Analysing Protest-related Tweets: An Evaluation of Techniques by the Open Source Intelligence Team

Laurens H.F. Müter[123] and Remco C. Veltkamp[2]

[1] Netherlands Police, Van Deventerlaan 50, 3528 AE Utrecht, the Netherlands,
WWW home page:
https://www.uu.nl/onderzoek/ai-labs/nationaal-politielab-ai
[2] Utrecht University, Nationaal Politielab AI,
Leuvenlaan 4, 3584 CE Utrecht
[3] l.h.f.muter@uu.nl

Abstract. The police's Open Source Intelligence (OSINT) team is constantly looking for better ways to analyze large corpora with high precision. This paper presents the evaluation by the OSINT team of more advanced methods than currently used to extract information from tweets related to an upcoming large-scale demonstration. An initial interview revealed the current way of working. Next, more advanced machine learning techniques such as sentiment analysis and network analysis are used to create visualizations that would better suit OSINT's needs. Finally, our proposed visualizations are evaluated by two OSINT analysts who state that the results are clear, actionable, and relevant, whereas completeness of information and privacy pose additional challenges.

Keywords: Natural Language Processing, Twitter, Protest, OSINT

1 Introduction

Civil protests are a well-known phenomenon in many societies [1]. Over the last decades, an upward trend in the number of protests can be perceived, where periods of economic and political crisis have a tendency to intensify the occurrences of protests [2]. Protests are an expression of social movements, where entities are formed by people who challenge the status quo [3]. Social movements are usually initiated by people who fall victim to political decisions, actions, or policy [4]. When a social movement starts to take shape, other people might join, thus strengthening the movement [5]. The emergence and expansion of social movements are difficult to predict but can have a considerable impact on public safety [6] if, for example, demonstrations take place in public space. Public safety is one of the reasons the Dutch police are interested in anticipating social movements and related demonstrations.

The Open Source Intelligence (OSINT) team of the Dutch police is estimating the social impact of large-scale social events, such as demonstrations. To find information regarding a demonstration, the police use open and closed sources. One of these open sources is the microblogging platform Twitter. While the OSINT team considers a range of platforms and sources, we decided to only use Twitter within this study, since the data from this platform is easily obtainable and provides insight into what people are discussing. Twitter is a frequently used information source within the Dutch police: two-thirds of the interviewed Dutch police officers reported having gained valuable information from Twitter [7].

Various multi-tool analysis techniques are designed to analyze tweets; however, many of these tools and techniques are based on commercially available software, limiting the OSINT team's ability to use a wider spectrum of machine learning and visualization techniques. To address this limitation, a case study is presented where information related to an upcoming large-scale protest is extracted from tweets using commercially available techniques. This exploratory case study explores existing visualizations and information extraction techniques and how these can be used to strengthen in information position of the police's OSINT team.

We present the evaluation by the police's OSINT team of more advanced methods than currently used to extract information from tweets related to an upcoming large-scale demonstration. An initial interview revealed the current way of working. Next, more advanced machine learning techniques such as sentiment analysis and pairwise co-occurrence network analysis are used to create visualizations that would better suit the OSINT team's needs. These are evaluated by the police's OSINT team, where the relevance, correctness, and explainability are considered to be strong, whereas completeness of information and privacy pose additional challenges.

2 Related Work

Various studies have been conducted in the research field regarding information spreading via Twitter during demonstrations.

Tonkin et al. [8] show that during the London riots in August 2011, irrelevant tweets died out, whereas the number of retweets from popular individuals increased exponentially. This has the potential that relevant information can be found within tweets, which is also pointed out by Procter et al. [9], who take a more detailed approach to studying information sharing during demonstrations. They ranked tweets by retweet counts to find influential tweets and users, where four stages of information spreading were distinguished: (1) public announcement, (2) the idea is picked up by a group, (3) dedicated accounts are set up to coordinate activities, and (4) actors with many followers are involved. Burch et al. [10] took things a step further by also including additional information from tweets such as usernames, geographic location, tweet timestamps, and tweet content in the context of the Vancouver riots on June 15, 2011. These measurables were manually reduced to a single term to describe each tweet (e.g. "Fan Reflection"). Next, these terms were used to find overlapping themes, which resulted in five categories (Fan Reflection, Fandom, Riot Propagation, Global Perspectives, and Shame on Vancouver) that provide insights into what was happening during these riots [10].

The work of Burch et al. shows that each overlapping theme adds value, whereas the complement of these themes contains more information than each individual theme. To further study the effect of complementing measures within tweets, Our paper also builds upon the premise that the combination of multiple measurables from tweets can strengthen their individual results, hence multiple tweet measurables are used: textual body, keywords, hashtags, emojis, and retweets. Each of the measurables is described within the context of the case study to show how the extracted demonstration-related information can strengthen the information position of the police.

2.1 Hashtag

A hashtag consists of a hash symbol followed by a word that bookmarks a tweet (e.g. #Weather), thus linking a tweet to a specific topic. However, hashtags can also serve as the symbol of a community by linking users with similar ideas [11]. For example, a user can declare to be part of a community by using a specific hashtag (even if the user might not be aware of this).

Although only approximately 13% of all tweets contain one or more hashtags [12], they can provide a quick overview of popular topics. Moreover, tweets often contain multiple hashtags, so highly related hashtags can be found by analyzing the occurrences of hashtags in tweets. A way to cluster hashtags is by creating a topic network using co-word analysis, which can give insights into how hashtags appear together [13]. Within a topic network, hashtags that contain related topics are often in close proximity to hashtags that directly describe an incident. This way, topic networks can be used to filter hashtags related to incident-related topics, even when these tweets contain hashtags that are indirectly related to the incident.

2.2 Emojis

An emoji refers to a pictograph that is used to express emotions and to add color and playfulness to messages, as stated by Wijeratne et al. [14]. Similar to words, emojis can have different meanings depending on the context of a message.

Emojis are assigned to an international standard UTF-8 encoding [15]. Due to a constant need to create and extend the stock of available emojis (e.g. applying different skin colors to a smiley) [16], the Unicode consortium has composed strict guidelines to which emoji suggestions must comply. The Unicode consortium also releases all newly approved emojis every year [17].

Emojis have been used in tweet-related studies. For example, Ljubešić and Fišer [18] study the spatial distribution of emojis in a data set of more than 17 million tweets. By using cluster analysis, they find correlations between emoji distributions and the differences in living conditions in various parts of the world.

Another way to look at emojis in the context of tweet analysis is sentiment analysis, as described by Chen et al. [19]. In their study, they adopted attention-based Long Short-Term Memory Networks (LSTM) in combination with a Python library for text classification named fastText [20] to describe emojis as a special token. The fastText representation of emojis is then used to obtain information about the sentiment of an emoji with respect to its context (e.g. a smiling emoji can represent joy or sarcasm). With their LSTM encoder model that includes emojis, they are reportedly able to classify the sentiment of tweets with better accuracy than using an LSTM encoder on text only.

Fede et al. [21] state that emojis can also be represented as a network that connects emojis that occur in the same tweet. This network can then be used for multiple sets of tweets that relate to topics such as sports and large news events. Form a directed network, where each emoji is connected to the emojis from the same tweet. They also analyzed their directed emoji network using different techniques (Frequency, Page rank, and Node betweenness), which resulted in a top ten emojis list per data set.

2.3 Retweets

Retweeted messages on Twitter can tell something about the spreading of information within a social network [22]. With this information, it is possible to select candidate messages that are likely to be retweeted and obtain insights into how public opinion will evolve. In other words, by finding retweet candidates, near-future social acumen can be obtained. Finding retweet candidates can be done in different ways. For example, by determining the most influential users in a network [22] or building a logistic regression model based on user topic distribution and whether a message has been retweeted before [23].

As such, retweet counts by itself can hold additional value for this study in combination with sentiment analyses. This way, highly retweeted messages can give insights into public opinion about specific topics, as stated by Lashari and Wiil [24].

2.4 Combinations

Some studies show some combinations of measurables. Zhou and Ai [25] show a strong correlation between hashtags and emojis in tweets, where emojis and hashtags strengthen their expressions of sentiment, display identity, and highlight keywords. They even expanded by generating a model to predict hashtags from emojis in a tweet. Other findings include emojis being less ubiquitous in combination with hashtags, despite cultural differences, although in most cases users choose either to include hashtags or emojis [26], where hashtags more often appear in tweets than emojis [27]. Furthermore, hashtags rise and fall rapidly in popularity [28] and emojis have limitations in expressiveness. However, the combination of hashtags and emojis can be used in a way that emojis can help to find expressive hashtags, that can be used in finding meaningful context from tweets.

Another way of combining different measurables is by including GEO data, which is present in 1% of the tweets [29]. Li et al. [30] used emojis to obtain sentiment scores of tweets, which also contained geographical information. By plotting these sentiment scores, Li et al. [30] are able to estimate the impact of events and visualize the spreading of such events.

Previous studies show promising results in gaining insights from tweets using single measurables such as retweets, hashtags, emojis, and keywords. However, uncertainty remains regarding the value of combined measurables. Moreover, little is known about the practical implications of using these measurables in complement within the context of the police, with real-world scenarios.

3 Legal and Ethical Considerations

When collecting and handling data, multiple perspectives should be considered. First, from an ethical perspective, the consequences are discussed on collecting information from people who have not given their consent to use their data for the purpose of making estimations about upcoming protests. Second, the legal perspective on handling the data relates to considerations for the police about collecting, analyzing, and storing data. Finally, from a practical perspective, questions are considered on which information the police need to make workable analyses.

3.1 Ethical Perspective

In the existing literature, multiple challenges are described in using tweets for research. For example, Webb et al. [31] conducted interviews with experts about possible ethical implications regarding the usage of tweets in academic research. Based on these interviews, Webb et al. [31] constructed surveys that were set out to a wider range of Twitter analysts about their current practices and opinions on topics such as privacy and storage of tweets. From these surveys, they extracted multiple topics to consider when working with tweets: (1) absence of academic consensus, (2) informed consent, (3) minimizing harm, (4) anonymization, (5) commercial vs regulatory vs. academic practice, and (6) research integrity.

To cope with the first of these challenges, academic consent, we obtained ethical approval for this study from appropriate institutional review boards and ethics committees [32]. However, since the usage of tweets in the context of protest prediction is somewhat controversial, additional considerations are also described. Informed consent is impossible to obtain from Twitter users. However, in order to use the services of Twitter, users have to agree on a statement that information from the platform is shared with third parties. Moreover, a recent study showed that 73% of Twitter users are aware of this [33]. Thus, it would be likely to assume that Twitter data can be used for the purpose of this study.

To minimize possible harm caused by this study, considerations are taken on what information is publicized. First, specific information about groups or individuals, for example, names, location indications, or direct text that would allow tracing individuals, is excluded. Second, detailed descriptions of conversations or unique expressions and citations in tweet texts are excluded since cited tweet text can lead to the user that posted the tweet. Third, results are presented in an abstract and aggregated manner. Anonymization is realized by hashing usernames and mentions where guidelines on academic best practices [34] are followed.

3.2 Legal Perspective

How to handle data is bound by legal constraints, which mainly focus on privacy. Platforms might differ in what is allowed in relation to processing data from users. The policies of platforms can change, which might have consequences on how the data is handled during or after the study. The OSINT team uses the GDPR (General Data Protection Regulation) [35] to decide on what data they can use. For example, profiles to monitor individuals are not allowed, even if it involves anticipation of potential crimes and riots [35]. For this study, no additional legal constraints regarding the usage of tweets when the platform policy is updated in the future are applicable, as the GDPR states that processing for scientific and historical research, statistical purposes, and archiving public interest are compatible with the original collecting purpose [35].

3.3 Practical Perspective

Based on the work of Rivers and Lewis [36], we describe three suggestions for handling data obtained from Twitter. First, the information available for an analyst should not expose private sensitive information. It is worth noting that username mentions carry a higher risk of revealing such information compared to hashtags. Second, only high-level information such as descriptive statistics, location, time, and expected number of protesters should be presented to the OSINT team. Third, studies should focus on groups rather than individuals in order to cope with privacy concerns. Also, researchers should consider possible risks of social disruptions before the analysis, for example by using DETA for an ethical data assessment [37].

4 Method

We draw upon Wieringa's design cycle encompassing the activities of designing and investigating a software project [38]. The first phase, known as problem investigation, is centered on clarifying the problem's context. We selected an upcoming social event as a case study and conducted a semi-structured interview to gain insights into the processes currently employed by the OSINT team to gather knowledge about the event. The second phase (treatment design) involved collecting tweets related to the selected event by querying the open Twitter API. Once the tweets were collected, they underwent comprehensive analysis resulting in various artifacts within the problem context. The third phase, treatment validation, involves evaluating the artifacts through a second semi-structured interview with two OSINT analysts.

4.1 Case Study

In order to test and investigate if the needs of the OSINT team can be fulfilled with automated text analysis and visualizations from tweets, a case study is selected that comprises a range of OSINT-specific tasks that come together in a single event.

When selecting an event, multiple factors were taken into consideration. First, the event had to occur in the near future, which would provide the possibility of collecting data in advance while reducing the possibility of other disruptive societal changes that might introduce bias in the tweets. Second, a large number of protesters should be expected to join the protest, at least 10,000 to be classified as large scale [39]. Third, data about related topics should be widely available on Twitter and in other sources (e.g. news articles). Fourth, the event should be based on popular media topics and should be exposed to a large number of people (including people who do not have any interest in the topic). Fifth, the event should be clearly distinguished from other possible events during the same period in time, to prevent information cluttering.

Three large-scale upcoming social events were pre-selected for this study. One of these events was the farmer's protest in the city of The Hague on March 11, 2023, against policies that restrict the use of nitrogen in the agricultural sector. This event was widely reported in the media with an expected number of protests between 10,000 [40] and 25,000 [41, 42]. Second, Extinction Rebellion announced a protest on the same day in The Hague with an expected number of 2,000 protesters [41]. third, the provincial elections which occurred on March 15, 2023. From these upcoming protests, the farmers' protest on 11 March 2023 was eventually selected because it met all the criteria.

The Farmers' Defence Force announced to organize a protest on February 8, 2023, in The Hague [43–46]. The chairman of Farmers Defence Force aimed to organize the largest demonstration ever, addressing various issues including the nitrogen crisis, the gas crisis, subsided houses, and water damage in the province of Limburg.

4.2 Initial Interview

The initial interview lasted one hour and was focused on the inner workings of the current OSINT process. The team interviewed consists of two members, one responsible for managing the tools used by the OSINT team and the other specialized in extracting information from social media, particularly Twitter. Although the OSINT team is open about its workflow, it is not allowed to disclose any operational information. Therefore, undisclosed prerequisite information needed to set up the experiment is missing such as queries and data collection parameters, which are chosen by the authors.

The questions to discuss the OSINT team's workflow were organized into four categories: (1) preparation and relevant topics, (2) sources of information, (3) current tooling and visualization, and (4) result presentation. Preparation and relevant topics were explored using generic questions such as "What is your process for collecting information regarding a large-scale event?". The two interviewees were asked to describe the information they hoped to find and the specific aspects of the information that they were interested in. The interviewees were also questioned about the current visualization techniques that are used in their workflow, especially for upcoming events, the sources upon which this information was based, and the techniques used to analyze the information. Additionally, the interviewee was asked about the process used to verify the accuracy of (visualized) information. Finally, questions were posed regarding the presentation of the findings. The interviewer asked who would be reading the findings, whether or not a specific format for the findings is applied, the level of technical detail that should be included in the results, and how the results would be used.

As the interviews could not be recorded, the interviewer took detailed notes during the interview, including citations of parts of the responses.

4.3 Tweet Collection

To collect tweets related to the upcoming farmer's protest in The Hague, we used the open API of Twitter, with queries based on articles from mainstream media available one month prior to the event. Additional Twitter filters were included to exclude retweets and non-Dutch tweets. Since Twitter is not open about its filter implementations, the results are checked by hand. The query used to collect the tweets is divided into multiple categories, where each category represented a different aspect of the upcoming protest; (1) generic includes abstract queries that are applicable to a wide range of protests (for example "demonstration"), (2) specific includes queries that are specific for the farmers' protest (e.g. "nitrogen mediator"), (3), related groups, with includes the groups directly involved with the protest (e.g. "Farmers Defence Force"), (4) counter groups, who oppose the protest-related groups (e.g. "climate activists"), (5) related topic, includes topics discussed in the media indirectly related to the protest (e.g. "childcare allowance affair"), (6) supporting parties, political parties or spokesman who support the protest (e.g. "PVV"), (7) opposing parties, political parties or spokesman who oppose the protest (e.g. "Kaag").

4.4 Tweet Analysis

Drawing upon the findings obtained from the initial interview, a range of analyses and visualizations are employed in the collected data set of tweets.

Descriptive Statistics Initially, descriptive statistics are applied to the collected data set, including the total number of tweets, the number of tweets collected for each search term, the temporal distribution of tweets, and the frequency of tweets leading up to the demonstration. Additionally, the top ten rankings for retweets, favorite tweets, users with the most followers, highest listed, and favorites counts are also included.

Hashtags The twenty most frequently used hashtags are identified and their counts are presented. A network visualization is then constructed to display these hashtags, with each hashtag represented as a node and the edges indicate co-occurrences of the hashtags within the same tweet. Notably, only hashtags that co-occur in more than 120 tweets are selected to be displayed to improve readability.

Emojis Similar to the analysis of hashtags, emojis are examined by identifying the top twenty most frequently used emojis. Then, these emojis are presented in a graph where a threshold of 40 is set, considering that the frequency of emojis is generally lower than that of hashtags [27].

Websites Websites were extracted from the collected tweets, however, due to Twitter's URL shortening, an additional preprocessing step was necessary to obtain the full URL pointing to an external source. To avoid including media files (e.g., quoted tweets or images/videos in a tweet), a regex filter was applied. The top twenty external sources with the highest number of directed websites are presented.

Sentiment In this study, expressions of discomfort or negative sentiment were predicted using four labeled data sets from previous demonstrations. To create the prediction model, a TF-IDF vectorizer was used on the data, in combination with a standard scaler, since it was successfully applied in previous studies [47–49]. Emojis extracted from tweets are used as bag-of-words on a logistic regression model to predict the probability scores of tweets expressing discomfort, as described in [47–49]. This allowed the tweets to be sorted and provided some extreme examples of highly discomforting tweets versus non-discomforting tweets. Additionally, we calculated a ratio of discomfort versus non-discomfort tweets. For tweets that also contained a hashtag, a discomfort score could be calculated for the hashtag.

4.5 Evaluation Interview

After the collection and analysis of tweets based on the initial interview, an evaluation interview was conducted to discuss the findings. It should be noted that the OSINT team did not disclose any of its own findings with respect to this case study, which is analyzed in preparation for the protest by the OSINT team. More specifically, the OSINT team mainly assessed the presented findings from this study. As information is extracted on multiple levels, various analyses are presented, ranging from descriptive statistics to more advanced techniques such as sentiment analyses. The main focus of this interview is to obtain feedback on these analyses and visualizations.

5 Findings from the Initial Interview

5.1 Information Collection

Information is collected from various open sources, including social media platforms. Most collected information is used to gain insights about specific topics. The OSINT team is not interested in who sends or reads tweets. However, the

OSINT team is interested in content derived from tweets, such as how many people discuss a specific topic and what topics relate to the discussed topic. Furthermore, the OSINT team wants to know within what context a certain topic is used. The same applies to topics related to upcoming social events, where some messages might lead to polarization, whereas other messages mainly serve the purpose of entertainment. Some people violate the law by posting offensive or even racist messages. In such cases, permission can be obtained to target a specific person, but this does not happen very often and is not considered a task of the OSINT team. The main aspects the OSINT team is interested in are related to obtaining a global view of matters discussed in relation to large-scale social events. For example, when people are attempting to move a group by proposing specific actions, such as "Let's move to the station". Another aspect of information includes expressions of discomfort towards a specific (political) person. For example, "The president is unsuited and should retire". When a large group of people is critical of a public spokesman, some of these commands may result in violent behavior, and this may be a reason to safeguard this person.

5.2 Current Analysis and Visualisations

The OSINT team is constrained to the utilization of commercially available tools. These tools are primarily not designed for intelligence purposes, as they are tailored for tasks such as brand or product reception analysis. Therefore, analysts are constantly seeking alternative methods to extract useful information. Currently, OSINT analysts have access to several information sources, including temporal fluctuations in the number of tweets and retweets, which can be utilized to identify popular topics. Depending on the requests, word clouds are generated to visualize common words and textual representations, such as emojis. However, accessing raw data can be challenging, as external parties perform the analysis and provide only the results, making it difficult to verify their accuracy. Moreover, many insights are often derived through indirect analysis. Analysts typically gauge the potential for violent behavior by scrutinizing the tone and language used in messages related to specific topics. However, relying on analysts' "gut feeling" lacks sufficient support for evidence-based decision-making, especially when scarce and costly police resources are needed to handle an upcoming event. In such cases, only hard evidence can justify the allocation of manpower. However, a lack of hard evidence to support analysts' suspicions regarding potential unrest related to upcoming events is not unusual. As a consequence, in the past, the police have either overreacted or reacted too late to contain and possibly prevent riots that analysts had already anticipated. This suggests a need to develop testable and reproducible results that can be generalized to other upcoming events.

5.3 Expected Results

The insights obtained through OSINT analysis are often utilized by decision-makers who oversee the deployment of police resources. Unlike OSINT analysts, these individuals typically lack extensive technical backgrounds and seek clear and concise reports that can be quickly processed due to time constraints. As a result, the research questions handled by the OSINT team are usually quite specific, since specific research questions usually lead to concise reports. A standard format for presenting results is not available for the OSINT team, however, analysts are typically limited to using no more than one A4 paper.

6 Results

6.1 Number of Tweets

A total of 340,658 tweets were collected prior to the demonstration. Due to the consistent use of queries during the collection period, time series could be plotted, see Figure 6.1. Based on the initial interview for the purpose of this study, the following observations are considered to be possibly relevant for OSINT analysts.

- On February 20, the number of tweets rapidly increased to a peak at 10:00 AM and remained elevated for the rest of the day.
- On February 21, a slight increase in activity was visible around 9:30 AM. Although the pattern was similar to other tweets, the activity remained slightly elevated throughout the day.
- On February 26, increased activity in the morning occurred, which rapidly decreased and returned to normal levels around 12:00 PM.
- On March 5, at the end of the morning, extra activity is visible which is peaking at around 10:00 AM, but rapidly decreasing thereafter. From 3:00 PM onwards, the activity was not significantly different from normal.
- On March 7, slightly more activity in the afternoon and evening can be detected in comparison to patterns observed in other tweets. Normal behavior usually shows a slight increase in the evening, but on the 7th of March, this was particularly evident.

6.2 GEO Locations

From the collected tweets, 246 contained geographical information (.07% of the total number of tweets). The extracted locations are plotted on a map, as shown in Figure 6.2. The majority of tweet activity was concentrated in large cities (e.g. Amsterdam and Rotterdam). We observe an increased level of activity in the small subsidy areas such as Veenendaal, Doetinchem, and Deventer.

6.3 Top Tens

Due to privacy concerns, it is not feasible to disclose the users who have received the highest number of retweets or favorites. Nevertheless, the most frequently occurring themes in the highly retweeted messages include criticism of Klaver and Kaag, who are both active members of the Dutch parliament, as well as the

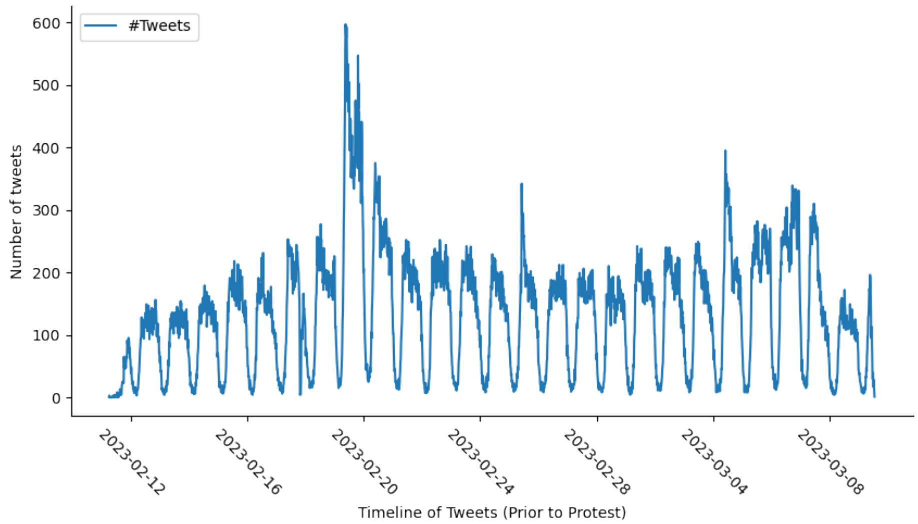

Fig. 1. Time series on the number of collected tweets prior to the upcoming farmers' protest. The x-axis represents a timeline of tweets, with each point or interval corresponding to a specific time frame. The tweets are binned in 15-minute intervals, showing the number of tweets (y-axis) created at those particular times before the protest occurred.

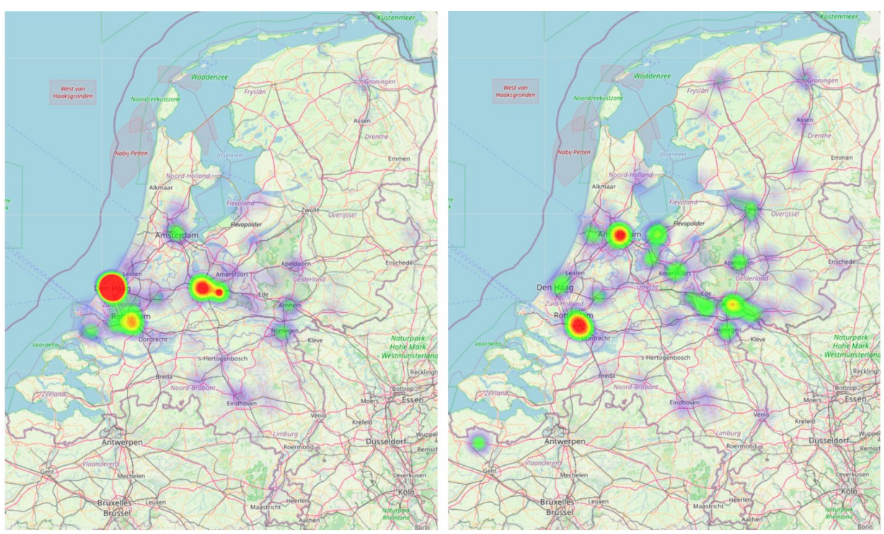

Fig. 2. Heatmap displaying the Twitter activity in the Netherlands. The left map displays the number of protest-related tweets over a two-year time span (2020 - 2022). The tweet activity related to the farmer's protest is displayed on the right map. The data are scaled while other parameters (such as zoom level, pixel ratio, and view port) are kept constant between both maps.

issue of rising prices. The top ten retweets also cover topics such as COVID-19 conspiracy theories and the nitrogen crisis. Tweets that have received the most favorites also predominantly contain criticism of Klaver and Kaag, which extend to societal issues such as housing shortages and price increases. We observe considerable differences in the number of collected tweets per query, ranging from 2,296,206 related to "Geert Wilders" versus "farmers' point of view" with only 7 hits. The topmost hits include "Forum for Democracy' (175,554), "Kaag" (174,960), "Police" (130,235), and "Nitrogen" (119,836).

6.4 Hashtags

The hashtag network presented by Figure 6.4, contains the top twenty most used hashtags and their connections in terms of coincidence across tweets. The hashtag "stem ze weg" (translates to "vote them away') is frequently associated with various political parties such as "D66", "PVV", "BBB", and "FvD", and individuals like "Kaag" (a former parliament member who represented D66), and in lesser extent with the parties "VVD" and "CDA" and former prime minister "Rutte". Additionally, related political topics such as the "toeslagenaffaire" (translates to "childcare allowance scandal'), "stikstof" (translates to "nitrogen'), and "verkeizing" (translates to "the March 15, 2023 elections') are notably present in connection to hashtags related to "D66" and "PVV". This suggests a strong relationship between political parties "FvD" and "JA21" and FvD's leader "Baudet".

6.5 Emojis

A network of the top twenty most used emojis is displayed in Figure 6.5. The most used emojis express amusement (laughing emojis and a clown), thumbs up and winking emojis also appear in the top ten. Additionally, the top ten include emojis expressing discontent (such as the thinking face, face vomiting, and angry expressions). The connections between emojis show a cluster of expressions of amusement (monkey, grinning face with sweat, and various laughing expressions). Surrounding this cluster there are emojis with expressions of frustration (rolling eyes, angry or vomiting expressions). In addition to the amusement cluster, there are thematic emojis (fist, red heart, Dutch flag) that are not directly related to either amusement or criticism clusters.

6.6 Websites

The analysis reveals that news-related websites are the primary sources of reference, wherein a news outlet with right-wing political inclinations (e.g. "Nieuwrechts') appears to have the highest popularity (with 158 references). This is followed by multiple mainstream media sources such as nu.nl, Volkskrant, Telegraaf, and NOS. Specific articles from these media sources are linked less frequently. The themes commonly associated with the linked articles include issues surrounding nitrogen, farmers' protests, and voting guidelines.

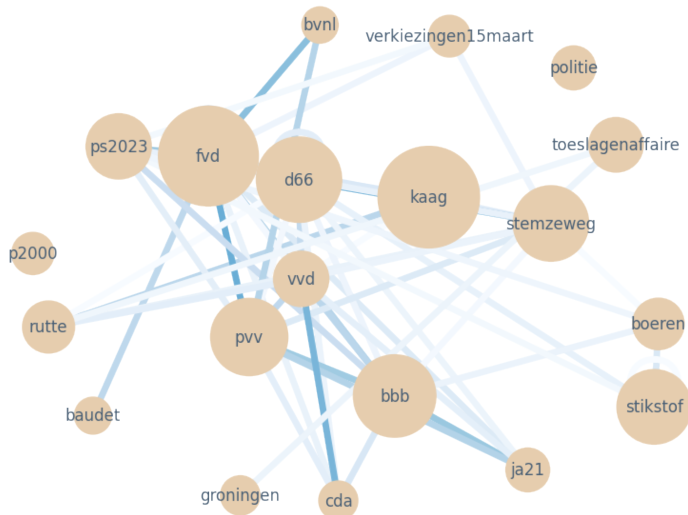

Fig. 3. Pairwise co-occurrence network most common hashtags, where the size of the node represents the number of occurrences (e.g. large nodes correspond to frequently used hashtags), and the edges represent the co-occurrences of the hashtags, where a darker color indicates a greater co-occurrence.

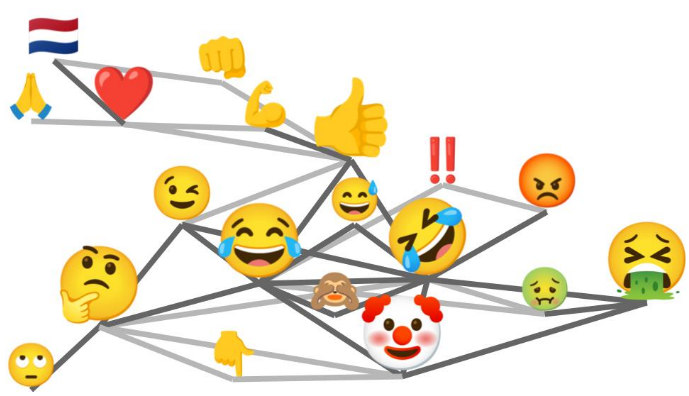

Fig. 4. Network of the most common emojis. The size depends on the occurrence of the emojis. The smallest emojis occur less than 400 times, the middle sizes occur more than 400, but less than 600 times, and the largest emojis occur more than 600 times. The edges indicate the number of tweets where the emojis co-exist, where a light color indicates a co-existence of less than 40 times, and the dark color indicates that the two connected emojis occur together in at least 40 tweets.

6.7 Sentiment

Applying the logistic regression model based on a vectorized representation of emojis in tweets from previously labeled tweets resulted in a model to estimate sentiment. Themes surrounding the most intense expressions of discomfort include criticism of Kaag, Paternotte, D66, and VVD. Additionally, various political themes reoccur, such as nitrogen, the distribution of poverty, and refugees. Hashtags associated with high levels of discomfort encompass themes such as nitrogen, refugees, politics, and conspiracy theories, as summed in Table 1.

Table 1. Most negative hashtags per category. Terms are translated from Dutch, the original are: (1) stront-moeras, (2) zurepoep, (3) sourplanet, (4) sourpeople, (5) sourdutchies, (6) millieu, (7) asielstroom, (8) illigal, (9) kaagplaag, (10) overkill, (11) pschycrembel, (12) greatreset, (13) openbaring, and (14) nord stream 2 aanval. Note that the spelling of some hashtags is incorrect.

Theme	**Most negative hashtags**
Nitrogen	"shit swamp"[1], "sour poop"[2], "sourplanet"[3], "sourpeople"[4], "sourdutchies"[5], "environment"[6]
Refugees	"asylum influx"[7], "illegal"[8]
Politics	"kaagplague"[9], "overkill"[10]
Conspiracy theories	"pschycrembel"[11], "greatreset"[12], "revelation"[13], "Nord Stream 2 attack"[14]

7 Evaluation

The evaluation interview conducted with the OSINT team revealed that certain aspects were consistent with OSINT analysts' research objectives, while others raised concerns. Each visualization was assessed individually to provide a comprehensive evaluation. Overall, the information presented was found to be in line with the findings of the OSINT team. The results were described as "clear" and "actionable", and certain aspects even offered new perspectives for the OSINT analysts. However, it was noted that not all information discovered by OSINT analysts was reflected in this analysis.

GEO Location Drawing reliable conclusions from the geographic information of tweets is challenging due to the limited number of geolocated tweets available, as they represent less than .10% of the total number of tweets. Despite this limitation, the presented findings remain valuable, especially in identifying activities originating from rural areas with small populations. Such observations may strengthen the assumption that farmers are engaged in the topics related to the upcoming event, warranting further investigation into the sentiment towards these subjects. Future research endeavors could explore the possibility of extracting geographic information from the tweet text.

Number of Tweets The analysis has revealed peaks and valleys in Twitter activity that raise further questions about the reasons behind increases or decreases in activity. Currently, OSINT analysts only provide information on the volume of tweets, but comparing activity on different days can offer valuable insight into Twitter behavior relative to other days, a feature not currently available in the tools that are used. As such, it is recommended that future investigations explore the feasibility of establishing a baseline against which the number of tweets can be contextualized.

Top Ten Tweets The top ten tweets offer insights, but the OSINT analysts stated that their usage is subject to legal constraints. When the level of granularity is such that it can be traced back to a single tweet or user, the privacy of that user may be compromised. Nevertheless, general descriptions of the most frequent topics among the top favorited and retweeted messages are "useful in specific cases" when safety is preferred over privacy, according to the OSINT team.

Hashtags The present tooling includes the use of word clouds to visualize topics and their respective frequency of occurrence. However, the proposed approach is more comprehensive, since it involves connecting hashtags that occur in the same tweets, allowing for a deeper understanding of the relationships between topics. By analyzing these connections, it is possible to identify distinct clusters of common topics, as well as examine the connections between a singular topic and a variety of other less common topics. Such visualizations offer valuable information for discovering the underlying relationships and connections between different topics.

Emojis The use of emojis in isolation may not provide substantial insights; however, when combined with the previously mentioned hashtag network, this approach may hold potential value. Typically, textual constructs, such as emojis, are excluded from OSINT analyses, because the same sentiment can be extracted from the surrounding text alone.

Websites During the COVID-19 pandemic, websites were frequently used to investigate the spread of false information on social media platforms. The presented Table of websites illustrates the prevalence of certain platforms in the analyzed tweets, which could potentially shed light on the demographics of the participating members.

Sentiment The expression of discomfort is a useful metric for gaining insights into groups that require additional attention. Typically, models provide only a score without further explanation. By presenting example tweets along with their sentiment scores, OSINT analysts can extract additional information, enhancing the model's transparency.

8 Concluding Remarks

This study displays multiple visualizations extracted from tweets related to an upcoming protest with the aim of strengthening the information position of the OSINT team. The presented results show potential for future research in a more practical setting. However, we acknowledge several limitations to this study. During the evaluation interview, the OSINT analysts identified some missing elements, i.e. more GEO data is required to display a reliable map, the number of tweets should be presented with a baseline, and the combination of sentiment, hashtags, and emojis might strengthen the visualizations. Moreover, not all information discovered by the OSINT team was reflected in this analysis. This could be due to limitations in the proposed techniques or simply because the data was not present in the tweets.

We did not use a framework or pre-established questionnaires, such as the System Usability Scale (SUS) to contextualize our approach within existing techniques and to assess the usability of our method, since the scope of this research is mainly focused on exploring the information needs rather than presenting this information, which will serve as a focal point for future research. Another limitation is the scope of our study is restricted to the involvement of only two OSINT analysts. Consequently, the insights and feedback provided may be influenced by the unique perspectives and expertise of these two individuals, potentially not fully representing the entire landscape of OSINT practices within the Dutch police. It's worth noting that the interviewed analysts are affiliated with the national police, which deals with nationwide events impacting the entirety of the Netherlands. Nevertheless, the Dutch police encompass various specialized areas within OSINT, despite sharing common ground across these disciplines.

Future research should be addressed to discover with information is presented in tweets, and what information should be extracted from other sources. From an ethical and legal perspective, there is still room for improvement on how to cope with the constraints set to the OSINT analysts and whether the possibilities of these visualizations extend the power of the OSINT team to a level that is no longer acceptable to the public. The visualizations provide valuable information for discovering the underlying relationships and connections between different topics. Future research in this area may benefit from the integration of additional visualizations, such as a multi-color schema to address sentiment scores per topic. The utilization of particular emojis may indicate group characteristics that would help to identify violent groups, hence further research could explore this direction. The sentiment scores improve the models' transparency, however, this measure's efficacy is limited when used in isolation from other visualizations, such as hashtag networks. One potential approach is to highlight the most commonly used hashtags in the network and allow users to click through them to view tweets with underlying sentiments.

Acknowledgments

This work was supported in part by the Dutch Police, who provided insights into their open-source intelligence work. Our gratitude goes especially to the OSINT team, who provided valuable insights into their activities through the two interviews.

References

1. David S Meyer and Sidney Tarrow. A movement society: Contentious politics for a new century. *The Social Movement Society: Contentious Politics for a New Century*, pages 1–28, 1998.
2. José-Manuel Sabucedo, Cristina Gómez-Román, Mónica Alzate, Jacquelien van Stekelenburg, and Bert Klandermans. Comparing protests and demonstrators in times of austerity: Regular and occasional protesters in universalistic and particularistic mobilisations. *Social Movement Studies*, 16(6):704–720, 2017.
3. Kathryn Rose Ranney. *Social media use and collective identity within the occupy movement*. PhD thesis, Honolulu, University of Hawaii at Manoa, 2014.
4. Mark Richard Bauermeister. Social capital and collective identity in the local food movement. *International Journal of Agricultural Sustainability*, 14(2):123–141, 2016.
5. Sarah A Soule and Susan Olzak. When do movements matter? The politics of contingency and the equal rights amendment. *American Sociological Review*, 69(4):473–497, 2004.
6. Marco Giugni. How social movements matter: Past research, present problems, future developments. *How Social Movements Matter*, pages xiii–xxxiii, 1999.
7. Albert J Meijer. *Politie en sociale media. Van hype naar onderbouwde keuzen.* Reed Business Information, 2013.
8. Emma Tonkin, Heather D Pfeiffer, and Greg Tourte. Twitter, information sharing and the London riots? *Bulletin of the American Society for Information Science and Technology*, 38(2):49–57, 2012.
9. Rob Procter, Farida Vis, and Alex Voss. Reading the riots on Twitter: Methodological innovation for the analysis of big data. *International Journal of Social Research Methodology*, 16(3):197–214, 2013.
10. Lauren M Burch, Evan L Frederick, and Ann Pegoraro. Kissing in the carnage: An examination of framing on Twitter during the Vancouver riots. *Journal of Broadcasting & Electronic Media*, 59(3):399–415, 2015.
11. Lei Yang, Tao Sun, Ming Zhang, and Qiaozhu Mei. We know what@ you# tag: Does the dual role affect hashtag adoption? In *Proceedings of the 21st international Conference on World Wide Web*, pages 261–270, 2012.
12. Carolin Gerlitz and Bernhard Rieder. Mining one percent of Twitter: Collections, baselines, sampling. *M/C Journal*, 16(2), 2013.
13. Babak Amiri, Ramin Karimianghadim, Navid Yazdanjue, and Liaquat Hossain. Research topics and trends of the hashtag recommendation domain. *Scientometrics*, 126:2689–2735, 2021.
14. Sanjaya Wijeratne, Lakshika Balasuriya, Amit Sheth, and Derek Doran. Emojinet: Building a machine readable sense inventory for emoji. In *Social Informatics: 8th International Conference, SocInfo 2016, Bellevue, WA, USA, November 11-14, 2016, Proceedings, Part I 8*, pages 527–541. Springer, 2016.
15. Sanjaya Wijeratne, Lakshika Balasuriya, Amit Sheth, and Derek Doran. Emojinet: An open service and API for emoji sense discovery. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11, pages 437–446, 2017.
16. Philippe Kimura-Thollander and Neha Kumar. Examining the “global” language of emojis: Designing for cultural representation. In *Proceedings of the 2019 CHI conference on Human Factors in Computing Systems*, pages 1–14, 2019.

17. Bethany Berard. I second that emoji: The standards, structures, and social production of emoji. *First Monday*, 2018.
18. Nikola Ljubešić and Darja Fišer. A global analysis of emoji usage. In *Proceedings of the 10th Web as Corpus Workshop*, pages 82–89, 2016.
19. Yuxiao Chen, Jianbo Yuan, Quanzeng You, and Jiebo Luo. Twitter sentiment analysis via bi-sense emoji embedding and attention-based LSTM. In *Proceedings of the 26th ACM International Conference on Multimedia*, pages 117–125, 2018.
20. fastText. URL: https://fasttext.cc/ (accessed: 12 July, 2023). [Online].
21. Halley Fede, Isaiah Herrera, SM Mahdi Seyednezhad, and Ronaldo Menezes. Representing emoji usage using directed networks: A Twitter case study. In *International Conference on Complex Networks and their Applications*, pages 829–842. Springer, 2017.
22. Tauhid R Zaman, Ralf Herbrich, Jurgen Van Gael, and David Stern. Predicting information spreading in Twitter. In *Workshop on Computational Social Science and the Wisdom of Crowds, Nips*, volume 104, pages 17599–601. Citeseer, 2010.
23. Liangjie Hong, Ovidiu Dan, and Brian D Davison. Predicting popular messages in Twitter. In *Proceedings of the 20th International Conference Companion on World Wide Web*, pages 57–58, 2011.
24. Intzar Ali Lashari and Uffe Kock Wiil. Monitoring public opinion by measuring the sentiment of retweets on Twitter. In *3rd European Conference on Social Media*, pages 153–161. Academic Conferences and Publishing International, 2016.
25. Yuhang Zhou and Wei Ai. # Emoji: A study on the association between emojis and hashtags on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 1169–1180, 2022.
26. Mayank Kejriwal, Qile Wang, Hongyu Li, and Lu Wang. An empirical study of emoji usage on Twitter in linguistic and national contexts. *Online Social Networks and Media*, 24:100149, 2021.
27. Mark Alfano, Ritsaart Reimann, Ignacio Ojea Quintana, Anastasia Chan, Marc Cheong, and Colin Klein. The affiliative use of emoji and hashtags in the Black Lives Matter movement in Twitter. *Social Science Computer Review*, page 08944393221131928, 2022.
28. Su Mon Kywe, Tuan-Anh Hoang, Ee-Peng Lim, and Feida Zhu. On recommending hashtags in Twitter networks. In *Social Informatics: 4th International Conference, SocInfo 2012, Lausanne, Switzerland, December 5-7, 2012. Proceedings 4*, pages 337–350. Springer, 2012.
29. Yabing Liu, Chloe Kliman-Silver, and Alan Mislove. The tweets they are a-changin': Evolution of Twitter users and behavior. In *Eighth International AAAI Conference on Weblogs and Social Media*, 2014.
30. Mengdi Li, Eugene Ch'ng, Alain Yee Loong Chong, and Simon See. Multi-class Twitter sentiment classification with emojis. *Industrial Management & Data Systems*, 2018.
31. Helena Webb, Marina Jirotka, Bernd Carsten Stahl, William Housley, Adam Edwards, Matthew Williams, Rob Procter, Omer Rana, and Pete Burnap. The ethical challenges of publishing Twitter data for research dissemination. In *Proceedings of the 2017 ACM on Web Science Conference*, pages 339–348, 2017.
32. Utrecht Universiy Institute of Information and Computing Sciences. Ethics and Privacy. URL: https://www.uu.nl/en/research/institute-of-information-and-computing-sciences/ethics-and-privacy (accessed: 21 July, 2023), 2023. [Online].

33. Luke Sloan, Jeffrey Morgan, Pete Burnap, and Matthew Williams. Who tweets? deriving the demographic characteristics of age, occupation and social class from Twitter user meta-data. *PloS One*, 10(3):e0115545, 2015.
34. Utrecht University. Research Data Management Support. URL: https://www.uu.nl/en/research/research-data-management/guides/policies-codes-of-conduct-and-laws (accessed: 12 July, 2023), 2023. [Online].
35. Bart W Schermer, Dominique Hagenauw, and Nathalie Falot. Handleiding Algemene verordening gegevensbescherming en Uitvoeringswet Algemene verordening gegevensbescherming, January 22, 2018. URL: https://www.rijksoverheid.nl/onderwerpen/privacy-en-persoonsgegevens/documenten/rapporten/2018/01/22/handleiding-algemene-verordening-gegevensbescherming (accessed: 16 February, 2023).
36. Caitlin M Rivers and Bryan L Lewis. Ethical research standards in a world of big data. *F1000Research*, 3:38, 2014.
37. Mirko Tobias Schäfer, Aline Franzke, Gemeente Utrecht, and Redmar Fransen. De ethische data assistent (DEDA). https://deda.dataschool.nl/wp-content/uploads/sites/415/2022/11/DEDA-NL.handbook.V3.1.pdf, 2022.
38. Roel J Wieringa. *Design science methodology for information systems and software engineering*. Springer, 2014.
39. Michael Biggs. Size matters: Quantifying protest by counting participants. *Sociological Methods & Research*, 47(3):351–383, 2018.
40. AD. Verkeershinder verwacht op 11 maart vanwege demonstraties. URL: https://www.ad.nl/den-haag/verkeershinder-verwacht-op-11-maart-vanwege-demonstraties~a17f2841/ (accessed: 11 March, 2023), 2023. [Online].
41. RTL Nieuws. Den Haag wacht gespannen protestdag met boeren en klimaatactivisten: Vijf vragen. URL: https://www.rtlnieuws.nl/nieuws/nederland/artikel/5370786/demonstranten-extinction-rebellion-a12-boerenprotest (accessed: 11 March, 2023), 11 March, 2023. [Online].
42. VRT. Duizenden boeren protesteren in Den Haag tegen stikstofbeleid, 700 klimaatactivisten opgepakt na protestactie op snelweg. URL: https://www.vrt.be/vrtnws/nl/2023/03/11/ondanks-verbod-trekken-tractoren-in-kolonne-naar-den-haag-voor-p/ (accessed: 21 July, 2023), 11 March, 2023. [Online].
43. NRC. Farmers Defence Force wil op 11 maart weer demonstreren in Den Haag. URL: https://www.nrc.nl/nieuws/2023/02/08/farmers-defence-force-wil-op-11-maart-weer-demonstreren-in-den-haag-a415\\6582 (accessed: 12 July, 2023), 8 February, 2023. [Online].
44. Hart van Nederland. Boeren kondigen 'grootste demonstratie ooit' aan op 11 maart in Den Haag. URL: https://www.hartvannederland.nl/nieuws/politiek/boeren-kondigen-grootste-demonstratie-ooit-aan-op-11-maart-in-den-haag (accessed: 12 July, 2023), 8 February, 2023. [Online].
45. DVHN. Den Haag zet zich schrap: stad dreigt op 11 maart volledig vast te lopen door protesten. URL: https://dvhn.nl/binnenland/Den-Haag-zet-zich-schrap-stad-dreigt-op-11-maart-volledig-vast-te-lopen-door-\\protesten-28282312.html (accessed: 12 July, 2023), 6 March, 2023. [Online].
46. Mark Jongeneel. Burgemeester Jan van Zanen (VVD) bereidt zich voor op mogelijke verkeers- en protestchaos op 11 maart in Den Haag. URL: https://www.dagelijksestandaard.nl/binnenland/burgemeester-jan-van-zanen-vvd-bereidt-zich-voor-op-mogelijke-verkeers-en-protest\\chaos-op-11-maart-in-den-haag (accessed: 12 July, 2023), 2023. [Online].

47. Satyendra Singh, Krishan Kumar, and Brajesh Kumar. Sentiment analysis of Twitter data using TF-IDF and machine learning techniques. In *2022 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COM-IT-CON)*, volume 1, pages 252–255. IEEE, 2022.
48. Fika Hastarita Rachman et al. Twitter sentiment analysis of Covid-19 using term weighting tf-idf and logistic regression. In *2020 6th Information Technology International Seminar (ITIS)*, pages 238–242. IEEE, 2020.
49. Md Rakibul Hasan, Maisha Maliha, and M Arifuzzaman. Sentiment analysis with NLP on Twitter data. In *2019 International Conference on Computer, Communication, Chemical, Materials and Electronic Engineering (IC4ME2)*, pages 1–4. IEEE, 2019.