Characterizing the Toxicity of the Brazilian Extremist Communities on Telegram

Athus Cavalini
athus.cavalini@ifes.edu.br
Universidade Federal do Espírito Santo
Vitória, Brazil

Thamya Donadia
thamya.donadia@edu.ufes.br
Universidade Federal do Espírito Santo
Vitória, Brazil

Giovanni Comarela
gc@inf.ufes.br
Universidade Federal do Espírito Santo
Vitória, Brazil

ABSTRACT

Telegram has become a central element in discussions related to the ecosystems of information disorder and extremism on social networks. Present on 70% of smartphones in Brazil, the application presents itself as a safe communication space, which began to be used by deplatformed individuals and groups, including extremist groups, who saw the application as a space for building communities and maintaining contact with their audience. In this context, this study presents a characterization of Brazilian extremist communities on Telegram based on the analysis of over 2 million messages broadcast on 128 chats on the platform, focusing on the analysis of the toxicity observed in the content shared in these spaces and its relationship with the conversational dynamics of the groups. The results reveal that these communities share highly toxic messages, including manifestations of hate speech and conspiracy theories, and that the toxicity of the content reflects on its popularity and consequently its spread across the network.

KEYWORDS

Extremism, Toxicity, Telegram, Brazil

1 INTRODUCTION

Telegram emerged in 2013 as a tool designed to provide a secure communication space, particularly for individuals residing in authoritarian states, with a strong emphasis on user privacy [6]. Over the years, Telegram has experienced remarkable growth, reaching approximately 800 million monthly active users (MAU) by the year 2023 [19]. Originally intended as a platform for safe communication, Telegram has, however, garnered attention from historically deplatformed extremist groups [16], including white supremacists [6] and terrorist organizations such as the Islamic State and al-Qaeda [21].

In the period from 2020 to 2021, [17] identified a significant increase in radicalization in speeches present in Telegram messages. Research on this subject highlights that extremist channels have adopted more popular approaches, allowing unrestricted access to users interested in participating in discussions, which become increasingly violent and fueled by hate, propagating extremist ideologies, disseminating materials related to these ideologies, and promoting attacks to minority groups, as documented by [20].

In Brazil, Telegram's popularity is also observed, where it is currently installed on the majority of the smartphones.[1] Motivated by this popularity, many studies are beginning to shed light on its adoption by extremist communities, encompassing groups on the far-right [2] and anti-vaccine communities [11].

As the app's popularity continues to grow considerably, concerns emerge regarding its role in facilitating the spread of extremist ideologies and misinformation and its impact on social dynamics. The application's resistance to comply with legal decisions, notably in Brazil, adds another layer of complexity to this issue. Hence, understanding the dynamics of these groups within the Telegram ecosystem becomes crucial, and researchers face the challenge of addressing the implications of Telegram's evolving role in the formation of extremist communities and the spread of their discourse.

Considering this context, this study aims to examine channels and groups that compose Telegram's extremist communities in Brazil. We focus on two main Research Questions (RQ): **RQ1:** How toxic is the content shared by the extremist community, and how does toxicity impact content popularity? **RQ2:** What are the main topics of discussion among the extremist community and how are these topics related to the observed toxicity levels?

Answering these questions is a challenging task. First, to the best of our knowledge, there is no readily available and suitable dataset for the necessary analyses. Second, messages can be shared in different ways (e.g., text and audio), and may, possibly, contain several artifacts, such as grammar errors, informal language, and emojis. To tackle these challenges, we carefully devised collection strategies and we identified, and put together, a set of tools capable of extracting information from messages, assessing toxicity, and handling the aforementioned issues.

By answering **RQ1** and **RQ2**, the main contribution of this work is to promote an overview of the communication dynamics of these communities, and how extremist movements have adapted to digital spaces to get closer to young people and expand their audience.

2 RELATED WORK

There are already several tools and methodologies aimed at studying Telegram and its content, some of them highlighted in the literature. Telegram Monitor, developed by [9], stands out in this scenario, allowing the collection and processing of data in the application. This resource has been used by the author to monitor political groups and channels in the Brazilian context. On the other hand, [2] presents an analysis of the use of Telegram as a vehicle for

In: Proceedings of the Brazilian Symposium on Multimedia and the Web (WebMedia'2024). Juiz de Fora, Brazil. Porto Alegre: Brazilian Computer Society, 2024.

ISSN 2966-2753

[1]Telegram's popularity in Brazil reached its peak in August 2022, and since then it has seen a small adoption decline – https://www.forbes.com/sites/angelicamarideoliveira/2023/03/01/messaging-app-telegram-sees-decline-in-brazil/

spreading misinformation by far-right groups in Brazil, adding a critical perspective to the use of the platform for such purposes.

By expanding to an international perspective, [7] brings a comprehensive analysis of the globalization of the QAnon movement, exploring its communities on Telegram. Their study reveals the movement's approach to issues related to global politics, conspiracy theories, and even the anti-vaccine movement. Similarly, [20] offers a broader view by indicating the growth of extremist ideologies and the spread of misinformation campaigns on the app, opening discussions about the need to combat this practice on a global level.

In this context, [3] presents that social media offers an unmoderated space for communication among like-minded individuals, which may intensify extremist views. Furthermore, according to [12], the rise of right-wing extremism is marked by the emergence of an underground subculture consisting of young extremists who have developed their own symbols, myths, and language.

These papers highlight not only the potential of Telegram as an object of study but also the challenges and risks associated with its use in political and social contexts. As the platform continues to play a significant role in our communication model, the critical analyses presented in the literature contribute to a deeper understanding of the dynamics of the ecosystems found on Telegram.

However, there are still several issues to be addressed, especially involving detailed descriptions of the operating logic of these communities and their main strategies for gaining an audience and bringing their narratives closer to the public. In this way, this work contributes to the existing literature by exposing and quantifying the toxicity of the Brazilian extremist communities on Telegram and by characterizing their main topics of interest. To the best of our knowledge, ours is the first work to study this specific phenomenon.

3 METHODS

The corpus of this paper was collected through the official Telegram API. For the purpose of this work, we collect data relating to a period of six months, more specifically between April 1st, 2023 and September 30, 2023, which can be considered a noise-free period based on the Brazilian context, which has its platforms flooded with political-party agenda during electoral periods due to the political polarization experienced in recent decades.

The approach employed to obtain the Brazilian extremist ecosystem of Telegram was based on a controlled snowball sampling technique, widely used in research based on similar platforms [14]. We chose to use the variation proposed by [2], which utilizes PageRank to include only relevant actors and keep the collection within the chosen community. In this method, a small but very representative set of channels or groups are collected, and other channels that prove influential based on relationships with the initial set are selected and added to the collection in a new iteration.

For the initial set, we manually analyzed channels and groups mapped by [8] in the context of the anti-vaccine movement, from which five chats were selected where the presence of extremist content was identified. The relationships between channels were represented by the number of messages forwarded between them and the metric used to calculate influence was PageRank [13].

In addition to the main corpus, a baseline corpus was obtained using the same methodology from an initial set of 50 generic Telegram channels and groups obtained from public repositories found on Google Search. No manual filtering was applied to this set.

For the analysis, both text and transcribed voice messages were submitted to the Perspective API [10], a model that scores the toxicity levels of a text, in addition to identity attacks, insults, profanity, and threats. The voice messages present in the dataset were transcribed using Whisper [15], a speech recognition model trained on a large multi-language dataset, in order to be able to apply the same analyses used for text messages.

Additionally, topic modeling was applied to the content using BERTopic [5], a state-of-the-art technique that leverages transformers and c-TF-IDF in order to achieve coherent topics. The BERT model used in this technique was BERTimbau [18], which was pretrained on a large Brazilian Portuguese corpus, for 1,000,000 steps, using a whole-word mask.

Albeit new and robust topic modeling algorithms continue to emerge, their innovative advantages often diminish the significance of traditional model evaluation metrics such as perplexity and coherence [4], which can be inversely proportional to interpretability [1]. For instance, BERTopic has been noted to generate more outliers than expected, complicating its interpretability [4].

In that way, we chose to use OpenAI's ChatGPT 3.5 as the representation model of the BERTopic framework, in order to generate topic descriptions based on the representative documents outputted by BERTopic for each topic. This allows us to automate the interpretation of the topics without human bias or the necessity of an extensive (and potentially subjective) analysis of documents.

After the mining and processing phases, we conducted an in-depth quantitative analysis of the obtained results, while expanding our insights by performing qualitative analyses of data samples from multidisciplinary perspectives. The main objective was to obtain a detailed and comprehensive understanding of the extremist ecosystem on Telegram and its conversational dynamics.

4 RESULTS AND DISCUSSION

In this section, we present a description of our dataset and an extensive discussion about the obtained results.

4.1 Dataset Description and Network Dynamics

The final corpus[2] was composed of 128 Telegram chats. During the six-month period analyzed, 2,420,326 messages were shared, around 60% of which were plain text messages. Furthermore, 15,789 voice messages were identified, being downloaded and transcribed.

Concerning the network's audience, channels and groups have an accumulated number of 1.282.715 participants. Groups, that represent discussion spaces (a many-to-many conversational logic), have an average of 528,5 participants, while channels, which can be seen as spaces for broadcast (a few-to-many conversational logic), have an average of 11.574,5 participants.

On Telegram, the channels' participants lists are private, showing only the absolute number. From the groups, however, the complete list of participants can be obtained. This gives us the possibility to calculate the average number of groups per participant. In this

[2] Available at https://github.com/dsl-ufes/telegram-extremist-communities/.

case, each user participates in an average of 1.06 groups, indicating that the accumulated audience can approach the absolute audience, that is, the number of unique users on the network.

Table 1: Chats metrics. Average per channel/group.

		Views	Forwards	Replies	Toxicity
Dataset	Groups	3771.11	23.93	0.98	0.26
	Channels	1339.73	6.81	3,33	0.30
Baseline	Groups	4794.79	30.56	1.17	0.14
	Channels	396.12	1.46	0,.2	0.04

The discrepancy in engagement between the baseline channels and groups and the extremist community is notable, as shown in Table 1. The extremist community clearly shows a more verticalized communication, where the channels concentrate the volume of information and engagement. This trend is confirmed by the complexity of the messages shared on these channels, that is designed to be disseminated. They feature more elaborate and well-structured narratives, generally accompanied by images, videos, and external URLs. Also, while messages shared on channels have an average length of 253.61 characters, in groups the messages tend to be shorter, with an average length of 66.75 characters.

A critical distinction lies in the contrasting reply dynamics between channels and groups. Groups exhibit a conversational flow, characterized by sequential messaging with minimal references to older posts, while channels primarily utilize replies as the main and most comprehensive mode of public expression, given that only administrators can post original messages. Telegram's structure further distinguishes channels by displaying replies on separate screens rather than integrating them with the main interface. This difference in reply usage reflects the distinct communication patterns and engagement styles fostered by channels and groups.

Also, it is notable that, on the extremist dataset, users are more participative in discussions. In comparison to the baseline, in the extremist community users have shared an average of 65.56 messages, while in the baseline this average drops to 22.40. Additionally, in the extremist community, around 21% of users were responsible for the production of 90% of the messages, while in the baseline just 10% of users were necessary to produce the same percentage.

Furthermore, regarding toxicity, it is notable that groups exhibit higher levels of toxicity compared to channels. This is caused by the type of content observed across different chats, given that groups are conversational environments whereas channels are informational. In this way, messages shared in groups tend to carry a more expressive tone, closer to real dialogue and aligned with the linguistic characteristics associated with the participants' profile.

4.2 Toxicity and Engagement

The objective of this section is to carry out an analysis of the toxicity of the studied community and identify possible relationships between the toxicity of a content and its engagement.

In an overview, it is possible to observe that extremist channels and groups share more toxic content than the baseline. In the baseline set, the average toxicity score for text messages was 0.17, while in the extremist community, the average score was 0.29, an increase of 70%. The numbers also indicate that voice messages tend to be more toxic than text messages, even in the baseline dataset. This difference can be seen in detail in Figure 1.

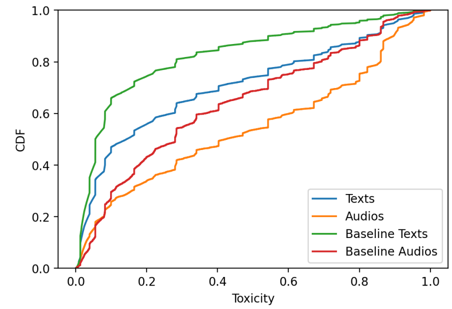

Figure 1: Toxicity of text and audio messages.

For this analysis, we only considered engagement as the number of forwards that messages received. This is because reactions on Telegram need to be categorized (as they can be positive or negative), while the possibilities are diverse. Furthermore, the number of replies may be underrepresented, since the conversational dynamics of the application do not suggest the use of the reply action, but rather the sequential sending of messages. Finally, the views metric also does not add value, since, in a conversation, all participants see almost all messages. This number only grows significantly when a message is forwarded to other channels and groups, accumulating views from all of them.

When comparing toxicity and engagement, we have interesting findings. For the majority of content circulating in the extremist community, characterized as conversational messages with few forwards (i.e., less than 50), non-toxic messages have a 97% higher number of forwards, which evidences their greater likelihood of being shared. Among the toxic and highly toxic, however, the highly toxic ones are more likely to be forwarded.

When we look at messages with high numbers of forwarding, the logic is reversed. Highly toxic messages start to receive more forwards, followed by toxic ones. In other words, this indicates that, for content that actually spreads across the network, reaching various channels and groups within the community (and potentially outside of it), more toxic content reaches a larger audience. Figure 2 shows this trend.

4.3 Discussion Topics

Upon observing the discussion topics extracted from the dataset, it is plausible to glean a comprehensive understanding of the most recurrent and pertinent subjects to the audience of this community.

For this analysis, we conducted four distinct modeling approaches, segregating text and voice messages, as well as non-toxic (with scores below 0.8) and highly toxic ones (starting from 0.8). Tables

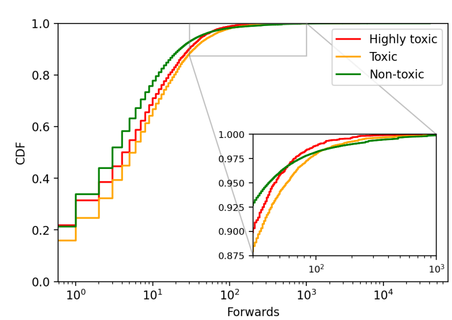

Figure 2: Forwards of non-toxic (score < 0.5), toxic (0.5 ≤ score < 0.8), and highly-toxic (≥ 0.8) messages.

2 and 3 delineate the topics that prominently emerge within the corpus of highly toxic text and voice messages, respectively. Certain topics have been omitted for the sake of visual clarity.

Table 2: BERTopic for toxic text messages.

#	Topic Description	# Messages
0	Jesus' Passion and Crucifixion in the context of Christian prayers and reflections	26247
1	Cultural Warfare and Political Commentary on Brazil	21467
2	Military conflict updates between Russia and Ukraine	16828
4	Discussion on vaccines and conspiracy theories surrounding COVID-19	5924
5	[Explicit content warning]	1949
6	Animal Welfare and Misconceptions	1270
7	Tensions in Niger - Military Coup, Uranium Ban, and International Intervention	1023
8	Manipulation of public opinion, digital privacy and challenges in the information age	926
9	Autism at school and the power of suggestion in teaching	861
10	Controversies in Space Exploration and Military Technology	523
13	Controversial opinions on women's football	250
14	Red Pill Movement Against Political Idolatry	228
17	Anti-New World Order Movement	96
18	Uganda's New Law Imposing Harsh Penalties on Homosexual Practices	64

It is noteworthy that the topics encompass clearly contentious agendas, such as wars, the "Red Pill" movement, and conspiracy theories involving vaccines, diseases, space exploration, and even the "New World Order", an elite wielding global control. Moreover, topics of potential hate speech are also highlighted. This is exemplified by Topic 13 in text messages, described as "Controversial opinions on women's football", as well as Topic 3 in voice messages. Among the most representative messages within these topics, one can observe discourses akin to that illustrated in Figure 3, which contains explicit expressions of sexism and misogyny.

Text message: It's very gratifying to see Brazilian-shit whores losing, they are beings devoid of any honor or loyalty competing just to validate their inflated ego. They will never understand the male world of football because they are there just for vanity.

Figure 3: Message from Topic 13 (from text messages).

It is important to highlight that, in the case of Topic 5 in text messages, the GPT model did not even generate a description, issuing a content warning for explicit material. This occurs when the content provided to the model exceeds the boundaries permitted in its privacy policy, which includes misinformation and hate speech.

Table 3: BERTopic for toxic voice messages.

#	Topic Description	# Messages
0	Brazilian societal issues and political commentary	367
1	Criticism of religious beliefs and societal norms	588
2	Racism and Race Relations, Cultural Assimilation, and Immigration in Society	105
3	Gender Dynamics and Relationship Perspectives	457
4	Virulent conversations about controversial relationships and behaviors	806

Another noteworthy point is the notable disparity between the tones of the topics extracted from text messages and voice messages. While the texts, even though part of conversational dynamics, present contents connected to more specific themes, the voice messages present broader topics that, in a way, can be interpreted as opinion topics or a self-expression space. In this sense, discussions about personal relationships and ideologies stand out, as observed in Topics 2, 3, and 4, representing more personal content associated with individuals and their reality, rather than topics of general interest and/or global impact. Figure 4 illustrates a message associated with Topic 3 of voice messages, reinforcing this pattern.

Voice message: No, totally accurate. Like, in the old days, well, not so much "old days", but in traditional society, it was the man who chose the husband for his daughter. And, if a guy wanted to marry his daughter, he had to bring something to the table. Because women don't have testosterone, they don't produce, they don't work. They don't generate wealth, at least not material wealth. [...] But yeah, dude, modern society, it's messed up, man, and he's gonna deal with a lot of crap.

Figure 4: Message from Topic 3 (from voice messages).

The topics generated from moderated or non-toxic messages, on the other hand, were even more linked to international geopolitical issues, in addition to being much more informational. While the audios were limited to Brazilian political-party discussions and religious reflections, the text messages presented topics such as "Global Geopolitical Tensions and Economic Concerns", "Chinese Market" and "Conspiracies Regarding the COVID-19 Pandemic".

5 FINAL REMARKS

The results of this study emphasize the direction of extremist communities toward the public, highlighting their wide reach. In response to RQ1, it is notable that the content disseminated by these communities is highly toxic compared to other content circulating on the same social network, reflecting the linguistic pattern of the audience present in these spaces.

Regarding RQ2, we observed that there is a clear difference between the content generated by the interactions of the audience, generally on the groups and which we called *conversational messages*, in relation to the fabricated content published on channels, that we called *informative messages*.

In conversational interactions, where toxicity is even more pronounced, the audience expresses themselves and discusses topics attached to everyday life. Analyzing text and voice messages, we highlight that voice messages are significantly more toxic and present distinct topics, with an even more personal tone. In other words, text is used for sharing information and debate, while voice messages serve as personal expression and counseling.

Therefore, even though they represent a small fraction of the total content in these communities, voice messages provide valuable insights into conversational, linguistic, and even ideological patterns of the audience present in these spaces. This underscores the construction of subjectivity within this social bubble, revealing the worldview that its members construct and share regarding the world, society, and their surroundings.

Additionally, returning to RQ1, when examining highly shared messages, most of which originate from channels, those with high toxicity receive more shares than those with moderate or non-toxicity. This suggests that, even concerning messages from a more hierarchical structure and of a more informative nature, there is a preference for more toxic content in terms of engagement and circulation. This implies that the content produced, when adjusted or shaped to the linguistic patterns of the audience, tends to be more widely accepted and thus more widely disseminated.

ACKNOWLEDGMENTS

This work was partially supported by: Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) – Finance Code 001; FAPES – 368/2022 - P: 2022-NGKM5, 129/2021 - P: 2021-GL60J, 1026/2022 - P: 2022-ZQX6F; and Fapesp/MCTI/CGI.br (#2020/ 05182-3).

REFERENCES

[1] Henrik Axelborn and John Berggren. 2023. Topic Modeling for Customer Insights : A Comparative Analysis of LDA and BERTopic in Categorizing Customer Calls. , s. 47 pages.

[2] Athus Cavalini, Fabio Malini, Fabio Gouveia, and Giovanni Comarela. 2023. Politics and disinformation: Analyzing the use of Telegram's information disorder network in Brazil for political mobilization. *First Monday* (May 2023). https://doi.org/10.5210/fm.v28i5.12901

[3] Olga V. Chyzh, Mark David Nieman, and Clayton Webb. 2019. The Effects of Dog-Whistle Politics on Political Violence. https://api.semanticscholar.org/CorpusID:195433166

[4] Roman Egger and Joanne Yu. 2022. A Topic Modeling Comparison Between LDA, NMF, Top2Vec, and BERTopic to Demystify Twitter Posts. *Frontiers in Sociology* 7 (2022). https://doi.org/10.3389/fsoc.2022.886498

[5] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).

[6] Jakob Guhl and Jacob Davey. 2020. A Safe Space to Hate: White Supremacist Mobilisation on Telegram.

[7] Mohamad Hoseini, Philipe Melo, Fabricio Benevenuto, Anja Feldmann, and Savvas Zannettou. 2023. On the Globalization of the QAnon Conspiracy Theory Through Telegram. In *Proceedings of the 15th ACM Web Science Conference 2023* (Austin, TX, USA) *(WebSci '23)*. Association for Computing Machinery, New York, NY, USA, 75–85. https://doi.org/10.1145/3578503.3583603

[8] Instituto Capixaba de Ensino, Pesquisa e Inovação em Saúde. 2024. Projeto Observa ICEPi. https://icepi.es.gov.br/projeto-observa-icepi Accessed on: January 2, 2024.

[9] Manoel Júnior, Philipe Melo, Daniel Kansaon, Vitor Mafra, Kaio Sa, and Fabricio Benevenuto. 2022. Telegram Monitor: Monitoring Brazilian Political Groups and Channels on Telegram. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media* (Barcelona, Spain) *(HT '22)*. Association for Computing Machinery, New York, NY, USA, 228–231. https://doi.org/10.1145/3511095.3536375

[10] Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. 2022. A New Generation of Perspective API: Efficient Multilingual Character-level Transformers. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) *(KDD '22)*. Association for Computing Machinery, New York, NY, USA, 3197–3207. https://doi.org/10.1145/3534678.3539147

[11] Fabio Luiz Malini, Fabio Gomes Goveia, Athus Cavalini, and Francis Sodré. 2023. Padrões da desinformação vacinal no Telegram. In *Proc. of IX Congresso Brasileiro de Ciências Sociais e Humanas em Saúde*, Suely Deslandes and Marta Verdi (Eds.). Associação Brasileira de Saúde Coletiva, Recife, 46626.

[12] Cynthia Miller-Idriss. 2009. *Blood and Culture.* Duke University Press, New York, USA. https://doi.org/doi:10.1515/9780822391142

[13] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1998. *The PageRank Citation Ranking: Bringing Order to the Web.* Technical Report. Stanford Digital Library Technologies Project. https://doi.org/10.1.1.31.1768

[14] Stijn Peeters and Tom Willaert. 2022. Telegram and Digital Methods: Mapping Networked Conspiracy Theories through Platform Affordances. *M/C Journal* 25, 1 (Mar. 2022). https://doi.org/10.5204/mcj.2878

[15] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust Speech Recognition via Large-Scale Weak Supervision. https://doi.org/10.48550/ARXIV.2212.04356

[16] Richard Rogers. 2020. Deplatforming: Following extreme Internet celebrities to Telegram and alternative social media. *Media Studies* (2020). https://doi.org/10.1177%2F0267323120922066

[17] Heidi Schulze, Julian Hohner, Simon Greipl, Maximilian Girgnhuber, Isabell Desta, and Diana Rieger. 2022. Far-right conspiracy groups on fringe platforms: a longitudinal analysis of radicalization dynamics on Telegram. *Convergence* 28, 4 (2022), 1103–1126. https://doi.org/10.1177/13548565221104977 arXiv:https://doi.org/10.1177/13548565221104977

[18] Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2020. BERTimbau: pretrained BERT models for Brazilian Portuguese. In *9th Brazilian Conference on Intelligent Systems, BRACIS, Rio Grande do Sul, Brazil, October 20-23 (to appear)*.

[19] Statista. 2023. Telegram messenger global MAU 2014-2023. https://www.statista.com/statistics/234038/telegram-messenger-mau-users/. https://www.statista.com/statistics/234038/telegram-messenger-mau-users/ Accessed: 2024-01-31.

[20] Samantha Walther and Andrew McCoy. 2021. U.S. Extremism on Telegram: Fueling Disinformation, Conspiracy Theories, and Accelerationism. *Perspectives on Terrorism* 15, 2 (Apr 2021). https://pt.icct.nl/article/us-extremism-telegram-fueling-disinformation-conspiracy-theories-and-accelerationism

[21] Gabriel Weimann. 2016. Terrorist Migration to the Dark Web. *Perspectives on Terrorism* 10, 3 (2016), 40–44. http://www.jstor.org/stable/26297596