

Online to Offline Crossover of White Supremacist Propaganda

Ahmad Diab
University of Pittsburgh
Pittsburgh, USA
ahd23@pitt.edu

Bolor-Erdene Jagdagdorj
Carnegie Mellon University
Pittsburgh, USA
bjagdagd@andrew.cmu.edu

Lynnette Hui Xian Ng
Carnegie Mellon University
Pittsburgh, USA
lynnetteng@cmu.edu

Yu-Ru Lin
University of Pittsburgh
Pittsburgh, USA
yurulin@pitt.edu

Michael Miller Yoder
Carnegie Mellon University
Pittsburgh, USA
yoder@cs.cmu.edu

ABSTRACT

White supremacist extremist groups are a significant domestic terror threat in many Western nations. These groups harness the Internet to spread their ideology via online platforms: blogs, chat rooms, forums, and social media, which can inspire violence offline. In this work, we study the persistence and reach of white supremacist propaganda in both online and offline environments. We also study patterns in narratives that crossover from online to offline environments, or vice versa. From a geospatial analysis, we find that offline propaganda is geographically widespread in the United States, with a slight tendency toward Northeastern states. Propaganda that spreads the farthest and lasts the longest has a patriotic framing and is short, memorable, and repeatable. Through text comparison methods, we illustrate that online propaganda typically leads the appearance of the same propaganda in offline flyers, banners, and graffiti. We hope that this study sheds light on the characteristics of persistent white supremacist narratives both online and offline.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**; • **Social and professional topics** → *User characteristics*; • **Information systems** → *Social networking sites*.

KEYWORDS

white supremacist, hate speech, propaganda, geospatial, temporal, narratives

ACM Reference Format:

Ahmad Diab, Bolor-Erdene Jagdagdorj, Lynnette Hui Xian Ng, Yu-Ru Lin, and Michael Miller Yoder. 2023. Online to Offline Crossover of White Supremacist Propaganda. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*, April 30-May 4, 2023, Austin, TX, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3543873.3587569>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WWW '23 Companion, April 30-May 4, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9419-2/23/04...\$15.00
<https://doi.org/10.1145/3543873.3587569>

1 INTRODUCTION

In 2021, top officials from the US Department of Justice and Department of Homeland Security named white supremacist extremism as the most significant domestic terror threat [55]. This ideology has inspired extremists to commit mass murders in Pittsburgh [45], El Paso [48], and Buffalo [54] in recent years. Far-right groups with white supremacist and fascist elements have recently led or planned violent attacks against governments in the United States, Brazil, and Germany. To spread their ideology and build communities of hate, white supremacist groups disseminate propaganda using memorable slogans like “end all immigration” in both the online and offline spaces [33]. With this backdrop, it is increasingly important to study the spread of hate by white supremacist groups in the online medium, and how it relates to offline propaganda.

Online hate speech is a frequent topic of attention in the cyber social threat space [2, 13, 22]. Some quantitative and computational work has focused specifically on white supremacist ideology. Alatawi et al. [1] develop methods for detecting white supremacist hate speech, while Eddington [18] draw network connections between white supremacy hate groups within the United States and United Kingdom in 2016. However, fewer studies examine the connections between offline events and online hate. Lupu et al. [37] observed that offline trigger events, such as protests and events, often leads to increase in online hate speech; Hirvonen [29] studied the affinity between online white supremacy hate speech and offline hostility; and Keum et al. [33] identified that the constant dissemination of ideology by white supremacist groups online can impact the mental health of their surrounding community. Given these potential dangers, it is important to characterize and study the patterns of messaging that draws people into white supremacism and investigate how this messaging connects to offline violence [23].

This paper aims to improve our understanding of the correlation between online and offline white supremacist messaging. Using an overlay of offline and online sources, we identify matching white supremacist propaganda and investigate the following research questions:

RQ1: What spatial and temporal patterns characterize the spread of offline white supremacist propaganda? Which types of propaganda spread the furthest and last the longest?

Through a spatial and geographic analysis of events reported as containing offline white supremacist propaganda (flyers, graffiti,

banners, etc), we identify when the same propaganda is used across events and analyze the reach of each propaganda phrase.

RQ2: What temporal patterns characterize the online-to-offline crossover of white supremacist propaganda?

We directly correlate online and offline white supremacist “quotes”, which are phrases representing ideologies and propaganda promoted by individuals and extracted from online posts and offline event descriptions. We do so via text similarity matches, serving as indicators to identify the persistence of the quotes in both spaces. Similar methods to identify persistence of text across multiple platforms were conducted in studies of the spread of news [35] and disinformation [40]. Using information on quote persistence, we characterize the geographical, time-series and narrative distributions of quotes that persist in both online and offline environments.

1.1 Contributions

To our knowledge, this study is the first of its kind to combine the analysis of white supremacist propaganda across both offline and online environments, studying a US-wide reach of white supremacist ideology. Using a combination of geographical, temporal and text analysis, we make the following contributions:

- (1) We measure the persistence of offline white supremacist propaganda messages through text analysis techniques and quantitatively estimate the spatiotemporal distribution of particular white supremacist slogans.
- (2) We analyze the prevailing common narrative themes between online and offline white supremacist propaganda.
- (3) We find preliminary evidence within 2 years of overlapping data that white supremacist propaganda quotes begin online before appearing in offline events, supporting the hypothesis that online activity can translate into offline activity.

2 RELATED WORK

2.1 White Supremacist Extremism

Merriam-Webster defines white supremacy as “the belief that the white race is inherently superior to other races and that white people should have control over people of other races”¹. Social movements organized under this principle are recognized as white supremacist extremism. Key beliefs of this ideology include essentialist, “natural” race and gender hierarchies with white men at the top [10]. White supremacists believe that white male power is threatened in today’s world and that violent action is necessary to protect the white race [3, 21]. White supremacist extremism has been demonstrated in premeditated violent attacks against people of other races, which provides white supremacists a shared sense of accomplishment, solidarity and power [57]. Larger scale incidents include the Tree of Life synagogue attack in Pittsburgh, Pennsylvania, where the shooter believed Western Civilization was facing “extinction” and that Jewish refugees were “invaders”; the 2017 “Unite the Right” rally in Charlottesville, North Carolina, where an

anti-racist protester was killed with a car [9]; anti-Muslim shootings in Christchurch, New Zealand; and the anti-Latinx shooting in El Paso, Texas inspired by the Christchurch event.

Though these movements are regarded as fringe, scholars such as sociologist Jessie Daniels [17] do not view them as isolated from structural white supremacy, a system in which white people control the material resources and major institutions of societies [4, 12]. These social movements exploit bigotries widely held in societies with structural white supremacy [8, 20, 46]. Ideologies that are implicit in structural white supremacy are explicitly stated in white supremacist social movements. Studying white supremacist communication is thus important not just to understand a fringe social movement, but also as a window into broader societies organized under white supremacy.

Offline white supremacist propaganda is a small piece of the larger communication infrastructure of white supremacist extremism. To better understand white supremacist extremism, we must also look online. White supremacist groups have recognized the potential of the Internet in spreading their message. Donald Black, the founder of white supremacist forum Stormfront, encouraged his followers to use social networking sites to “reach new people and bring them here” [43]. Similarly, David Duke, former Ku Klux Klan leader, mentioned in 2007 that the Internet could give “millions access to the truth [...] of racial enlightenment” [17].

White supremacist groups have made extensive use of social media and online forums for information provision, seeking donations, recruitment and networking [58]. They have also united around political campaigns, discursively creating and connecting their organizations and lingo to the campaigns [18], and made use of online chat channels to distribute and mobilize extremist movements [27].

Therefore, it is important to trace the presence of the same white supremacist propaganda in both online and offline environments.

2.2 The Intersection between Online and Offline Speech

The spread of online propaganda and offline activism has been observed to be positively related. Online platforms facilitate the spread of white supremacist ideology through community building and the development of shared realities and emotions. Empirical histories suggest that the sowing of discord among individual and online communities can create an impetus for violent offline expressions, such as protests and shootings [26].

The formation of online communities is rarely isolated: online communities signal a social connection in the offline world. Literature provides a mixed view of the connection between the online to offline ecology: a supporting relationship and a negative relationship. Spread of ideologies online can be dismissed as unproductive and inhibit the incidents of offline activism [56]. For example, activists online have been observed to distance themselves from offline riots, contrasting away the online-offline domains [34]. However, online efforts such as the production of video clips can crossover to offline action, where these video clips inspire extremism or teach the ways of extremism (e.g., how to fire a gun) [14]. Others find that both online speech and offline events can reinforce each other in a cycle of radicalization [24, 31].

¹<https://www.merriam-webster.com/dictionary/white%20supremacy>, accessed 24 January 2023

Data Field	Description
quote	white supremacist propaganda text
location	city and state of event
event	event details (i.e. people/organizations involved)
timestamp	time/date of event occurrence

Table 1: Data fields extracted per event for offline dataset

Given the close relationship and the dynamic interplay between the online and offline domains, it is clear that one domain influences the other, and vice versa. As such, it is important to understand the intersection between online and offline domains.

2.3 Computational Propaganda

Computational propaganda is the “use of algorithms, automation and human curation to purposefully distribute misleading information” [59]. The use of computational methods increases the ease and rate of disseminating propaganda information, and provides flexibility for both broad efforts and targeted attempts at distributing propaganda materials. The presence of computational propaganda has been observed most prominently in the political sphere, where automated agents have attempted to influence the UK-EU Brexit referendum [30], the US elections [59] and the Brazil elections [5], to name a few instances.

In recent years, detection of computational propaganda has made significant progress. Several shared tasks involving detection and classification of annotated propaganda datasets have been constructed for the development of detection and analysis techniques within the community [38, 50]. Detection methods involve constructing machine learning models for identification of the use of propaganda techniques [16] and the classification of the types of techniques employed [6] within texts, and the use of network analysis approaches to identify propaganda through the presence of coordinated inauthentic action [15].

Our work focuses on how online white supremacist propaganda of all types also appears in offline propaganda.

3 DATASETS

In this study, we use two datasets that contain white supremacist propaganda: one offline dataset and one online dataset. We elaborate on the details of these datasets in the following section.

3.1 Offline Dataset

We gathered 43,154 events involving white supremacist propaganda from the ADL H.E.A.T. Map². Collected by the Anti-Defamation League, an anti-hate non-profit organization, this map aggregates and plots incidents of hate, extremism, antisemitism and terrorism from a variety of sources, including news reports, government documents and victim reports. Commonly, events include propaganda such as flyers and banners, with a description of the incident and metadata including location, date, group, and ideology of the group. The data fields that are extracted from the events are reflected in Table 1. The data considered in this study spans from January 2008, until July 2022.

²<https://www.adl.org/resources/tools-to-track-hate/heat-map>

Data Field	Description
quote	white supremacist propaganda text
dataset	data source of posts (cf. Table 3)
platform	online platform source (cf. Table 3)
timestamp	time/date of event occurrence

Table 2: Data fields extracted per post in online dataset

Location Annotation. For each offline white supremacist propaganda phrase, or “quote”, we annotated it with latitude/longitude (lat/lon) coordinates by parsing the provided location information (city, state) through the Nominatim API³. The API takes in a (city, state) pair and returns the (lat/lon) coordinates to which the pair is geographically located. Errors in the output due to misspelled cities (e.g. Mattson, IL instead of Matteson, IL) and ambiguous locations (e.g. national park addresses) were manually corrected.

3.2 Online Dataset

The online dataset contains 4,371,453 posts after removing duplicates, it spans from October 2001 through November 2019. The data was collected from explicitly white supremacist domains and organizations, assembled from a variety of datasets and data dumps by [61]. Table 2 lists the data fields extracted per post from the online dataset. Further details about the online dataset are reflected in Table 3.

This dataset includes forums and websites dedicated to white supremacism, including Stormfront, Iron March, and the Daily Stormer. Tweets from organizations that the Southern Poverty Law Center labels as white supremacist hate groups [19, 47] are also included. This dataset filters 4chan /pol/ data, an imageboard known for white supremacy, to posts from users choosing Nazi, Confederate, Fascist, and White Supremacist flags (in the dataset from Papasavva et al. [42]), as well as posts in “general” threads with fascist and white supremacist topics (in the dataset from Jokubauskaitė and Peeters [32]). We also include smaller datasets manually annotated for white supremacist ideology from Rieger et al. [49], Siegel et al. [52], and Alatawi et al. [1].

4 METHODS

4.1 Matching Propaganda Quotes

For this work, we refer to individual propaganda slogans or phrases used by individuals or organizations as “quotes”. Through this paper, we report these quotes within double quotation marks (“ ”). For example, an incident where Patriot Front, a white supremacist group, distributed propaganda with the slogan “America is not for sale” and the incident where an unknown person drew the slogan “America is not for sale” in graffiti are incidents that share the same propaganda, though they differ in nature, location and people involved. Mapping out such connections can aid in identifying the spatial and temporal spread of similar white supremacist propaganda. Identifying quotes offline and locating them online helps establish connections between the same propaganda ideology spreading in both domains. We detail our process for matching quotes in this section.

³<https://nominatim.org/>

Platform	Data source	# Posts	# matched offline quotes
4chan	Papasavva et al. [42]	2,686,267	137,744
4chan	Jokubauskaitė and Peeters [32]	578,650	20,546
4chan, 8chan, Reddit	Rieger et al. [49]	361	65
Stormfront	Stormfront dump	751,980	57,761
Iron March	Iron March dump	179,468	10,559
Twitter	Qian et al. [47]	84,695	2409
Twitter	ELSherief et al. [19]	3,480	261
Twitter	Alatawi et al. [1]	1,098	46
Twitter	Siegel et al. [52]	170	37
Discord	Patriot Front dump	39,577	675
Daily Stormer	Calderón et al. [11]	26,099	5058

Table 3: Details of online dataset.

We start with the set Q of all offline quote occurrences, extracted from events in the offline dataset. For each pair of quote occurrences $q_1, q_2 \in Q$, we perform two steps of matching comparison to determine whether they are similar. The first step simply compares if the texts of the two quotes (lowercased and tokenized) are exactly the same. Should the first step return negative, the quotes are compared in a second step. The second step further pre-processes both quotes by removing punctuation and English stopwords. The pre-processed quotes are then compared with each other again; if they match exactly, they are considered as matches.

This matching flow was performed in three iterations: (1) offline-offline quotes, to identify matching quote clusters in the offline dataset; (2) offline-online, to identify quotes that had crossed over and appeared in both medium; and (3) online-online quotes, to identify matching quote clusters in the online dataset. In each iteration, we performed an all-pairs matching flow, meaning every quote was compared to all other quotes in the dataset. After each matching iteration, the generated unique set of quotes were manually inspected and combined if necessary.

In the offline dataset, we identified 1798 unique white supremacist propaganda quotes. Not all reported quotes are what would generally be regarded as propaganda, however. Some quotes are very short (e.g. “aryan”) or non-ideological (e.g. “send a message”). In order to filter these out, two of the authors annotated quotes for being ideological, not just descriptive. Specifically, annotators marked whether each phrase portrays a white supremacist ideological position or evaluation (such as “race traitors”) instead of being descriptive and likely also used by those without a white supremacist worldview (such as the phrase “white nationalists”). Annotators discussed difficult cases and came to a consensus. 1655 out of 1798 unique quotes (92.0%) were annotated as propaganda.

For the online dataset, we found 188,861 online posts containing exact matches of offline quotes. Searching for near matches (removing punctuation and stopwords) provided an additional 46,701 posts, yielding a total of 235,562 online posts that mention either the exact or near matches of propaganda quotes in offline events.

4.2 Spatial Coverage of Offline Propaganda Quotes

To calculate the spatial coverage of the offline quotes, we borrow the concept of a “radius of gyration” from physical science. The

radius of gyration, R_g , of a propaganda quote represents the typical distance travelled by the quote, weighted by number of occurrences in different locations. In our calculation, we weight each quote instance equally (the “mass” of each quote is 1 in the original physical science terms). We start with each unique quote cluster $C \in Q$, the result of our matching process described in 4.1. For each set of N matching propaganda quotes within a quote cluster $C = \{q_1, q_2, \dots\}$, we calculate the location centroid q_c of C as the average latitude/longitude of the quotes in C . R_g is the root mean square Euclidean distance of q_i to the centroid q_c . Formally, R_g is represented as Equation 1 and illustrated in Figure 1.

$$R_g = \sqrt{\frac{\sum_{i=1}^N (q_i - q_c)^2}{N}}, \quad (1)$$

where $(q_i - q_c)$ = Euclidean distance of q_i to q_c ,

q_c = centroid of quotes,

N = num matching quotes

Equation 1: Calculation of Radius of Gyration




Figure 1: Illustration of parameters used in calculating the radius of gyration for $N = 3$ matching quotes $q_1, q_2,$ and q_3 . The blue circle is the geographical cover (gyration) of the quote.

4.3 Temporal Coverage

We analyze at what points in time offline events and the same propaganda quotes occur in online posts. In this manner, we can observe the temporal trends per medium and analyze the temporal trends across each other.

We first calculate how long the quotes appear in reported offline events across the US, simply the number of days from the first occurrence to the final occurrence.

To get an idea of how long white supremacist propaganda quotes circulate in their first environment before crossing over to the second environment, for each quote cluster *C* of matching quotes, we examine the temporal relationship between first appearances online and offline. We calculate the average time difference in terms of days between its first appearance in the online dataset and its first appearance in the offline dataset, or vice versa.

5 RESULTS

5.1 Geographic Distribution of Offline Propaganda Quotes

Figure 2 depicts the geographical distribution of white supremacist propaganda quotes as extracted from the offline dataset. The blue circles are the centers of gyration for each quote cluster, sized proportionally by the number of occurrences of the quote. This propaganda is spread widely across the United States. To test how the distribution of propaganda compares with the overall population distribution of the United States, we first calculate the overall center of propaganda as the average latitude and longitude across all centers of quotes, weighted by quote frequency. If propaganda is evenly distributed across the US population, this center would match the mean center of population, the mean latitude and longitude of all population in the United States. We find that the center of propaganda is further to the northeast than this mean center of population, which suggests that propaganda has a slight tendency toward the northeast section of the US.

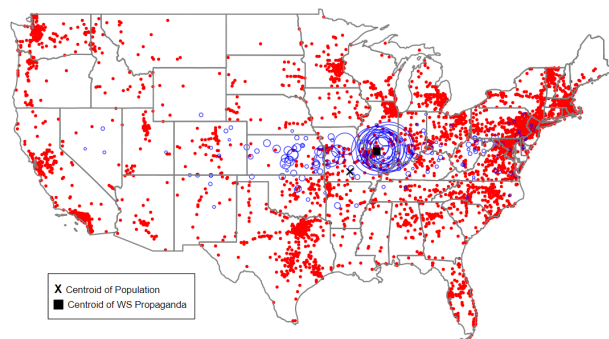


Figure 2: Geographic distribution of propaganda quotes from the offline dataset. Each red dot is an incident. Each blue circle is centered around the average geographical locations of a quote, its size is proportional to the quote’s frequency.

5.2 Spatiotemporal Distribution of Quotes

The spatiotemporal distribution of the offline propaganda quotes is visualized in Figure 3. In Region A, quotes have both high spatial and temporal coverage. These quotes, such as “America First”, “America is not for sale” are repeated extremely often and over a wide spread of area. Region B are quotes that cover a wide spread of geography but are not long-lasting (i.e., “March against sharia”). Region C depicts quotes that are long-lasting but do not cover a large area (i.e., “Defending our heritage”).

We observe that propaganda messages with high spatiotemporal coverage (i.e., Region A at the top right corner of the graph) are quotable, memorable, and could apply in many contexts. These properties of successful propaganda are in line with the findings about the discourse of propaganda in the Persian Gulf War [41]. We find that the most commonly repeated propaganda are also appeals to patriotism. An example of such a message is:

Message: “One nation, against invasion”

Event description: Patriot Front, a white supremacist group, distributed propanganda flyers at Northwest Arkansas Community College that read “America is not for sale”, “Reclaim America”, and “One nation, against invasion”

Frequency of quote: 942

Life span: 820 days




Figure 3: Spatiotemporal distribution of quotes, with the most popular quotes of each region listed to the right.

Looking at patterns in other regions of the spatiotemporal distribution, we observe much propaganda is time and location specific.

Some propaganda quotes are associated with the locations that they are encompassed in. Some quotes are associated with a small radius, possibly because the narrative theme within the message is fairly targeted towards the geographical community. For example, “which way western man” has the smallest radius of 2.6, centered in Philadelphia, Pennsylvania. This is a catchphrase that is misquoted from the 1978 book by William Simon and is mostly used by the European Heritage Association. William E. Simon, a businessman and philanthropist and the Secretary of the Treasury in 1974, was born in the city of Philadelphia. The use of the catchphrase around his birthplace could be a deliberate attempt to appeal in a local area. The quote appeared in 47 online posts and 26 offline events.

Another such message is “Justice for Cannon Hinnant”, which has a small radius of gyration covering North Carolina to Virginia.

The shooting of five year old Cannon Hinnant took place in August 2020 in Wilson, North Carolina [51]. Hinnant was a white American who was shot at point-blank by his Black neighbor, Darius Sessoms. The limited radius of the propaganda quote could be an attempt to arouse white supremacist feelings around the area of the incident by appealing to the commonality of geography. This specific quote appeared offline 40 times and was observed in 9 online posts.

Some narratives appear at a specific time, in tandem with world events. The narrative “stop coronavirus, deport illegal aliens, close borders, stop immigration” was first recorded on 21 March 2020 until 17 May 2020, coinciding with the beginning of the worldwide coronavirus pandemic. Another narrative rooted in conspiracy theories is “every single aspect of the covid agenda is jewish”. This narrative arose on 18 September 2021 until 31 March 2022 as part of an effort to link anti-government sentiment (“the COVID agenda”) with antisemitic extremist views [25]. This national and local targeting warrants further investigation into white supremacist communication tactics.

Figure 4 illustrates the temporal distribution of white supremacist quotes online and offline. This is for the 2-year period from 2018 through 2019, chosen for the presence of significant activity in both online and offline datasets. The online dataset does not contain data past 2019, and there are relatively few reports in the offline dataset before 2018. Overall, we find that white supremacist propaganda online leads appearances offline. 87% of white supremacist quotes first appear online before appearing offline. The average time difference between the appearance of quotes across both medium are 242 ± 212 days. That is, quotes circulate online for an average of 7 months before making it to the offline space. Examples of such quotes are: “white power world wide”, “troops to the border”, “build the wall deport them all”, “our race is our nation”, “only two genders”.

13% of the quotes appear offline first before making their appearance online. Examples of such quotes are: “protect your heritage”, “hate speech is free speech”, “wake up white America”, “Reclaim America”, “the holocaust was a good thing”. Figure 5 plots the first appearances of each quote in offline and online spaces. These quotes also appear within less than two months (49 ± 158 days) online after their first offline occurrence, a much shorter period for crossover than those that appear online first.

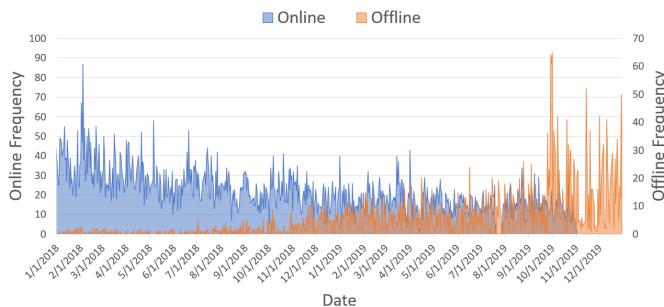


Figure 4: Frequencies of propaganda quote appearances of offline and online.




Figure 5: Visualization of the first appearances of white supremacist propaganda quotes: online or offline. The size of each bubble indicates the number of unique quotes appearing on the same day. It can be observed that the vast majority of quotes first appear online

5.3 Comparison of Online and Offline Propaganda Use

We perform a comparison of propaganda use in online and offline platforms through identifying popular quotes in both environments. Figure 6 compares the most popular quotes in both online and offline platforms. We observe that the most popular quotes, measured by the frequency of quote occurrences, differ between online and offline environments. In fact, propaganda that is popular online is not popular offline, and vice versa.

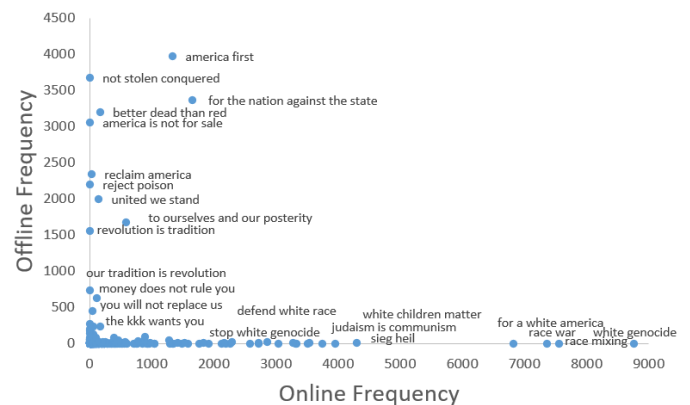


Figure 6: Frequency of most popular quotes in offline and online domains.

Propaganda popular in offline spaces are more vague and applicable in many contexts. They are more likely to give plausible deniability of racism and some seem to act as dogwhistles [28]. Propaganda popular online are more directly ideological, naming minoritized groups such as Jews and specifying that they are for a “white America” and “white children”.

6 DISCUSSION

In this study, we examined the spatiotemporal prevalence of white supremacist propaganda in both online and offline spaces.

We find that short, decontextualized, and easily repeatable phrases like “America First” have the largest spatial and temporal coverage offline. They often appeal to patriotism and are more coded than

direct. Though white supremacist propaganda is prevalent throughout the US, we identify a trend toward the northeast area. There could be many reasons for this. This area could have a historical reputation for liberalism, and so radical right-wing hate groups may view this as a “battleground” where propaganda is needed in a way that is not in more conservative areas. One of the most prominent groups behind events in our dataset, contributing 84% of the offline propaganda quotes, Patriot Front, a vastly spread hate group across the United States and is known to be active in eastern cities such as Boston and Philadelphia [44]. Some of their slogans call for the formation of a white ethno-state and restoring America to the past, which may explain a focus on this region of the country.

In the online space, we examined posts from a wide variety of data sources, ranging from online chat logs to forums to blogs to social media platforms. This broadens our analysis beyond the idiosyncrasies of any particular platform. The presence of white supremacist propaganda across all these different genres of online platforms highlights the severity and spread of white supremacist ideology.

We observe white supremacist propaganda crossing from online environments to offline events. This is in line with other online-to-offline studies, where online activism facilitates offline protests [26]. This points to the importance of monitoring narratives online from hate groups.

A large majority (87%) of the quotes appear online first before appearing offline, during which they circulate for an average of seven months. This is a fairly long period of time, during which narratives can coalesce around certain messages before they are publicly distributed as offline propaganda. Quotes that appear offline before we observe them in our online dataset cross over much faster, less than two months on average. This could be due to the ease of quickly picking up and discussing narratives in offline propaganda in online spaces (e.g. blogs, social media).

Our observations further show that propaganda that is popular in the online space do not necessarily equate to popularity in the offline space. This indicates that both environments have separate information consumption spheres, with different communicative strategies from groups with the same ideological objectives. Such a case has been observed in the US Capitol Riots event, an event with significant white supremacist propaganda, where there were different disinformation narratives propagated across YouTube and website videos [39].

6.1 Limitations and Future Work

This study of online-to-offline crossover is not without limitations. Though the ADL draws from a team of expert annotators to curate its offline hate group event dataset, from which we extract our offline propaganda dataset, the annotation methodology is largely unknown. The coding scheme and the inter-annotator agreement could affect the categorization of events. Our online dataset does not contain posts past 2019, which limits our ability to compare with offline events. Collecting newer online white supremacist posts (perhaps from sources such as Telegram) would allow a more recent comparison. In both temporal and spatial analysis, more stable results could be found by better-handling outliers in the dataset.

We match propaganda quotes with literal exact matches and matches without punctuation and stopwords. These methods could be expanded to include other forms of near matches, such as topic modeling, segmentation [35], n -gram overlap [53], fuzzy search [7], paraphrasing [60], or rhetorical function matching [36]. These more sophisticated matching approaches may yield quotes that are similar in semantics or rhetoric, but do not have exact word overlap.

In addition, the current analysis is limited to the United States and to English propaganda. White supremacist ideology circulates worldwide, and future work should expand this analysis to other countries and contexts where messaging may vary.

An analysis of the network structure of those engaging with white supremacist propaganda offline and online may give insight to how this content is spread, and through what groups narratives cross from online to offline spaces.

7 CONCLUSION

In this work, we conduct analysis of the relationship between white supremacist propaganda in online and offline spaces. Understanding when and where white supremacist hate messages spread is critical for an effective and knowledgeable response by authorities, organizations, companies and academics.

We combine data from a hand-annotated dataset that records incidents of offline white supremacist propaganda spread with a curated dataset of white supremacist quotes from online blogs, forums and social media. We begin by analyzing the geographic spread of offline propaganda, identifying that incidents of white supremacist activity are generally concentrated around regions of historical significance in the United States. Through temporal analysis comparison of quotes from both online and offline environments, we find preliminary evidence that online white supremacy propaganda generally appears online before being distributed offline. Through text comparison methods, we observe that the most popular propaganda in online and offline environments differ, but narratives seem to be time and location-specific. We hope that this study illuminates patterns in strategic communication used by white supremacist hate groups.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support from the AFOSR awards and the Collaboratory Against Hate Research and Action Center (CAH). We thank the ADL Center on Extremism for curating and providing the offline dataset and for inspiration for this project. Any opinions, findings, and conclusions or recommendations expressed in this material do not necessarily reflect the views of the funding sources.

REFERENCES

- [1] Hind S. Alatawi, Areej M. Alhothali, and Kawthar M. Moria. 2021. Detecting White Supremacist Hate Speech Using Domain Specific Word Embedding with Deep Learning and BERT. *IEEE Access* 9 (2021), 106363–106374. <https://doi.org/10.1109/ACCESS.2021.3100435>
- [2] Wafa Alorainy, Pete Burnap, Han Liu, and Matthew L. Williams. 2019. “The Enemy Among Us”: Detecting Cyber Hate Speech with Threats-based Othering Language Embeddings. *ACM Transactions on the Web* 13, 3 (July 2019), 14:1–14:26. <https://doi.org/10.1145/3324997>
- [3] Tawia Ansah. 2021. Violent words: strategies and legal impacts of white supremacist language. *Virginia Journal of Social Policy & the Law* 28, 3 (2021), 305–340.

- [4] Frances Lee Ansley. 1989. Stirring the Ashes: Race Class and the Future of Civil Rights Scholarship. *Cornell Law Review* 74, 6 (1989), 993–1077. <http://scholarship.law.cornell.edu/crlrhttp://scholarship.law.cornell.edu/crlr/vol74/iss6/1>
- [5] Dan Arnaudo. 2017. Computational propaganda in Brazil: Social bots during elections. (2017).
- [6] Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Propy: Organizing the news based on their propagandistic content. *Information Processing & Management* 56, 5 (2019), 1849–1864.
- [7] Hannah Bast and Marjan Celikik. 2013. Efficient fuzzy search in large text collections. *ACM Transactions on Information Systems* 31, 2 (May 2013), 10:1–10:59. <https://doi.org/10.1145/2457465.2457470>
- [8] Chip Berlet and Stanislav Vysotsky. 2006. Overview of U.S. White Supremacist Groups. *Journal of Political and Military Sociology* 34, 1 (2006), 11–48.
- [9] Emily Blout and Patrick Burkart. 2020. White Supremacist Terrorism in Charlottesville: Reconstructing 'Unite the Right'. *Studies in Conflict & Terrorism* (2020), 1–22.
- [10] Christopher Brown. 2009. WWW.HATE.COM: White supremacist discourse on the internet and the construction of whiteness ideology. *Howard Journal of Communications* 20, 2 (April 2009), 189–208. <https://doi.org/10.1080/10646170902869544>
- [11] Fernando H. Calderón, Namrita Balani, Jherez Taylor, Melvyn Peignon, Yen-Hao Huang, and Yi-Shin Chen. 2021. Linguistic Patterns for Code Word Resilient Hate Speech Identification. *Sensors* 21, 23 (Jan. 2021), 7859. <https://doi.org/10.3390/s21237859>
- [12] Michelle Caswell. 2021. *Urgent Archives: Enacting Liberatory Memory Work*. Taylor & Francis. <https://books.google.com/books?id=76AsEAAAQBAJ>
- [13] Naganna Chetty and Sreejith Alathur. 2018. Hate speech review in the context of online social networks. *Aggression and violent behavior* 40 (2018), 108–118.
- [14] Katie Cohen, Fredrik Johansson, Lisa Kaati, and Jonas Clausen Mork. 2014. Detecting linguistic markers for radical violence in social media. *Terrorism and Political Violence* 26, 1 (2014), 246–256.
- [15] Giovanni Da San Martino, Stefano Cresci, Alberto Barrón-Cedeno, Seunghak Yu, Roberto Di Pietro, and Preslav Nakov. 2021. A survey on computational propaganda detection. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 4826–4832.
- [16] Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-Grained Analysis of Propaganda in News Article. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 5636–5646. <https://doi.org/10.18653/v1/D19-1565>
- [17] Jessie Daniels. 2009. *Cyber racism: White supremacy online and the new attack on civil rights*. Rowman & Littlefield Publishers.
- [18] Sean M Eddington. 2018. The communicative constitution of hate organizations online: A semantic network analysis of "Make America Great Again". *Social Media+ Society* 4, 3 (2018), 2056305118790763.
- [19] Mai ElSherief, Caleb Ziemis, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. 2021. Latent Hatred: A Benchmark for Understanding Implicit Hate Speech. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 345–363. <https://doi.org/10.18653/v1/2021.emnlp-main.29>
- [20] Abby L. Ferber (Ed.). 2004. *Home-grown hate: Gender and organized racism*. Psychology Press.
- [21] Abby L. Ferber and Michael Kimmel. 2000. Reading right: the Western tradition in white supremacist discourse. *Sociological Focus* 33, 2 (2000), 193–213.
- [22] Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)* 51, 4 (2018), 1–30.
- [23] Iginio Gagliardone, Danit Gal, Thiago Alves, and Gabriela Martinez. 2015. *Countering online hate speech*. Unesco Publishing.
- [24] John Gallacher and Marc Heerdink. 2021. Mutual radicalisation of opposing extremist groups via the Internet. <https://doi.org/10.31234/osf.io/dfc5>
- [25] Grace Garner, Madeleine McGrann, Maja Lynn, Rachel Kranson, Michael Miller Yoder, and Daniel Klug. 2022. *The Relationship Between Antisemitism and COVID-19 Conspiracy on Twitter*. Technical Report. Collaboratory Against Hate. <https://www.collabagainsthate.org/papers-presentations/antisemitism-and-covid-19-conspiracy-on-twitter>
- [26] Hedy Greijdanus, Carlos A de Matos Fernandes, Felicity Turner-Zwinkels, Ali Honari, Carla A Roos, Hannes Rosenbusch, and Tom Postmes. 2020. The psychology of online activism and social movements: Relations between online and offline collective action. *Current opinion in psychology* 35 (2020), 49–54.
- [27] Jakob Gulh and Jacob Davey. 2020. A safe space to hate: White supremacist mobilisation on Telegram. *Institute for Strategic Dialogue* 26 (2020).
- [28] R. Henderson and Elin McCready. 2018. How Dogwhistles Work. In *New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science)*, Sachiyo Arai, Kazuhiro Kojima, Koji Mineshima, Daisuke Bekki, Ken Satoh, and Yuiko Ohta (Eds.). Springer International Publishing, Cham, 231–240. https://doi.org/10.1007/978-3-319-93794-6_16
- [29] Katrina Hirvonen. 2013. Sweden: When hate becomes the norm. *Race & Class* 55, 1 (2013), 78–86.
- [30] Philip N Howard and Bence Kollanyi. 2016. Bots, # StrongerIn, and # Brexit: computational propaganda during the UK-EU referendum. *arXiv preprint arXiv:1606.06356* (2016).
- [31] Jade Hutchinson, Muhammad Iqbal, Mario Peucker, and Debra Smith. 2022. Online and Offline Coordination in Australia's Far-Right: A Study of True Blue Crew. *Social Sciences* 11, 9 (Sept. 2022), 421. <https://doi.org/10.3390/socsci11090421>
- [32] Emiliya Jokubauskaitė and Stijn Peeters. 2020. Generally Curious: Thematically Distinct Datasets of General Threads on 4chan/pol/. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 863–867. <https://ojs.aaai.org/index.php/ICWSM/article/view/7351>
- [33] Brian TaeHyuk Keum, Xu Li, and Michele J Wong. 2022. Hate as a system: Examining hate crimes and hate groups as state level moderators on the impact of online and offline racism on mental health. *International Journal of Intercultural Relations* 91 (2022), 44–55.
- [34] Rebecca Kay LeFebvre and Crystal Armstrong. 2018. Grievance-based social movement mobilization in the # Ferguson Twitter storm. *New Media & Society* 20, 1 (2018), 8–28.
- [35] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. 497–506. https://dl.acm.org/doi/pdf/10.1145/1557019.1557077?casa_token=G2JLP0oyyS8AAAAA:Q0od5ta0fvnLV-3eXclr7idBqNJJEQ5NaLpqYjcPW40EUsdHJcX71DgsRiXyTmXun-0B_VParTpMw
- [36] Yu-Ru Lin and Wen-Ting Chung. 2020. The dynamics of Twitter users' gun narratives across major mass shooting events. *Humanities and Social Sciences Communications* 7, 1 (Aug. 2020), 1–16. <https://doi.org/10.1057/s41599-020-00533-8> Number: 1 Publisher: Palgrave.
- [37] Yonatan Lupu, Richard Sear, Nicolas Velásquez, Rhys Leahy, Nicholas Johnson Restrepo, Beth Goldberg, and Neil F Johnson. 2023. Offline events and online hate. *Plos one* 18, 1 (2023), e0278511.
- [38] Preslav Nakov, Giovanni Da San Martino, and Firoj Alam. 2022. Fact-Checking, Fake News, Propaganda, Media Bias, and the COVID-19 Infodemic. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (Virtual Event, AZ, USA) (WSDM '22)*. Association for Computing Machinery, New York, NY, USA, 1632–1634. <https://doi.org/10.1145/3488560.3501395>
- [39] Lynnette Hui Xian Ng, Iain J Cruickshank, and Kathleen M Carley. 2022. Cross-platform information spread during the January 6th capitol riots. *Social Network Analysis and Mining* 12, 1 (2022), 133.
- [40] Lynnette Hui Xian Ng and Jia Yuan Loke. 2020. Analyzing public opinion and misinformation in a COVID-19 telegram group chat. *IEEE Internet Computing* 25, 2 (2020), 84–91.
- [41] John Oddo. 2019. *The Discourse of Propaganda: Case Studies from the Persian Gulf War and the War on Terror*. Penn State Press.
- [42] Antonis Papasavva, Savvas Zannettou, Emiliano De Cristofaro, Gianluca Stringhini, and Jeremy Blackburn. 2020. Raiders of the Lost Kek: 3.5 Years of Augmented 4chan Posts from the Politically Incorrect Board. *Proceedings of the International AAAI Conference on Web and Social Media* 14 (May 2020), 885–894. <https://ojs.aaai.org/index.php/ICWSM/article/view/7354>
- [43] Barbara Perry and Ryan Scrivens. 2016. White pride worldwide: Constructing global identities online. *The globalisation of hate: Internationalising hate crime* (2016), 65–78.
- [44] Kevin C Peterson, Phillip Martin, and Michael Edison Hayden. 2022. "Children of the kkk": White supremacist Patriot Front marches through Boston, attacks Black Artist. https://www.democracynow.org/2022/7/6/neo_nazi_nsc_131_patriot_front
- [45] Matthew Phillips, Arunkumar Bagavathi, Shannon E Reid, Matthew Valasik, and Siddharth Krishnan. 2018. The daily use of Gab is climbing. Which talker might become as violent as the Pittsburgh synagogue gunman. *The Washington Post* (2018).
- [46] Meredith L. Pruden, Ayse D. Lokmanoglu, Anne Peterscheck, and Yannick Veilleux-Lepage. 2022. Birds of a Feather: A Comparative Analysis of White Supremacist and Violent Male Supremacist Discourses. In *Right-Wing Extremism in Canada and the United States*. Palgrave Macmillan, 215–254. https://doi.org/10.1007/978-3-030-99804-2_9
- [47] Jing Qian, Mai Elsherief, Elizabeth Belding, and William Yang Wang. 2018. Hierarchical CVAE for Fine-Grained Hate Speech Classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 3550–3559.
- [48] Natasha Quek. 2019. El-Paso Shootings: Growing Threat of White Supremacists. *RSIS Commentary* (2019). <https://dr.ntu.edu.sg/handle/10220/49921>
- [49] Diana Rieger, Anna Sophie Kumpel, Maximilian Wich, Toni Kiening, and Georg Groh. 2021. Assessing the Extent and Types of Hate Speech in Fringe Communities: A Case Study of Alt-Right Communities on 8chan, 4chan, and Reddit. *Social Media and Society* 7, 4 (2021). <https://doi.org/10.1177/20563051211052906>

- [50] Shaden Shaar, Firoj Alam, Giovanni Da San Martino, Alex Nikolov, Wajdi Zaghouni, Preslav Nakov, and Anna Feldman. 2021. Findings of the NLP4IF-2021 Shared Tasks on Fighting the COVID-19 Infodemic and Censorship Detection. In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*. Association for Computational Linguistics, Online, 82–92. <https://doi.org/10.18653/v1/2021.nlp4if-1.12>
- [51] Brittany Shammass. 2020. What we know about the killing of 5-year-old Cannon Hinnant. <https://www.washingtonpost.com/nation/2020/08/14/cannon-hinnant-killing/>
- [52] Alexandra A. Siegel, Evgenii Nikitin, Pablo Barberá, Joanna Sterling, Bethany Pullen, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. 2021. Trumping Hate on Twitter? Online Hate Speech in the 2016 U.S. Election Campaign and its Aftermath. *Quarterly Journal of Political Science* 16 (2021), 71–104. <https://www.nowpublishers.com/article/Details/QJPS-19045>
- [53] Richard So, Hoyt Long, and Yuancheng Zhu. 2019. Race, Writing, and Computation: Racial Difference and the US Novel, 1880-2000. *Journal of Cultural Analytics* (2019), 1–30. <https://doi.org/10.22148/16.031>
- [54] Inés Bolaños Somoano and Richard McNeil-Willson. 2022. *Lessons From the Buffalo Shooting: Responses to Violent White Supremacy*. Technical Report. International Centre for Counter-Terrorism.
- [55] Eileen Sullivan and Katie Benner. 2021. Top law enforcement officials say the biggest domestic terror threat comes from white supremacists. *The New York Times* (May 2021). <https://www.nytimes.com/2021/05/12/us/politics/domestic-terror-white-supremacists.html>
- [56] Denise J Wilkins, Andrew G Livingstone, and Mark Levine. 2019. All click, no action? Online action, efficacy perceptions, and prior experience combine to affect future collective action. *Computers in Human Behavior* 91 (2019), 97–105.
- [57] Steven Windisch, Pete Simi, Kathleen Blee, and Matthew DeMichele. 2018. Understanding the micro-situational dynamics of white supremacist violence in the United States. *Perspectives on Terrorism* 12, 6 (2018), 23–37.
- [58] Meghan A Wong, Richard Frank, and Russell Allsup. 2015. The supremacy of online white supremacists—an analysis of online discussions by white supremacists. *Information & Communications Technology Law* 24, 1 (2015), 41–73.
- [59] Samuel C Woolley and Philip Howard. 2017. Computational propaganda worldwide: Executive summary. (2017).
- [60] Muheng Yan, Yu-Ru Lin, and Wen-Ting Chung. 2022. Are Mutated Misinformation More Contagious? A Case Study of COVID-19 Misinformation on Twitter. In *14th ACM Web Science Conference 2022 (WebSci '22)*. Association for Computing Machinery, New York, NY, USA, 336–347. <https://doi.org/10.1145/3501247.3531562>
- [61] Michael Miller Yoder, Ahmad Diab, David West Brown, and Kathleen M. Carley. 2023. A Weakly Supervised Classifier and Dataset of White Supremacist Language. Under review.