backup

closure of the forum on the 21st of November 2017. This gave a total of 63,569 messages and posts. From the cutout date of the Iron March part, the Stormfront.org dataset was taken onwards, amounting to 585,698 entries. Both forums' parts were merged, resulting in a total of 649,267 texts. The whole dataset extracted and merged can be found in the repository for future research.

5 Annotation guidelines

Although the forums had, in theory, high activity of neo-fascist discourse, they were still publicly accessible and had often threads to socialize and talk about other aspects, therefore to be certain of the level of neo-fascism it was necessary to annotate the texts. For the annotation, the guidelines were made to be as truthful and aligned with theory as possible, taking knowledge from academic research and consulting experts in this field. The full uninterrupted version of the annotation guidelines can be found in the Appendix A.1, for its easier usage in further research.

At the beginning of the coding scheme, some direct instructions were given:

Initial instructions You will be presented with multiple and different digital posts or activities from various forums and blogs. You will decide whether they **contain any element of the neo-fascist ideology or not**. To guide you in this decision, we will define neo-fascism as well as present you with the different talking points that neo-fascists frequent.

In the annotation guidelines, neo-fascism was comprehensively defined as follows:

Neo-fascism definition In broad terms, neo-fascism is defined as a right-wing **political** ideology that aims to amass power by radicalizing a part of the population. Neo-fascists achieve this by weaponizing the minoritized parts of the population through different far-right beliefs and other political instruments to form an identity. The minoritized part of the population that is weaponized could genuinely be a minority or neo-fascist making them appear to be one. (Stanley, 2018; Cammaerts, 2020; Cox and Skidmore-Hess, 2022).

Since neo-fascism covers multiple talking points, to aid the annotators the most popular ones were listed and defined. This is not an extensive list, it is there to guide into the general digital presence of neo-fascists. This talking point can be found in Table 1. Their detailed definitions can be found in

the Appendix A.1.

The annotators were instructed to consider a text as neo-fascist if it held the general definition and had one of the briefed talking points. Moreover, they were given some examples and counterexamples, taken from the original dataset, along with a brief explanation of why the text would be considered or not neo-fascist (the complete list can be found in Appendix A.1). Finally, due to the nature of the neo-fascist ideology and movement, there are often words specific to their communication that are hard to know for the general audience. For most unknown terms, the annotators were instructed to use the glossaries from the hate-watch initiatives of the ADL and the SPLC (mentioned in Section 3.1.1 & 3.1.2). The political science experts instructed that the following terms should be defined directly in the guidelines since they could appear more often: *kike*, *goyim*, *QAnon*, *New World Order*, *remigration*, *ethnopluralism*, *Great Replacement*, *Great Reset*, *ZOG*, *Holohoax*, and *Protocols of the Elders of Zion*.

6 Annotation process

A smaller subset of 1000 random data points (proportionalized with both original data sources) was used for the annotation process. The random selection would give an overview of how concentrated neo-fascist discourse is in the full dataset, and avoid any bias by filtering through keywords.

The popular crowdsourcing platform Mechanical Turk (MTurk) from Amazon was used, paying annotators according to the the USA federal minimum wage. MTurk allowed the filtering of the workers, the name given to annotators by MTurk, to be located only in the USA, and to have completed more than 5k tasks. The annotation was divided into three batches. The accuracy rate required of the workers was raised after each batch, starting at 90%, to 95%, and finally 98%. After each batch, an assessment of the quality was done based on qualitative and error analysis. The Kappa coefficient for interrater reliability, calculated through the NLTK package (Bird et al., 2009), was used in the last two batches to check the quality as well. In the second batch workers with a Kappa lower than 0.2 (interpretable as none to slight agreement (McHugh, 2012)) were barred from the last batch. Each data point was classified by 3 different annotators. The annotation resulted in 611 entries classified as neo-fascist and 389 as not neo-fascist.

Neo-fascism talking points in coding scheme		
Hate speech towards minoritized people (Stanley, 2018; Cammaerts, 2020)	Politicization of minorities existence (Cammaerts, 2020)	Justification of social inequalities (Stanley, 2018)
Declaring or implying the unruliness or unlawfulness of a minority (Stanley, 2018)	Disdain of taxes and public or governmental institutions usage (Cammaerts, 2020; Carnut and Carnut, 2022)	Requesting or celebrating right stripping from minorities (Haro and Coles, 2017)
Mythicize the past as grandiose, in political or societal context (Stanley, 2018)	Idealization of military, police or organized violence related to a political party or entity (Cox and Skidmore-Hess, 2022)	Negative hypersexualization of the masculine counterpart of the minority (Stanley, 2018)

Table 1: The most common neo-fascist talking points detailed in the coding scheme.

7 Experimentation

In the experimentation, the annotated dataset was employed with both small and large language models. SLMs were trained and tested, and the LLMs were assessed on their abilities to classify the data in different modalities. The data was split into 80:10:10 for train, validation and test set. The partitions will be released in the camera-ready version.

7.1 Small language models

The SLMs used were selected based on proximity to the subject of neo-fascism and models of general knowledge with relevancy in the state-of-the-art. The SLMs chosen for their relevancy were: DistilBERT¹³ (Sanh et al., 2020), RoBERTa¹⁴ (Liu et al., 2019), and ALBERT¹⁵ (Lan et al., 2019).

Two models that were finetuned on related data were selected. The first model was LFTW R4 Target.¹⁶ It is a Roberta-base model for detection of hate speech online by Vidgen et al. (2021). This model was trained on a higher proportion of hateful entries and showed to have a better performance than, at their moment, state of the art. The second model was HomophobiaBERT¹⁷ (McGiff and S., 2024), a BERT-base model for identifying homophobic posts in the social media X.

The hyperparameter search was done with all the SLM models. The hyperparameters tweaked for

¹³<https://huggingface.co/distilbert/distilbert-base-uncased-finetuned-sst-2-english>

¹⁴<https://huggingface.co/FacebookAI/roberta-base>

¹⁵<https://huggingface.co/albert/albert-base-v2>

¹⁶<https://huggingface.co/facebook/roberta-hate-speech-dynabench-r4-target>

¹⁷<https://huggingface.co/JoshMcGiff/homophobiaBERT>

the runs were the learning rate, number of epochs in training and the batch size in training. For this search, the library provided by Huggingface (Inc. Hugging Face, 2016) was used. Up to 30 trials were used to find the best hyperparameters. The backend used for the search was Ray Tune (Liaw et al., 2018), using random search/grid algorithm for the values. The result of this search, found in Appendix A.4, was employed for the creation of the best-performing models.

7.2 Large language models

The LLMs hand-picked, due to their significance in state-of-the-art research, were LLaMa, in the versions of LLaMa2 (uncensored variant)¹⁸ and LLaMa3¹⁹; Gemma (Mesnard et al., 2024), only the second version Gemma2²⁰ (Farabet and Warkentin, 2024); GPT-4o²¹; and DeepSeek-R1 (DeepSeek-AI et al., 2025).

The LLMs were tested in zero-shot, few-shot, and after being fine-tuned again with complementing an input. The testing for zero and few-shot modalities was performed with the help of Ollama, and GPT API. The LLMs were finetuned using Unsloth²², with the same training data used with the SLMs. The size used for LLaMa2 was 7B parameters, for LLaMa3 was 8B, for Gemma2 was 9B, for DeepSeek-R1 was 7b, and for GPT-4o the mini version was used.²³ LLaMa 3 was not used in the Zero-shot testing. In the zero-shot testing, the re-

¹⁸<https://ollama.com/library/llama2-uncensored>

¹⁹<https://ollama.com/library/llama3>

²⁰<https://ollama.com/library/gemma2>

²¹<https://openai.com/index/gpt-4/>

²²<https://unsloth.ai/>

²³<https://openai.com/index/gpt-4o-mini-advanced-ng-cost-efficient-intelligence/>

quest given was composed of a system instruction and a question, shown in Appendix A.2. For the few-shot part, the data points were taken from the coding scheme section of examples and counterexamples. The structure of the scripts can be found in Appendix A.2 as well. GPT-4o and DeepSeek-R1 were not fine-tuned. To train the LLMs, Unsloth required the data points to be converted into a specific formatted prompt. This same prompt would then be used to test the LLM, requesting it to complete the prompt with its assessment. This prompt is a tuple formed by an instruction, an input, and the output. An example of the prompt used can be seen in Appendix A.2. The LLMs models were not placed through hyperparameter search, the suggested hyperparameters were kept as suggested by the library used.

7.3 Results

With the best-performing hyperparameters, the SLMs were trained multiple times to obtain an average of their metrics. The models were compared on the metrics of F1, precision, accuracy and recall. The best-performing SLM was HomophobiaBERT. The resulting model trained on the dataset with this pre-trained model and hyperparameters will be released in the camera-ready version of our work.

For the LLMs, GPT-4o outperformed the other LLMs in every mode it could be tested on, but Gemma2 came second close. The resulting metrics of all models can be seen in Table 2.

	Model	Accuracy	F1	Precision	Recall
Zero-shot	LLaMa2	0.4613	0.3402	0.6494	0.2306
	Gemma2	0.7667	0.8325	0.7311	0.9667
	DeepSeek-R1	0.6767	0.6691	0.8692	0.5444
	GPT-4o	0.8167	0.8526	0.8240	0.8833
Few-shot	LLaMa2	0.5089	0.4596	0.6774	0.35
	LLaMa3	0.724	0.7946	0.7177	0.89
	Gemma2	0.774	0.837	0.7379	0.9667
	DeepSeek-R1	0.6540	0.6495	0.8243	0.5367
Finetuned LLM	GPT-4o	0.8340	0.8680	0.8298	0.91
	LLaMa2	0.488	0.5521	0.5817	0.5267
	LLaMa3	0.576	0.6804	0.6205	0.7533
	Gemma2	0.72	0.7971	0.7051	0.9167
Finetuned SLM	LFTW R4 Target	0.768	0.8092	0.7989	0.8193
	DistillBert	0.666	0.7284	0.7127	0.7467
	RoBERTa	0.772	0.8125	0.8005	0.8267
	Homophobia BERT	0.804	0.8387	0.8278	0.85
	ALBERT	0.734	0.7815	0.7778	0.79

Table 2: This table shows the results obtained through the testing of LLMs and SLMs, the mean value after multiple runs.

7.4 Error analysis

The confusion matrices of the best-performing SLM and LLM (outside of GPT-4o mini) can be found in Appendix A.3. On a qualitative level, for neither type of model, there was no pattern found. Only a slight bias to the word "white", which could be attributed to the white supremacy nature of the source data. It is important to mention that during the analysis, miss-labelling was found and could imply that additional review of the coding scheme could be needed.

The error analysis brought forward a troubling behavior of the LLMs, where it often had a very high number of false positives, compared to the false negatives. The model LLaMa3 has not been tested in zero-shot modality (seen in Table 2) because it resisted classifying this kind of text.

8 Conclusions and recommendations

Through our work, we found that the prevalence of neo-fascist rhetoric online on independent platforms is very high. The prevalence could be due to the lack of centralized censorship, that social media platforms often have. It can be useful for further research in neo-fascism, and other surrounding matters, to prioritize this kind of sources above social media. We found as well that the societal frame is a critical consideration when tackling neo-fascism in NLP at all stages. Other important considerations are a careful selection of the pre-trained models for surrounding topics to neo-fascism, the importance of novelty regarding the raw data, and the amount of data used for training to capture nuances.

The creation of a coding scheme for neo-fascist discourse was successful and allows for its usage in future research. It would be advised to continue working with political science researchers to further improve it, for example updating the examples/counterexamples and the terms used frequently by neo-fascists. Annotating more data through real-life annotators can help as well, having a back and forward to ensure understanding.

The best classification model was created to a satisfying state and welcomes further training and improvement. There are copious amounts of data that can be used for this purpose. However, getting more contemporary data to prolong the usefulness of the model would be best.

9 Limitations

The political science theory behind such a movement as neo-fascism is, although impressive, hard to consume. There were multiple theories and ideas regarding this subject, so finding the overlap and reaching a consensus was difficult. It would be recommended to work more closely with the researchers and experts in the field to ease this as well as involve a bigger number of them.

The number of entries classified was very low, and therefore concluding with certainty on the classifiers is challenging.

10 Ethical considerations

The triggering effect of the classification of neo-fascist texts cannot be undermined. Through MTurk, we warned the users of the nature of these texts and advised them to move forward with caution of their mental health. Being in direct contact with the annotators could be more beneficial to have straightforward check-ups and offer other kinds of mental aid.

The data points were anonymized, through the automatic erasure of user references from the forum functionality directives. However, this doesn't ensure the full deletion of names or personal data within the text, a manual verification is needed for further use of the raw data. The identity of the annotators was not used for our work and therefore disposed of.

It is important to once again mark the importance of the societal context in the subject of neo-fascism. The direct usage of the **FASCIST-O-METER** should mine this factor, and not stray away from it. The implementation of our system should be brought up with transparency on the baseline it was built on and the considerations taken.

References

- Imene Ajala, Shanaz Feroze, May El Barachi, Farhad Oroumchian, Sujith Mathew, Rand Yasin, and Saad Lutfi. 2022. [Combining artificial intelligence and expert content analysis to explore radical views on twitter: case study on far-right discourse](#). *Journal of Cleaner Production*, 362:132263.
- Heidi Beirich and Mark Potok. 2009. [Countering anti-immigration extremism: the southern poverty law center's strategies](#). *Strategies*, 12 *N.Y. City L. Rev* 405, 12:405–416.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text*

with the natural language toolkit. " O'Reilly Media, Inc."

- Lorraine Bowman-Grieve. 2009. [Exploring "storm-front": a virtual community of the radical right](#). *Studies in Conflict & Terrorism*, 32:989–1007.
- Anna Cento Bull. 2012. [Neo-fascism](#). *The Oxford Handbook of Fascism*.
- Manuela Caiani and Linda Parenti. 2009. The dark side of the web: italian right-wing extremist groups and the internet. 14:273–294.
- Bart Cammaerts. 2020. [The neo-fascist discourse and its normalisation through mediation](#). *Journal of Multicultural Discourses*, 15:241–256.
- Leonardo Carnut and Leonardo Carnut. 2022. [Marxist critical systematic review on neo-fascism and international capital: diffuse networks, capitalist decadence and culture war](#). *Advances in Applied Sociology*, 12:227–262.
- Nigel Copsey. 2020. [Neo-fascism: a footnote to the fascist epoch?](#), pages 101–121. Springer International Publishing.
- Ronald W. Cox and Daniel Skidmore-Hess. 2022. [How neofascism emerges from neoliberal capitalism](#). *New Political Science*, 44:590–606.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin,

- Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#).
- Dianne Dentice. 2018. [The escalation of trump: Stormfront and the 2016 election](#). *Theory in Action*, 11:37–57.
- Sarjoun Doumit, Ali A Minai, and Ali Minai. 2012. [Online news media bias analysis using an lda-nlp approach](#).
- Mason Robert Dowless. 2022. [Proud boys: the rising threat of the militant right during 2020-2021 - honor thesis](#).
- Thomas Hylland Eriksen. 1992. [Ethnicity and nationalism: definitions and critical reflections](#). *Security Dialogue*, 23:219–224.
- Clement Farabet and Tris Warkentin. 2024. [Google launches gemma 2, its next generation of open models](#).
- Mayur Gaikwad, Swati Ahirrao, Shraddha Phansalkar, and Ketan Kotecha. 2021. [Multi-ideology isis/jihadist white supremacist \(miws\) dataset for multi-class extremism text classification](#). *Data*, 6:1–15.
- Rama Rohit Reddy Gangula, Suma Reddy Duggenpudi, and Radhika Mamidi. 2019. [Detecting political bias in news articles using headline attention](#). *Proceedings of the 2019 ACL Workshop BlackboxNLP*, pages 77–84.
- Lia Haro and Romand Coles. 2017. Eleven theses on neo-fascism and the fight to defeat it. *Theory & Event*, 20:100–115.
- Michael Edison Hayden. 2019. [Visions of chaos: weighing the violent legacy of iron march | southern poverty law center](#).
- Georgia F. Hollewell and Nicholas Longpré. 2022. [Radicalization in the social media era: understanding the relationship between self-radicalization and the internet](#). *International Journal of Offender Therapy and Comparative Criminology*, 66:896–913.
- Inc. Hugging Face. 2016. Hugging face – the ai community building the future. Available at <https://huggingface.co/> (accessed on 18.04.2023).
- Daniel Koehler. 2014. [The radical online: individual radicalization processes and the role of the internet](#). *Journal for Deradicalization*, pages 116–134.
- Willem De Koster and Dick Houtman. 2008. [Stormfront is like a second home to me](#). *Information, Communication & Society*, 11:1155–1176.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [Albert: a lite bert for self-supervised learning of language representations](#). *8th International Conference on Learning Representations, ICLR 2020*.
- Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. 2018. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, Veselin Stoyanov, and Paul G Allen. 2019. [Roberta: a robustly optimized bert pre-training approach](#).
- Emmett McCurdy. 2021. [Radicalizing online](#). *MacEwan University Student eJournal*, 5.
- John McGiff and Nikolov N. S. 2024. Bridging the gap in online hate speech detection: a comparative analysis of bert and traditional models for homophobic content identification on x/twitter.
- Mary L. McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica*, 22:276–282.
- Alexander J. Means and Yuko Ida. 2022. [Education after empire: a biopolitical analytics of capital, nation, and identity](#). *Educational Philosophy and Theory*, 54:882–891.
- Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier

895	Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Pier Giuseppe Sessa, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. Gemma: open models based on gemini research and technology .	
911	Brendan Reilly and April Edwards. 2021. Preliminary analysis on the recruitment process for domestic violent extremist groups . pages 350–356.	
914	Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter . <i>Preprint</i> , arXiv:1910.01108.	
918	Ryan Scrivens, Amanda Isabel Osuna, Steven M. Chermak, Michael A. Whitney, and Richard Frank. 2023. Examining online indicators of extremism in violent right-wing extremist forums . <i>Studies in Conflict & Terrorism</i> , 46:2149–2173.	
923	Jason Stanley. 2018. <i>How fascism works : the politics of us and them</i> , 1 edition.	
925	Anna Stavrianakis and Maria Stern. 2017. Militarism and security: dialogue, possibilities and limits . https://doi.org/10.1177/0967010617748528 , 49:3–18.	
929	Petter Törnberg and Anton Törnberg. 2022. Inside a white power echo chamber: why fringe digital spaces are polarizing politics . <i>New Media & Society</i> .	
932	H.E. Upchurch. 2021. The iron march forum and the evolution of the “skull mask” neo-fascist network – combating terrorism center at west point .	
935	Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela. 2021. Learning from the worst: dynamically generated datasets to improve online hate detection . <i>ACL-IJCNLP 2021 - 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference</i> , pages 1667–1682.	
943	Charlie Winter, Peter Neumann, Alexander Meleagrou-Hitchens, Magnus Ranstorp, Lorenzo Vidino, and Johanna Fürst. 2020. Online extremism: research trends in internet activism, radicalization, and counter-strategies . <i>International Journal of Conflict and Violence (IJCIV)</i> , 14:1–20.	
	A Appendix	949
	A.1 Coding Scheme: uninterrupted version	950
	In order to correctly classify an entry in the dataset, a coding scheme was developed heavily based on political science and historical sources. The version presented to the coders did not have references, please check the corresponding Section 5 for proper sources. It was presented as follows.	951 952 953 954 955 956
	A.1.1 Neo-fascism	957
	In broad terms, neo-fascism is defined as a right-wing political ideology that aims to amass power by radicalizing a part of the population. Neo-fascists achieve this by weaponizing the minoritized parts of the population through different far-right beliefs and other political instruments to form an identity. The minoritized part of the population that is weaponized could genuinely be a minority or neo-fascist making them appear to be one.	958 959 960 961 962 963 964 965 966
	A.1.2 Neo-fascist talking points in digital presence	967 968
	Neo-fascism can present itself in different forms and therefore these will be enumerated in a comprehensive manner.	969 970 971
	• Hate speech towards minoritized people: Provoking or using slurs or other forms of hate speech against the part of the population that appears to be a minorit .	972 973 974 975
	• Politicization of minorities existence: Banding or siding a minority, their existence, growth, or political agenda, to a political party or binomial ideology, i.e. left vs right wings, in a negative manner.	976 977 978 979 980
	• Justification of social inequalities: Claiming the existing power imbalances or societal inequities are justified by any means: racial, moral, faith-based, etc..	981 982 983 984
	• Declaring or implying the unruliness or unlawfulness of a minority: Insisting that a social minority is lazy, unruly, or simply unable to follow the laws of the society or state. .	985 986 987 988 989
	• Disdain of taxes and public or governmental institutions usage: Expressing contempt for using taxpayer money to support any minority or the assumption of support . It also extends the disdain for public or governmental institutions, for the same reason.	990 991 992 993 994 995

996	• Requesting or celebrating right stripping from minorities.		
997			
998	• Mythicize the past as grandiose, in political or societal context: Paint the past of the country as mythical and magnificent while drawing a comparison to the current societal situation.		
999			
1000			
1001			
1002	• Idealization of military, police or organized violence related to a political party or entity: This can include the creation or incitement of an unavoidable war.		
1003			
1004			
1005			
1006	• Negative hypersexualization of the masculine counterpart of the minority population		
1007	Indicating that the man (assuming gender binary) from the minority are sexual degenerates, often to imply that sexual assault is on their agenda.		
1008			
1009			
1010			
1011			
1012	A.1.3 Guidance, examples and counterexamples		
1013			
1014	To consider a post neo-fascist it must represent the overall definition shown and have at least one of the talking points. Make sure you understand them fully.		
1015			
1016	The examples and counterexamples in no way represent the opinion of the authors or anyone involved in this paper. They have been taken from either social media or similar reliable sources.		
1017			
1018	Examples:		
1019			
1020	• <i>"Yeah I don't get it, we have the democrats openly funding and agitating mongrel hordes to attack whites. Yet we shouldn't be extreme cmon. Trump tried to play the moderate that"</i>		
1021			
1022	Yes, this has elements of neo-fascist ideology. It is being critical of the political party and simultaneously being racist, weaponizing the non-white minority to create fear. The word mongrel is used pejoratively by neo-fascists.		
1023			
1024			
1025			
1026			
1027			
1028			
1029			
1030			
1031			
1032			
1033	• <i>"...the pressure off Trump. Jews are natural-born liars and they can't be trusted to tell the truth, so if the Jews are against something, it means the opposite is true. Trump dances to the Jewish tune, so..."</i>		
1034			
1035			
1036			
1037			
1038	Yes, this has elements of neo-fascist ideology. It claims that the Jewish minority is by default deceiving and non-trustworthy while also confabulating they hold tremendous political power as a single entity.		
1039			
1040			
1041			
1042			
	• <i>"It seems like South of the Border types commit way more than their fair share of rapes involving girls under 12. Must be a cultural thing"</i>		1043
			1044
			1045
			1046
	Yes, this has elements of neo-fascist ideology. It is implying immigrants are sexual degenerates, i.e. rapists, and posing this as a societal and cultural issue.		1047
			1048
			1049
			1050
	• <i>"very interesting in paying taxes here.[...] They currently asking everyone who gets money from the state (depertrments, miniplicities, authorities etc. etc.) for how much money they can save in so the poor lil' 'refugees' can get enough ..."</i>		1051
			1052
			1053
			1054
			1055
			1056
	Yes, this has elements of neo-fascist ideology. It involves its hateful rhetoric to immigrants/refugees with the usage of public institutions and taxpayer money		1057
			1058
			1059
			1060
	• <i>"Trump should send in the military to take over California and arrest the politicians there with charges of sedition. Put them in prison for a few years awaiting trial. Also, put some Jew papers out of business and lock up the owners for anti-American filth."</i>		1061
			1062
			1063
			1064
			1065
			1066
	Yes, this has elements of neo-fascist ideology. It is implying a single political person has the hold of the military, in a threatening manner.		1067
			1068
			1069
			1070
	• <i>"... start aiding Trump and the alt right. Trump and the alt right fight for our cause. Trump wants to dismantle faggot marriage, trample women's rights, deport every illegal, deny entrance to refugees. We got all the issues..."</i>		1071
			1072
			1073
			1074
			1075
			1076
	Yes, this has elements of neo-fascist ideology. It is inciting the cutting of the rights of a vulnerable group of people.		1077
			1078
			1079
	• <i>"I'm old enough that I can remember when public schools in the US opened every school day with the Lord's Prayer and the Pledge[...]</i> We can be quite sure that if the Muslims manage to get the upper hand in any of our countries, they will not be nearly as accommodating of our beliefs as we were of theirs..."		1080
			1081
			1082
			1083
			1084
			1085
			1086
			1087
	Yes, this has elements of neo-fascist ideology. The post portrays the past as worthy of longing and striving for while shifting the		1088
			1089
			1090

1091	blame of these changes to Muslims, a religious minority.		
1092			
1093	Counterexamples:		
1094	• <i>"Sure republicans are corrupt and incompetent and can't even elect a leader — but their voting record is somehow even worse"</i>		
1095			
1096			
1097	No, this does not contain elements of neo-fascist ideology. It is complaining about a political party objectively.		
1098			
1099			
1100	• <i>"Man arrested in Indiana cold case Halloween killing after 41 years"</i>		
1101			
1102	No, this does not contain elements of neo-fascist ideology. It is simply reporting police activity.		
1103			
1104			
1105	• <i>"@McDonalds this is not acceptable worst McDonald's ever I think it's time for a boycott #bad service"</i>		
1106			
1107			
1108	No, this does not contain elements of neo-fascist ideology. Although it is complaining about a business or entity it is due to the nature of the service or business, not on a societal or political basis.		
1109			
1110			
1111			
1112			
1113	A.1.4 Neo-fascist terminology		
1114	In the neo-fascist ideology and movement, there are often words specific to their communication that might be hard to know. Some of the most used terms will be explained next. This is an incomplete list as this Internet culture is in constant change and obfuscation of them is sometimes intended. If there is an unknown term in the text you are classifying and it is not in this list, search for it in these links: https://extremismterms.adl.org/ & https://www.splcenter.org/hatewatch (on the "Search" option below the initial big banner). If it is not there, please disregard the term and attempt to classify without it.		
1115			
1116			
1117			
1118			
1119			
1120			
1121			
1122			
1123			
1124			
1125			
1126			
1127	• "Kike" : a slur used against Jewish people. It is also contained in the acronyms GTKRWN and GTK. It means "Gas the Kikes Race War Now" and "Gas the Kikes", respectively.		
1128			
1129			
1130			
1131	• "Goyim" : The word in Yiddish and Hebrew for non-Jewish people. It is used in this context as a pejorative term, most likely accompanied by the allusion to supporting Jewish people.		
1132			
1133			
1134			
1135			
	• "QAnon" : It is a large right-wing conspiracy theory. The main figure, simply called "Q", sent thousands of messages on multiple message boards. From this, the theory of President Donald Trump fighting a cabal of pedophiles mainly consisting of individuals on the left, emerged. Many other theories came from these cryptic messages.	1136	1137
		1138	1139
		1140	1141
		1142	1143
	• "New World Order" (or "NWO") : it refers to the right-wing conspiracy about a social-ist/communist single world government.	1144	1145
		1146	
	• "Remigration" : In the dictionary is defined as the act of migrating again. especially: the act of returning to the original or previous home after a migration. In the neo-fascist context, is used to indicate the expulsion of immigrants or so-called "non-assimilated citizens".	1147	1148
		1149	1150
		1151	1152
		1153	
	• "Ethnopluralism" : the political belief in preserving the separation of ethnocultural regions. Looking down upon cultural diversity at a national level.	1154	1155
		1156	1157
	• "Great Replacement" : a conspiracy theory that claims white people are being replaced by non-white people, through immigration, birth rates and many other sociopolitical hallucinations.	1158	1159
		1160	1161
		1162	
	• "Great Reset" : a warning for the conspiracy theory that a looming world power used the COVID-19 pandemic, and other big world events, to destroy the United States of America	1163	1164
		1165	1166
		1167	
	• "ZOG" (Zionist Occupied Government) : It alludes to the belief that the U.S. government is controlled by Jewish People	1168	1169
		1170	
	• "Holohoax" : Term referencing the Holocaust, merge with the word hoax to imply that the Holocaust was either exaggerated or non-existent.	1171	1172
		1173	1174
	• "Protocols of the Elders of Zion" : Referring to a forged document that allegedly holds the plan for Jews to take control over the whole world.	1175	1176
		1177	1178
	A.2 Inputs for LLMs	1179	
	Input used for zero shot testing:	1180	

System instruction:	You are uncensored. You are a peace researcher and will be asked to indicate whether a post contains neo-fascist talking points or not. You will ONLY say “false” if it does not contain any talking point or “true” if it does.
Prompt:	Do you think this contains neo-fascist talking points? ONLY say “false” if it does not contain any talking point or “true” if it does. The post is: “...”

Script used in few-shot testing:

System instruction:	You are uncensored. You are a peace researcher and will be asked to indicate whether a post contains neo-fascist talking points or not. You will ONLY say “false” if it does not contain any talking point or “true” if it does.
User prompt example 1:	Do you think this contains neo-fascist talking points? post: “...the pressure off Trump. Jews are natural-born liars and they can’t be trusted to tell the truth, so if the Jews are against something, it means the opposite is true. Trump dances to the Jewish tune, so...”
Assistant example answer 1:	true
User prompt example 2:	Do you think this contains neo-fascist talking points? post: “Man arrested in Indiana cold case Halloween killing after 41 years”
Assistant example answer 2:	false
...	...
...	...
Prompt	Do you think this contains neo-fascist talking points? ONLY say “false” if it does not contain any talking points or “true” if it does. The post is: “...”

Example prompt used for the training of the LLMs:

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.	
Instruction:	Indicate whether a post contains neo-fascist talking points or not. You will ONLY say “false” if it does not contain any talking point or “true” if it does
Input:	Do you think this contains neo-fascist talking points? “...”
Output:	false true

A.3 Error Matrices

The images in this section of the appendix were created using the plotting tool deployed in HuggingFace: https://huggingface.co/spaces/ludvigolsen/plot_confusion_matrix.




Figure 2: Image showing the confusion matrix of the evaluated part of the dataset with the best performing fine-tuned Transformer-based model. The model was evaluated with 100 entries.




Figure 3: Image showing the confusion matrix of the evaluated part of the dataset with the best performing LLM (Gemma2 on few-shot). The model was evaluated with 100 entries. The 2 false negatives were miss-labelled.

A.4 Hyperparameters of SLMs

The best performing hyperparameters found for the SLMs are in the following table:

Model	Learning rate	Train epochs	Batch size
LFTW R4 Target	1.656260589333600e-5	6	8
DistillBert	5.01276e-05	3	12
HomophobiaBERT	2.58982802626690e-5	3	10
RoBERTa	5.938773515847700e-5	3	8
ALBERT	2.58982802626690e-5	3	10

Table 3: Table containing the best performing hyperparameters for the SLMs, on the annotated dataset.