UTRECHT UNIVERSITY

MASTER THESIS

Unveiling protest-related social networks on Twitter

A comprehensive analysis of topic differences and temporal patterns within protest-related social networks in the Netherlands.

Author:
Jesse BORGHARDT
6458114

1st Supervisor:
Prof. Dr. Remco VELTKAMP
2st Supervisor:
Dr. Christof van Nimwegen
Company Supervisor:
Laurens Müter

Faculty of Science
Graduate school of Natural Science
Business Informatics

July 18, 2023

UTRECHT UNIVERSITY

Abstract

Faculty of Science
Business Informatics

Master of Science

Unveiling protest-related social networks on Twitter

by Jesse BORGHARDT

The number of protests in the Netherlands has risen significantly in recent years, posing challenges for the Dutch police to ensure safety during these events. The increasing frequency of unannounced protests and the personnel shortages within the organization further exacerbate the complexities of ensuring public safety during such events. This necessitates research to better comprehend the motivations behind protests and effectively prepare for future events.

This study addresses this need by utilizing open source Twitter data to perform social network analyses. The research is conducted in three phases: an exploratory phase, a method development phase, and an evaluation phase. In the exploratory phase, the possibilities of using social network analyses within the context of protest-related social networks are investigated. Building on the experiences from this phase, a method called PReSNA is developed for analyzing protest-related social networks. Following the design science research approach, the PReSNA method is then evaluated using a separate data set covering two large protests in March 2023.

The findings reveal significant differences in hashtag usage prior to protests, providing valuable insights into the possibility to detect topics that can trigger future protests. Moreover, significant differences are observed in the characteristics of groups engaged in protest-related Twitter conversations, shedding light on the factors driving activity within these groups.

The results of this study demonstrate the value of open source data in gaining insights into protest behavior. Furthermore, these findings highlight the potential for future research within this area.

Overall, this research contributes to the understanding of protest-related social networks and offers a methodological framework for analyzing such networks. By leveraging open source data and social network analysis techniques, this study provides valuable insights for researchers and practitioners in the field of protest management.

Keywords: Protest motivation, Group characteristics, Dutch protests, Social network analysis, Twitter, Hashtags

Contents

List of Abbreviations

D	**D**egree
ES	**E**dge **S**ource
ET	**E**dge **T**arget
EW	**E**dge **W**eight
MC	**M**odularity **C**lass
OSNs	**O**nline **S**ocial **N**etworking **s**ites
OSINT	**O**pen **S**ource **IN**ntelligence **T**eam
PICO	**P**opulation **I**ntervention **C**omparison **O**outcome measures
SNA	**S**ocial **N**etwork **A**nalysis
WD	**W**eigthed **D**egree

Chapter 1

Introduction

The number of protests in the Netherlands is rising each year. In 2021, it became clear that the number of protests in the four biggest cities of the Netherlands has almost doubled in the last five years (Pols, 2021). Furthermore, an interesting trend has emerged in protests that are not announced beforehand to the municipality (Groot, 2022). This causes the police to react instead of being prepared in advance. While officers are always available for unexpected situations, the lack of prior notification means that the police have less information regarding the size and area of the protest.

In addition to unannounced protests posing a challenge for the police, announced protests may proceed differently than previously announced. Factors such as weather, publicity and accessibility impact the number of protesters showing up (Stekelenburg, Klandermans, and Dijk, 2011). Considering these factors is crucial for the police, as a protest with limited media attention but organized on a sunny weekend day can still attract a large attendance. Conversely, a protest organized on a business day with extensive media coverage can also result in large turnout. These fluctuating factors make it difficult for the police to predict attendance in advance, necessitating real-time estimations based on observations of the protest, including estimations of surface area and density. Publicly available tools such as mapchecking.com are used by the police for a rough estimation of the protester count, relying on real-time observations due to their reliance on crowd density (NOS, 2021). These tools, however, are not applicable to estimate protest attendance before the start of the protest.

Additional to tools that use maps to estimate the number of participators, research has shown that other methods can be used to come to similar results. Sobolev et al. (2020) for instance, demonstrated that publicly available pictures of protests shared via social media, as well as cellphone location data, could be used to estimate protest size. The researchers argued however, that multiple methods result in multiple outcomes, which allows newspaper to pick the number that best fits their story.

The Dutch police is also interested in solutions that can help better estimate the size of a protest, but before implementing such solutions, the organization seeks a better understanding of the factors that contribute to protest attendance.

This study, therefore, is conducted in collaboration with the Dutch police, with the goal of identifying the factors that affect protest motivation, with a specific focus on media and social media.

While there is extensive research conducted on protest motivation, studies that incorporate media and social media have primarily focused on the role of social media in organizing protests or its impact on societal polarization. The primary objective of this study is therefore to identify differences in protest-related social networks that could help predicting protest motivation.

1.1 Problem statement

With the increasing number of protests, the Dutch police is forced to have more officers available. However, police forces in the Netherlands are facing significant personnel shortages. The police unit The Hague for instance, was short of 383 officers in 2021, which accounted for 6.7% of the intended operational formation (Enthoven and Hofstede, 2022). This shortage necessitates the deployment of personnel from other units to assist during large protests. From January 1st to June 15th, 2021, a total of 38,500 hours of assistance was provided to Amsterdam by other units solely for protests (Enthoven and Hofstede, 2022). In cities where protests are frequent, protest estimations can play a crucial role in distributing an appropriate number of officers to the right locations (Groot, 2022). This approach helps avoid an excess of officers in one area while facing shortages in another. Considering the primary objective of the police, which is to maintain public order, it is evident that ensuring an adequate number of officers are deployed at specific spots is essential (Nassauer, 2016). Therefore, reliable protest estimations are needed to enable the responsible authorities to plan and prepare more effectively in advance.

As mentioned in the introduction, there are external factors, such as weather, publicity, and availability, that can be utilized to predict protest attendance before the commencement of a protest (Walgrave et al., 2013). However, there also exist intrinsic protest motivators, such as grievances and group-based anger, which have a more substantial impact on protest motivation. Both intrinsic motivation and external factors will be incorporated in this study as relationships between factors may exist. The intrinsic motivation to protest varies for each individual and this is why a data source where human emotions and freedom of speech are commonly expressed is essential.

Twitter is chosen for this purpose, as it is a platform that highly values freedom of speech (Bahrami et al., 2018). Moreover, Twitter provides a practical data source as tweets are publicly available. Recognizing the usefulness of open source data, the Dutch police has already done collaborations with instances to study the how this could be used (Statistiek, 2023). Identifying social unrest can be relevant for the police as it may be the basis for protest. Tools such as Meltwater and Audiense are already capable of tracking the impact of communication strategies used by businesses by analyzing thousands of news items and social media channels (Meltwater, 2023; Audiense, 2023). Similar tools, if developed for the police, could aid in identifying social media activities before, during, and after protests. This can help clarifying why protest arose and what social networks were tweeting about during certain protests.

In addition to the identification of topics that cause social unrest, the police is also interested in understanding factors that contribute to the motivation to participate in protests. With the aim to provide insights in human behavior, this study will primarily focus on the human side of protest motivators, rather than the technical aspects of tool development.

1.2 Research objectives and aim

This research is conducted through a collaboration between the Utrecht University and the Dutch police. The Dutch police aims to utilize Twitter to gain insights into topics that are associated with protests. In addition to the insights in the topics, the police is also interested in the social context of groups that participate in protests, including their social network, roles within those networks, and media usage. Therefore, the primary aim of this study is utilizing social networks analysis to identify temporal patterns and network differences between topics and groups.

As mentioned in the introduction (1, previous research has explored the intersection of protests and (social) media, primarily focusing on the role of social media in organizing protests and the role of media and echo chambers in societal polarization (Quattrociocchi, Scala, and Sunstein, 2016; Wieringa et al., 2018; González-Bailón et al., 2011).

Social media as a contributing factor to societal polarization, makes it an indirect contributing factor to the rise of extremism, as discussed by Sunstein (1999). The link between social media networks and motivation for protests serves as the foundation for this study. It is evident that multiple factors play a role in motivating protests, making social network analysis a valuable tool for gaining insights into these motivators. In order to provide guidance for this study and future research, a method has is developed with the primary objective of analyzing protest-related social networks. As this method can provide significant guidance in future studies, the main objective of this study is defined as:

Develop a method to analyze protest-related social networks with Twitter data

To achieve this research objective, it is necessary to formulate research questions (RQs). Given the extensive existing research on protest motivation, conducting a thorough review of the current state of the art is crucial. This review will provide deeper insights in the factors that contribute to protest motivation and facilitate an understanding of how social media influences these identified factors. The first research question is therefore:

- **RQ 1** "Which motivating factors to protest can be identified from existing literature?"

In order to answer this research question, recent papers will be reviewed and an overview of the outcomes will be presented in the related work section. When this question is answered, the main protest motivators are identified and more specific, the role of media, social media, and social networks is known. This allows for the analysis of the protest-related social network to start.

Within these social networks, a lot of information and news is shared with each other (Larson et al., 2019). It is interesting to review the information that is shared within these social networks to identify possible differences between social networks. This will be done by analyzing the hashtags that are used as this will provide insight into the topics that are discussed at a certain time. In addition to the differences over time, the differences in hashtag use per identified group are analyzed. This results in the following research question:

- **RQ 2** "What differences in hashtag usage can be used to analyze protest-related social networks?"

When a topic analysis is done with the use of hashtags, a clustering will be made within the network of users that are related to protest. This allows for a more in depth analysis of how these motivators can be identified in social networks on Twitter. Literature found for instance, that a group of protesters consists of a diverse set of individuals that have a varying background (Chenoweth et al., 2022). This raises the question what distinguishes them from the persons that do not protest. For this question, we are especially interested in the differences or similarities between media consumption and the following research question was therefore formulated:

- **RQ 3** "What differences in media consumption can be identified when comparing clusters of a protest-related social network?"

This question will be addressed through an exploratory data analysis of protest-related social networks on Twitter, utilizing network visualization and exploration software. The objective behind generating these visualizations is to facilitate the examination of disparities that may arise when comparing media usage among distinct clusters within a protest-related social network.

In addition to media usage, literature has shown that social networks are one of the most important factors that contribute to the motivation for and awareness of protest (Larson et al., 2019). It is therefore important to investigate what individuals are part of which social networks and if shifts between these networks can be seen over time. Therefore, the following research question was formulated:

- **RQ 4** "What differences in group characteristics can be identified when comparing clusters of a protest-related social network?"

1.3 Added Value

1.3.1 Practical contribution

As mentioned in the chapter 1, the Dutch police are facing personnel shortages which forces them to make decisions on where to deploy their officers. When specifically looking at protest, this becomes even more relevant because large numbers of officers are involved in a single protest. To estimate the size of a protest, factors such as the weather and accessibility should be considered. However, the Dutch police want to go further by studying protest motivators and especially social networks with the use of Twitter data. By doing so, the organization hopes to get more insights into the reasoning behind protests. More specifically, the police is interested in the opinion of citizens towards protest-related topics. Eventually, they want to develop a tool which automatically provides overviews of protest-related "hot-topics" with overall opinions of the tweeters integrated in it by making use of text analysis algorithms that can identify expressions of discomfort.

Before this tool is developed, a better understanding of the motivation to protest is necessary, and additionally the possibilities to use social media for this purpose should be studied more extensively. This study contributes to this goal by identifying factors that affect protest motivation, and on top of that by identifying the possibilities to use social networks to predict protest motivation. In addition, the knowledge that is obtained with this study will be discussed with employees of the Dutch police to contribute to the development of the tool that they want to create.

1.3.2 Academic contribution

As earlier mentioned, research on protest motivation is already conducted. This study, however, is more focused on how the Dutch police can use social networks to gain insights into protest motivation. Since the study is specifically focused on the police operating in the Netherlands, Twitter data that will be used is limited to Dutch tweets regarding protest activities that take place in the Netherlands.

In addition to the practical contribution, the academic contribution will mainly be focused on the role that social networks can play in predicting protests. Not only will be studied what differences in clusters within this protest-related social network can be identified, but these differences will be evaluated afterwards with the use of another data set. Combining this knowledge will show if this knowledge can be used to predict protest motivation. This will eventually result in a method that can be used to systematically analyze protest-related social networks. Visualizing protest-related social networks and gaining new insights from these visualization to predict new protests differs from research that is done before.

Chapter 2

Related Work

Before selecting the appropriate research method, it is essential to evaluate the methods previously employed in the field. Numerous studies are already conducted in this domain, underscoring the significance of identifying the strengths and limitations associated with the utilized methods. Moreover, it is plausible that certain methodologies are not explored yet within this field, rendering them potentially intriguing for this study to gain novel insights. However, it is crucial to acknowledge that there might be reasons why some methods are not utilized, and it is imperative to ascertain these reasons and determine whether they are applicable to the current study.

2.1 Personality Traits

In order to get more insights in the context of protest motivation, understanding the role of personality traits in protest motivation is essential for gaining a comprehensive understanding of individual engagement in collective action.

2.1.1 Differences in personality traits

Diener and Lucas (2019) define personality traits as: "the reflection of people's characteristics patterns of thoughts, feelings, and behaviors. Personality traits imply consistency and stability - someone who scores high on a specific traits like Extraversion is expected to be sociable in different situations and over time" (p. 278).

To describe someone's personality, the Big Five personality dimensions are most used in the literature. These Big Five personality dimensions are: *Extraversion, Agreeableness, Conscientiousness, Neuroticism, and Openness/Intellect (OCEAN)* (John, Naumann, and Soto, 2008). Each person can be classified in each dimension which means that someone can be low, medium, or high on a specific trait (Diener and Lucas, 2019). There are three criteria that characterize personality traits: (1) consistency, (2) stability, and (3) individual differences. Consistency refers to the trait being consistent in different situations, e.g., at home and at school. Stability means that traits are equal over time (e.g., age 30 and age 40) and individual differences refers to degree of behavior differences between persons (Diener and Lucas, 2019).

To assess personality, the International Personality Item Pool (IPIP) was created by Goldberg et al. (1999). This measurement tool contains 485 5-point Likert scale items in which a distinction is made between the Big Five. Each of the five is broken down into nine facets, which are all assessed by 9-13 items (DeYoung, Quilty, and Peterson, 2007). Although these 485 items are accurate in retrieving someone's personality, the substantial number of items makes it time consuming to use. Therefore, Donnellan et al. (2006) came with the Mini-IPIP scales. A smaller version of the original IPIP. This Mini-IPIP is a form consisting of twenty items in which the factors are equal to the OCEAN dimensions mentioned before (Table 2.1). The tools for assessing someone's personality will be an essential part when this study will be expanded in the way were individuals or group characteristics need to be identified to make estimations regarding protests.

Item	Factor	Text
1	E	Am the life of the party
2	A	Sympathize with others' feelings
3	C	Get chores done right away.
4	N	Have frequent mood swings
5	O	Have a vivid imagination.
6	E	Don't talk a lot. (R)
7	A	Am not interested in other people's problems. (R)
8	C	Often forget to put things back in their proper place. (R)
9	N	Am relaxed most of the time. (R)
10	O	Am not interested in abstract ideas. (R)
11	E	Talk to a lot of different people at parties.
12	A	Feel others' emotions.
13	C	Like order.
14	N	Get upset easily.
15	O	Have difficulty understanding abstract ideas. (R)
16	E	Keep in the background. (R)
17	A	Am not really interested in others. (R)
18	C	Make a mess of things. (R)
19	N	Seldom feel blue. (R)
20	O	Do not have a good imagination. (R)

TABLE 2.1: Mini-IPIP items (adapted from Donnellan et al. (2006))

2.1.2 Exposing personality traits

In section 2.3, the role of personality traits on protest motivation will be explained. Different studies are done with the goal to identify personalities that are more likely to protest and this causes the question to arise whether an online user profile can be linked to a certain personality. Since this study is focusing on the use of Twitter data, the related literature that was focused on also uses data of social media.

As mentioned in section 1.3.1, this study could be used in the future to help estimate the responses of society to news topics. It is therefore important to first assess whether online users accurately reflect the real-world persons. Back et al. (2010) studied the relationship between someone's Facebook profile and their actual personality. This was done by using multiple personality reports which all measured the Big Five personality dimensions described by John, Naumann, and Soto (2008). Results showed that online social networking sites (OSNs) were not used to promote idealized virtual identity and this could be the reason for the popularity of social media according to the researchers.

Now that is clear that an online personality is an accurate reflection of someone's real personality, it is interesting to disclose different methods which retrieve this online personality. Golbeck et al. (2011) for instance studied the possibility of predicting the personality of a Twitter user. The intend was to use this prediction to improve the user experience. The Big Five personality dimensions were all taken into account to get a complete image of a person. To retrieve the personality, past studies were used to retrieve the different elements of the personality dimensions. Selfhout et al. (2010) for example, stated that extraversion, agreeableness, and openness all correlated with friendship selection which is easy to track on social media. Music preferences also relates to personality according to Rentfrow and Gosling

(2003). Furthermore, Jost, West, and Gosling (2009) found that persons that were voted for in elections revealed parts of the personality of the voter. Not only political preferences can be used to identify a personality, Gosling, Sandy, and Potter, 2010 mention that even the preferences for dogs or cats can help identifying someone's personality. To apply this knowledge, Golbeck et al. (2011) made use of multiple user data. Examples are: number of followers, number of followings, density of social network, number of hashtags, number of links, and words per tweet. These data sources will be taken into consideration for this study to identify the users who tweet about protests.

2.2 Herding behavior

In the previous section (2.1) is declared that each personality is different and how these differences can be measured. In addition to individual differences however, research is also done on the behavior of individuals when they are part of a group and the group behavior itself. A group is defined by Forsyth (2018) as: "two or more individuals who are connected by and within social relationships" (p. 4). When social media arose, people were able to literally join new groups with one click. In addition to the possibility to join new groups, social media enabled users to maintain already existing social contacts as well. Although the way of communicating is different when it takes place online, Forsyth (2018) mentions that these online groups have many of the same characteristics and processes compared to offline groups. In particular the group dynamics are similar to those of offline groups; they both develop norms, identify goals, experience conflicts, and have a division of roles within the group. What is especially interesting for this study is the role that groups have on the opinion or actions of the individual within that group. A study by Lazer et al. (2010) showed how political attitudes affected friendship formation and how friendship affected subsequent attitudes. Results showed no considerable evidence that friendship formation was based on political attitudes, however clear evidence showed that the political attitudes shifted in the direction of social ties over time. This process which is defined as "the alignment of thoughts or behaviors of individuals in a group (herd) through local interaction rather than centralized coordination" is called "herding behavior" (Raafat, Chater, and Frith, 2009)(p. 420). This phenomenon causes users within a certain network to shift towards a more equal political attitude in the long term, while at the same time, another network could shift towards an opposite political attitude which may be the basis for polarization.

2.2.1 Exposing herding behavior

Given the presence of herding behavior in social groups, it becomes possible to assess the level of cohesion within a group. As discussed in the previous section (Section 2.2), individuals within a group tend to converge towards a more equal attitude over the long term. This observation leads to the conclusion that groups that have existed for a longer period of time or exhibit an initially equal attitude can be regarded as having strong social bonds.

Shi et al. (2017) studied how retweeting behavior could be predicted within social networks. One of the crucial factors was the social tie strength which significantly impacted the change that someone will retweet a message. The source of information that is retweeted was found to have a trivial impact on the retweeting behavior. This strengthens the hypothesis that polarization can arise on social media because people that are part of a group retweet information of their connections without considering the source of the information.

These findings can be used to expose herding behavior and strong social networks. Twitter users retweeting information of each other are likely to have strong ties, regardless of the information that is retweeted. This is an effective instrument to measure the degree of connection within a group and will therefore be used in this study.

2.3 Contributing factors to protest motivation

For this study, the aim is to identify the factors that affect the motivation to protest. Before studying the role of media and social media, it is necessary to review what kind of studies are already performed within this field. The two previous sections (2.1.1 & 2.2) have described possible factors that could affect the motivation to protest. This section will elaborate on this role in protest motivation.

2.3.1 Identified contributing factors

In section 2.1, it became apparent that most literature use the Big Five personality traits as the main instrument to describe personality. The impact of each of the five traits on protest motivation was examined to understand whether the motivation to protest is depending on personality. Opp and Brandstätter, 2010 found that out of the Big Five personality traits, agreeableness and consciousness have a negative effect on protest participation, while openness was found to have a positive effect on protest participation.

In addition to personality traits, group dynamics should also be considered as the effect that groups have on the individuals within it is described in section 2.2. Group dynamics contribute to forming an opinion and when a group strongly disagrees with the current state of affairs, this could cause protests. This was proven by Klandermans (2002) because evidence was found that group identification fosters protest participation. These were the first signs that identified groups were motivated more to protest compared to individuals and this will be explained further in the following section.

2.3.2 Summarizing framework

Klandermans is investigating protests for decades and states that grievance is the main motivator for protesting since persons who are satisfied will not be motivated to protest. The first studies regarding protest motivation are focused on the cost on benefits of participation. Klandermans (1984) stated that participation was a good manner to change a state of affairs at a low cost. Later, four pathways to protest participation were explored: instrumentality, identity, ideology, and group-based anger (Stekelenburg, Klandermans, and Dijk, 2011). Different studies are done to measure the role of identity at protests. Results showed that the stronger someone's identification with a social category (e.g., elderly, obese, gay people, farmers), the stronger the motivation is to participate in a protest on behalf of that category (Stürmer and Simon, 2004; De Weerd and Klandermans, 1999; Mummendey et al., 1999). These four categories are applied in a model that displays the relation between each other and can be seen in Figure 2.1 (Stekelenburg et al., 2011). This model is the basis for this study because it can be seen as an overview of the current identified factors that affect protest motivation.

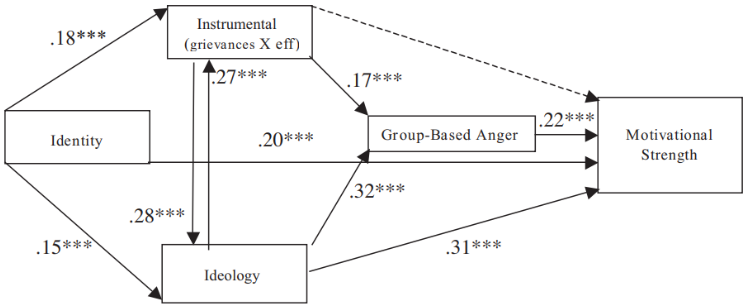

FIGURE 2.1: Structured Equation Modelling: Integrative model (adopted from Stekelenburg, Klandermans, and Dijk (2011))

2.4 Social networks

As mentioned in section 2.3, instrumentality, identity, and ideology are identified as the main motivators for protest participation. In addition to these motivators, Van Stekelenburg and Klandermans (2017) studied the role of the individual within movements. This role of the individual is already studied by Wright (2003). These researchers state that "It is simply obvious that in order to engage in collective action the individual must recognize his or her membership in the relevant collective" (p. 413). The action of recognizing the role of the individual within the broader collective can be seen as a form of depersonalization (Van Stekelenburg and Klandermans, 2017). This process ensures that self-perception and behavior come in line with the contextually relevant in-group prototype and therefore results in the transformation of individuals into group members (Hogg, Terry, and White, 1995).

The identification with a group seems to be one of the most important motivators to participate in a protest and since this study will focus on the role of social media, it is interesting to see what studies are performed to identify groups on these social platforms. However, before a group can be identified, it is necessary to define a group.

The distinction between groups and social networks is described by Wellman, 1997. This researcher state that "just as a local area network is only one kind of a computer network, a group is only one kind of a social network. More precisely, a group is a social network whose ties are tightly bounded within a delimited set and are densely-knit so that almost all network members are directly linked with each other" (p. 1). Therefore, they argue that when people who study online users as groups, they face the problem that one of the characteristics of a group is that they should know the membership and boundaries of the group, which is impossible with people continually joining and leaving (Wellman, 1997).

In addition to these social networks, in which people are linked to each other in the online world, a distinction can also be made between people with certain characteristics. Examples of these characteristics are age, gender, sexual orientation, political preferences, education, jobs and so on. For that reason, a person can belong to multiple groups at the same time, based on the group one wants to identify. Furthermore, there also exist relationships between these groups. For example, persons that are well educated, students, unemployed, or unreligious are more likely to be extreme left (Ooijevaar and Kraaykamp, 2005). To prevent misunderstandings, a distinction between groups is defined by Wellman, 1997, in which people are directly linked to each other, and groups based on person characteristics, will be made in this study. Groups that consist of persons that are directly linked to each within social networks are referred to as **groups**, and groups that consist of people with the same characteristics will be indicated by the specific characteristic in common (e.g., **political groups, age groups, education groups**).

For the aim of this study, these different group characteristics are important as protests are often a result of an incident that affects such a particular group. The identification of these groups can be done by making use of different elements that can be retrieved from online messages such as tweets. Nguyen et al. (2013) studied the use of language in order to make statements about the age and gender of tweeters. This study used six style variables that users are aware of and explicitly choose to use: *capitalized words, alphabetical lengthening, intensifiers, LIWC-prepositions, word length, and tweet length*. Additionally, the used words were analyzed per age group which resulted in the top features for younger and older people. The word daughter for example was used often by elder people, while school was used more by younger people (Nguyen et al., 2013).

It is therefore clear that user characteristics can be retrieved from individual tweets and this can be used to identify relations between certain groups. On the other hand, tweeters can also be used to identify relationships between topics. Es, Geenen, and Boeschoten (2015) studied how viewing patterns on television could be traced down with the use of Twitter data. The researchers made use of official and popular hashtags that were used in tweets that referred to television programs. The collection contained 135.882 tweets that were posted by 39.792 unique tweeters. The researchers made use of a method suggested in a book by Rogers (2013) in which new methods with the use of social networking sites to study social networks were explained. This method allowed the researchers to identify the relationships between programs by making use of the users who tweeted about multiple programs. The result by using the force-directed layout algorithm "ForceAtlas2" is shown in Figure 2.2.

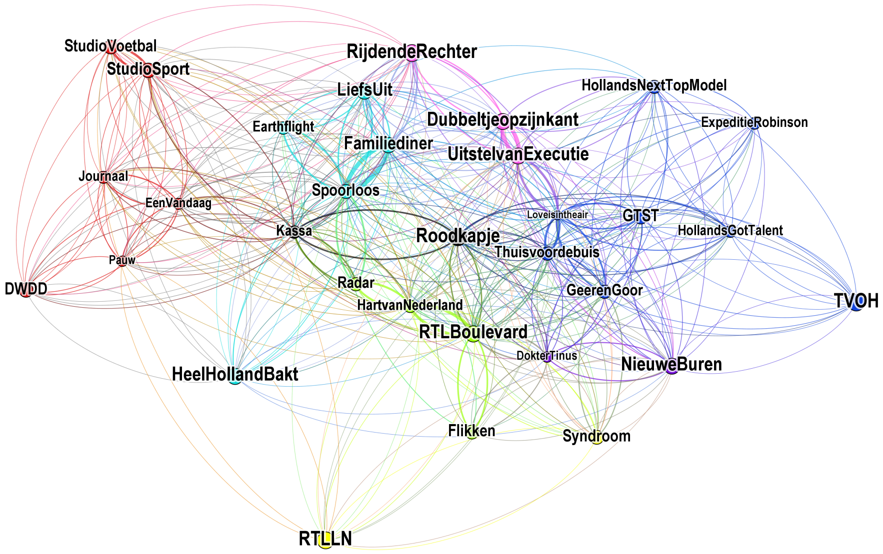

FIGURE 2.2: Overlap in tweeters among thirty-two television programs in the Netherlands (adopted from Es, Geenen, and Boeschoten (2015))

Another research, done by Veerbeek et al. (2022) on antisemitism was done with the use of automated classification models. These machine learning tools helped the researchers to analyze their data more efficient. The model was trained with 9,000 messages and should assess three 3 aspects; (1) does the message include antisemitism and if so, if it is explicit or implicit antisemitism, (2) if the antisemitism is expressed in the text, the figure, or in both, and (3) if the message contains antisemitism, what form of antisemitism is used. The best scoring model for this study was the Bidirectional Gated Recurrent unit (Bi-GRU), which is neural network with an internal long-term memory which can process series of inputs. Having the messages labeled, the researchers were interested in the persons behind them. Since Twitter users are not automatically classified into groups, they made use of a network analysis, in which interactions between users are used to identify social networks. Different methods exist to identify such a network and possible elements can be used for this. Examples to link users to each other are mentions, followers, or the user behind a tweet that is like by a specific person (Veerbeek et al., 2022). In this case, the researchers used retweets, since this can be seen as a strong indication of agreement (Veerbeek et al., 2022). As shown in Figure 2.3, the results showed interesting outcomes which tell something about the relationships between users who use antisemitism in their tweets. In this way conclusions can be drawn that for instance conservative right-conspiracy (green) is not connected to Pro-Palestine (blue) but is connected to Pro-Israel (red). The research method used by Veerbeek et al. (2022), will be the basis for one part of the study we will conduct. In this way a social network of the users tweeting about different protest in the Netherlands will be composed.

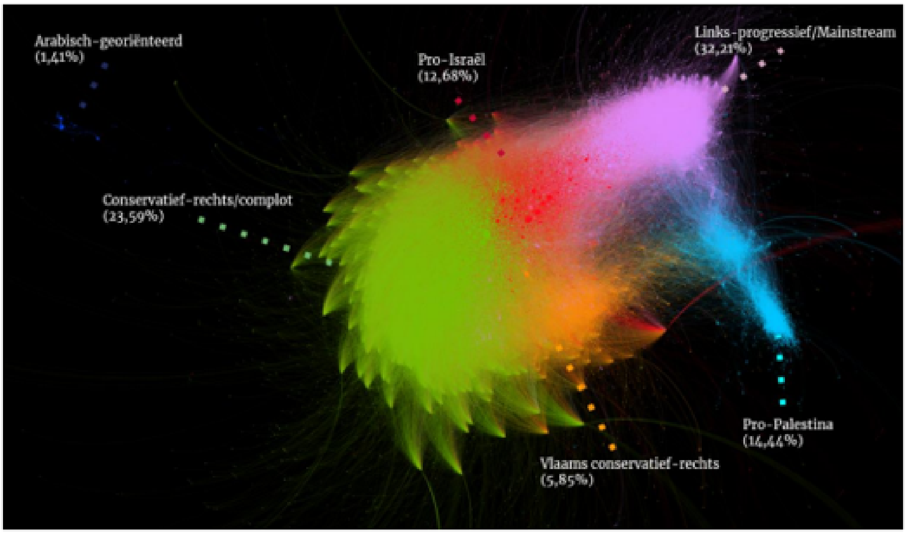

FIGURE 2.3: Social network according to retweets of Twitter users. Nodes are individual users, lines represent retweets. (Adopted from Veerbeek et al. (2022))

In Figure 2.3, the social network of users tweeting about antisemitism is visual. As already mentioned in section 2.2, within (online) groups, roles can be identified. A study of Welser et al. (2007) extensively explains the different roles within online discussion groups and as Twitter is used broadly for discussing, this is interesting to take into account. These researchers mention which roles can be identified in within online groups. Examples are local experts, answer people, conversationalists, and fans. The study elaborated on the differences in local networks for these roles and Figure 2.4 visualizes the different roles interact differently. These visuals can be used in this study as a reference when determining the roles of individuals tweeting about a protest.

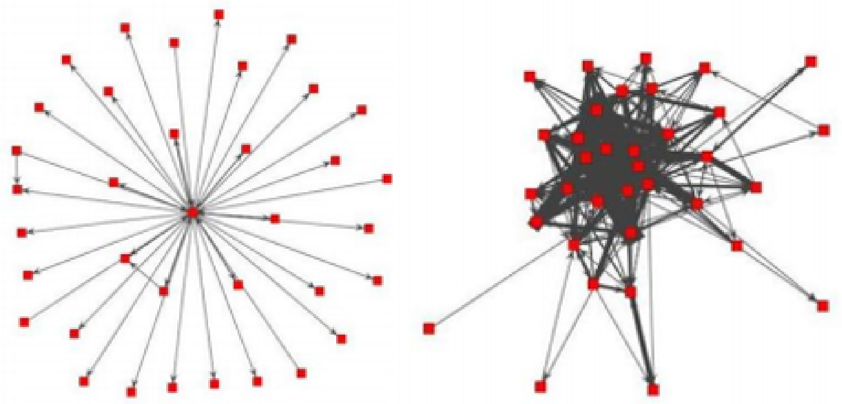

FIGURE 2.4: Examples of local networks for an Answer Person (left) and a Discussion Person (right). Directed tie = replied to & ; line thickness proportional to # of ties in the bi-directed relation. (Adopted from Welser et al. (2007))

2.5 Media and social media

As mentioned previously, the number of protests is rising each year (1) and additionally, during the previous subsections it became clear that one of the largest protest motivators is group-based anger. Especially when one can identify itself clearly with a certain affected group, the motivation to protest rises. But what role do media and social media have in the organization of and motivation to protest?

2.5.1 Media

Within this study, the media is referred to as both online media (e.g., online news sites) and offline media (e.g., newspapers). More specific, the focus is on media that has as their goal to spread news. In the next section (2.5.2), the effect of personalization on platforms is discussed. For media this is different since the news is simply posted and not different for each individual reader. This could reduce the chance that media polarizes society. However, it should be taken into consideration that the number of different media channels is enormous and that everyone will therefore choose the news channels that best suits their preferences. Each news channel tries to distinguish itself from other news channels by highlighting articles that fits the wishes of the readers. A study performed by Mitchell et al. (2014), found that liberals used different news sources compared to conservatives as they both choose sources that confirm their opinions. Stroud (2008) mentions that this lacking base of shared information between political groups lead to a polarized attitudes when talking about political matters. The way that news consumers choose the media sources they prefer could therefore be seen as the creation of their own echo chambers which contribute to the polarization between groups with different political interests.

2.5.2 Social media

This arising polarization when groups consume news from different sources is especially interesting to study within the field of social media. This because social media nowadays is focused on personalization to serve each individual user exactly what they like.

When social media started in the zeros, it was seen as a positive innovation because it improved the accessibility to the news, and made users able to produce media themselves (Jenkins, 2006; Shirky, 2009). Nowadays however, the goal to personalize a platform for each individual user poses a danger that users end up in an echo chamber, which results in users only seeing the news items that interest them and that confirm their opinions (Baumann et al., 2020). When one's beliefs are constantly confirmed, a normal discussion between persons with different perspective becomes harder and may lead to polarization. Sunstein (1999) state that group polarization helps to explain extremism and the behavior of religious organizations. This may also be the case for protest activities. A recent study, conducted in the context of the Utrecht Data School, studied the role of social media conversations on protests (Bakker et al., 2021). Results showed that different groups can be identified within tweets about a protest. In this study, we will elaborate on this by making the comparison between different protest instead of one protest.

not only does social media contribute to the polarization of society, it also is an effective tool to organize protest which lowers the barrier to do so. In the Netherlands, the first large example of mobilization via online platforms occurred in 2012, known as Project X Haren. Although it was not intended to mobilize people, the invitation for the sixteenth birthday of a girl living in Haren was accidentally set to be publicly visible, enabling all one billion Facebook users at that time to see and share the invitation. Although the police were notified beforehand, and youth were asked not to come, thousands of people came to party resulting in chaos, shop looting and overall destruction of the small town (Van Eck, 2017). This incident in Haren illustrates the power of social media to mobilize groups. In this case however, the organizer had no intention to mobilize thousands of people. This is obviously different for protest organization, where the goal is to convince people to join a protest.

A study by González-Bailón et al., 2011 focused on the effect of social media on protest organization. The researchers found patterns of recruitment and information diffusion in a movement that took place in Spain. It was concluded that exposure to repeated messages about the protest seemed to "activate" users. However, this was only the case when these messages were sent by multiple sources within the personal network. This study also found that users who were the most effective at the distribution of protest information were not necessarily the ones who had the most followers. More decisive was the position within the social network because the users with the most central positions and therefore the ones who were closely connected to other well-connected users spread the most information. The importance of the position of a person within a social network should be kept in mind as it was already described in section 2.2.1 that herding behavior had a positive correlation with the social tie strength between users. It therefore seems that people are more inclined to be motivated by people with whom they are strongly connected and these connections will be further investigated in this study when a Social Network Analysis (SNA) will be performed.

The aforementioned role of social media was mainly the direct effect it has on organizing a protest. Social media however, also has an indirect role in protest motivation. Mundt, Ross, and Burnett (2018) studied the role of social media on scaling of social movements in the context of Black Lives Matter (BLM) protests. Results showed that BLM groups perceive themselves to be part of a larger BLM movement. This was mainly due to relationships formed through social media platforms. This role of social media contributing to the feeling of belonging to a group can be seen as an indirect effect on protest motivation as sections 2.3.1 & 2.3.2 already described that the feeling of group identity increases protest motivation.

2.5.3 Protest-related social network analysis method

The research aim of this study is to develop a method than can be used to analyze protests. Related studies are already performed to provide guidance in the creation of such a method. In a study by Müter, Loerakker, and Veltkamp (2023), the possibility of information extraction of protest-related tweets was explored. Within this study, the method to collect tweets is emphasized. To retrieve tweets relevant to the protest under study, the query was defined based on categories. The defined categories, along with the keywords used in this paper, are visible in Table 2.2.

Category	**Terms**
Generic	'demonstration', 'protest', 'riots', 'demonstrators', 'unrest', 'social unrest', 'danger', 'occupation', 'on the streets', 'dissatisfaction', 'uprising', 'police', 'outbreak', 'hooligans', 'rioters', 'activists', 'arrested activists', 'demonstrating', 'taking action', 'rioting'
Specific	'March 11th', 'the biggest demo ever', 'Buying out farmers', 'nitrogen mediator', 'Torch', 'Tractor', 'Flag'
Related Groups	'FDF', 'Farmers Defence Force', 'farmer positions', 'farmer action groups', 'Mark van den Oever', 'Farmer activists'
Counter Groups	'Extinction Rebellion', 'climate activists', 'vulnerable natural areas', 'Hannah Prins', 'fossil subsidies', 'Environmental activists'
Related Topics	'childcare allowance affair', 'gas crisis', 'subsided houses', 'water damage', 'Nitrogen', 'Childcare allowance affair', 'Groningen gas', 'Immigration', 'Refugees', 'forced buyout', 'nitrogen'
Supporting Parties	'PVV', 'Geert Wilders', 'Groep Van Haga', 'Wybren van Haga', 'March 11th demo', 'Political parties', 'BBB', 'Farmer Citizen Movement', 'rural residents', 'Caroline van der Plas', 'FvD', 'Thierry Baudet', 'Gideon van Meijeren', 'Pepijn van Houwelingen'
Opposed Parties	'Tjeerd de Groot', 'Jan Paternotte', 'Kaag', 'Wierd Duk', 'PAS reporters', 'big agro', 'Johan Remkes'

TABLE 2.2: Keyword categories translated from Dutch as defined by Müter, Loerakker, and Veltkamp, 2023.

Chapter 3

Research Method

3.1 Research Method

The objective of this study is to develop a method that can guide future protest-related studies in doing analyses. In order to find the appropriate research method to fulfil this objective, the literature is consulted.

According to Hevner et al. (2004), much of the research within the field of Information Systems can be classified into two areas: behavioral science and design science. Where design science focuses on the creation of new artifacts to extend the capabilities of human and organizations, behavioral science seeks to verify existing- and develop new theories that explain or predict human or organizational behavior (Hevner et al., 2004).

Since this study focuses on both identifying the rationale behind protesting, and the development of a theory that explains what factors contribute to the motivation for humans to protest (artifact), it can be classified in the behavioral science area as well as in the design science. Therefore, both parts are distinguished in the following sections (3.1 & 3.1.1).

3.1.1 Behavioral Science

As research is conducted on the rationale behind humans protesting within this study. The study can be partly classified within the behavioral science. This form of science consists of multiple research designs that can be applied. Stangor (2014) explains three main approaches in his book: Descriptive-, Correlational-, and Experimental Research. Since the different factors and the relationships between them is investigated in this study, it can be identified as a Correlational Research. Stangor (2014) defines the goal of the correlational research design as: "To assess the relationship between and among two or more variables". If correlations between certain variables exist, predictions can be made when values of some variables are known by making use of Pearson product-moments correlation coefficient (r) (Stangor, 2014). The Correlation Research Method is executed by making use of the mixed methods approach. The mixed method approach is defined as a method in which researchers use methods of collecting or analyzing data from both the quantitative and qualitative research approaches in a single research (Creswell, 2003).

3.1.2 Design Science

As mentioned before, Hevner et al. (2004) defined design science as a form of science which creates and evaluates IT artifacts intended to solve identified organizational problems. In contribution to that, Hevner et al. (2004) also created seven guidelines to assist researchers to understand the necessary characteristics that an effective design-science research must meet. These guidelines show that the relevance of the artifact is particularly important since it must serve a certain identified goal which contributes to solving a problem (Hevner et al., 2004). For this study, the eventual problem that is trying to be solved can be defined as: "the Dutch police not being able to estimate the impact of social media and social networks on protest motivation, and eventually on protest size". Although this study will only partly contribute to solving this problem, it is still from utmost importance that this goal is kept in mind.

In addition to guidelines being composed by Hevner et al. (2004), there are also different process models created to use within the field of design science. One of these process models was created by Peffers et al. (2007) which consists of six activities. These activities are:

1. Problem identification & Motivation
2. Objective of a solution
3. Design & Development
4. Demonstration
5. Evaluation
6. Communication

In order to visualize the implementation of this method in this study, Figure 3.1 was created.

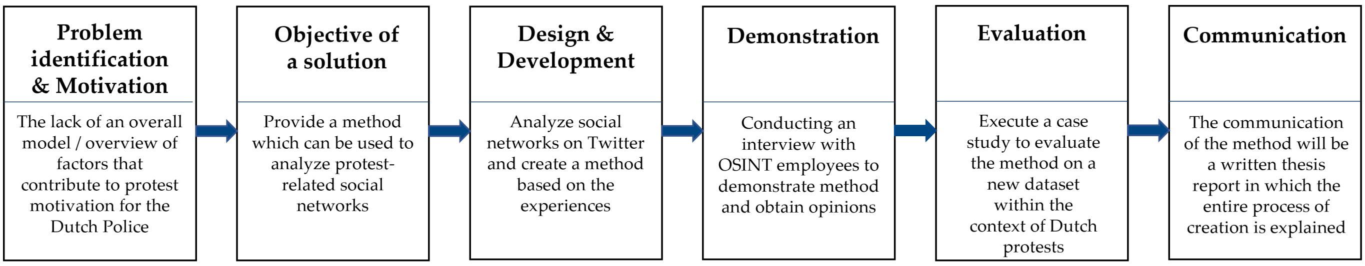

FIGURE 3.1: Design Science Process Model. Adopted from Peffers et al. (2007).

Problem identification & Motivation
During the first phase of this study, named the problem identification & motivation stage, the main problem is identified and the motivation and relevance of the research are articulated (Peffers et al., 2007). To identify the problem, the literature is studied to investigate whether a research gap can be found. When a gap is found, the motivation for the relevance of doing research within this area can be explained. Additionally, the importance of a new artifact within this field is elaborated on, as this will increase the chance of implementation the working field. The main problem and relevance are already stated in chapter 1 and will therefore not be further explained.

Objective of a solution
Once the problem is identified, a possible solution is outlined. With the knowledge obtained in the first phase of the study, it is known what is already studied and what artifact will contribute to the research. The objective for the desired solution is qualitative, which means that it will describe how the artifact is expected to support solutions to problems not addressed before (Peffers et al., 2007).

Design & Development
According to Peffers et al. (2007), this phase involves determining the functionalities and architecture of the artifacts. Therefore, it is crucial to have a comprehensive understanding of all the motivators identified in previous research and to explore the possibilities of social network analysis (SNA), as this forms the foundation of the artifact. In this study, the primary focus is on social networks as a key motivator, necessitating a detailed examination of this aspect. Once all the influencing factors have been explored, the next step involves investigating the relationships between these factors using exploratory social network analysis. The insights gained from this phase are then used to develop a systematic method for protest-related social network analysis.

Demonstration
This phase of the study aims to demonstrate the utility of the proposed solution, as emphasized by Peffers et al. (2007). The method is showcased to other researchers within the same working field to elicit critical feedback and insights. By demonstrating the method, it becomes an opportunity to not only showcase its effectiveness but also to engage in discussions about potential future work and the broader possibilities that this method presents.

Evaluation
The evaluation phase, as proposed by Peffers et al. (2007) in the design science methodology, is a crucial step in determining the effectiveness, usefulness and usability of the design solution in relation to the intended problem. This phase ensures that the solution is not only effective and efficient but also satisfactory for its users and stakeholders. To achieve this, it is necessary to subject the created method to a case study in order to evaluate whether the performance is acceptable. This iterative process allows for identifying any weaknesses in the artifact and making necessary improvements to enhance its effectiveness in solving the problem at hand (Peffers et al., 2007).

In order to fulfill these requirements, the results obtained from the original data set are subjected to scrutiny using a new data set. This verification process not only validates the results but also allows for generalizability within the context of protests to be substantiated. By conducting this evaluation, the method can be refined and its applicability can be enhanced, thereby improving its effectiveness in addressing the research problem.

Communication
In order to communicate the existence of the new developed model, this study is shared with the appropriate persons within the Dutch police organization. Additionally, the work will verbally be presented to the OSINT team. This ensures that the performed work is also recognized by the organization. This is even more important knowing that this work can be used for further research within or outside the organization.

3.2 Overview of used research method per research question

Research question	Literature	Data analysis
1. Which motivating factors to protest can be identified from existing literature?	X	
2. What differences in hashtag usage can be used to analyze protest-related social networks?	X	X
3. What differences in media consumption can be identified when comparing clusters of a protest-related social network?		X
4. What differences in group characteristics can be identified when comparing clusters of a protest-related social network?		X

TABLE 3.1: Research method used for each research question.

3.3 Social Network Analysis

For this study, visualizations are used as a tool to address the research questions. To identify social networks, the force-directed layout algorithm "ForceAtlas2", as used in prior studies, is utilized (Es, Geenen, and Boeschoten, 2015). This algorithm is used within the open source visualization software Gephi 0.10.1.

3.3.1 Data

The data utilized in creating the visualizations is obtained from Twitter. This data is publicly accessible and is automatically retrieved via the Twitter API. This API had to be triggered manually and was only able to retrieve the tweets up to two weeks in the past. This forced the researchers two start the algorithm each two weeks, but when no high Twitter activity was expected, this was postponed sometimes, resulting in gaps within the data set.

To ensure valuable and non-biased tweets for this study, the only key words that is used in the tweet retrieval query, is the Dutch word "demonstratie" (protest). Since this study does not focus on a specific type of protest, the less filters are used, the more non-biased findings can be retrieved. However, for the evaluation data set, more keywords are used to identify possible differences in outcomes. The list of keywords of the evaluation data set can be found in Appendix D.

The retrieved tweets contain all sorts of tweets (e.g., opinions of protests, protest announcements, possible reasons for protests, etc.). Using this single keyword for the original data set, it may be possible to identify topics that are frequently associated with protests. Various attributes that could be utilized for visualizations are extracted from these tweets. For each tweet, the following data is available:

- Created at
- Fiends count
- Followers count
- Full text
- Hashtags used in tweet
- User favorites count
- User listed count
- Hashed user name
- User statuses count
- Users mentioned

To effectively utilize the data, a distinction is made between tweets containing hashtags and tweets containing mentions. According to Bruns and Burgess (2011), a twitter hashtag can be defined as "a short keyword, prefixed with the hash symbol ‘#’ – as a means of coordinating a distributed discussion between more or less large groups of users, who do not need to be connected through existing ‘follower’ networks" (p.1). These hashtags can be used to identify topics that were discussed on Twitter. Additionally, mentions can provide insights into the social networks of users. Studying mentions is interesting as research has shown that social networks can evolve over time.

In this study, both hashtags and mentions are analyzed separately and in combination to identify patterns related to groups, topics, and events. The data set is divided into two subsets: one containing hashtags and the other containing mentions. To facilitate data transformation, Python is used to provide the combinations in statistical sense between multiple hashtags or mentions within a single tweet in an undirected manner. This linking process results in a data set where hashtags appearing in a tweet, such as A, B, and C, are represented with their relations, resulting in the connections A-B, A-C, B-C. The same process is applied to mentions. The resulting data set for the hashtags contains 352.397 relationships between hashtags, and 550.380 relationships between mentions. Since most retrieved tweets in the original data were tweeted in 2022, most relationships are also coming from that year. A further analysis of the distribution of retrieved data is visualized in the findings section (4).

Filtering

In order to get useful insights in the data, filtering is necessary after the ForceAtas 2 algorithm is ran because of the large data sets for visualization.

For the hashtags, the Weighted Degree (WD) distribution is computed and analyzed to compute an appropriate filter. The weighted degree is a component that is calculated based on sum of the weights of all edges connected to that node. The equation of the calculation can be seen below. W(v) represents the weighted degree of a node v, N(v) represents the set of nodes that are connected to v, and w(u, v) represents the weight of the edge between nodes u and v

$$w(v) = \sum_{u \in N(v)} w(u, v)$$

If node A for instance has two edges, one edge to node B (A-B, weight 3) and one edge to node C (A-C, weight 2), the weighted degree for node A is five. When looking at the weighted degree of nodes in the data set, it is a good representation of the usage of the hashtags. The higher the weighted degree of a hashtag, the more it has been connected to other hashtags. The maximum number of nodes that is used to keep topic analysis possible is set on three hundred. In order to reach this limit, the weighted degree filter is increased with steps of one hundred until less than three hundred nodes are remaining. The results of this analysis are shown in the findings section 4.

For the mentions, the strength between two persons is more important in order to classify the persons into social networks. Therefore, the Edge Weight (EW) minimum is set on 2, meaning that two persons should be mentioned in a single tweet at least twice.

Categorization

Since coloring helps getting insights in data quicker, categorization is applied to the data that remains after the filtering is done. The maximum number of colors that humans can distinguish easily is twelve (RockContent, 2022). RockContent (2022) defines different palettes for various types of visualizations. Three types of color schemes can be identified; sequential-, diverging-, qualitative schemes (ColorBrewer, 2021). According to ColorBrewer (2021), qualitative schemes are best suited for categorical data. For the qualitative schemes, the website defines a color palette of 12 colors that you can use to categorize your data. To be conform with this palette, a maximum of 12 categories is defined with the goal to classify the least number of nodes to the category "Other".

For the mentions data set, the 10 largest social networks are used to analyze. Gephi is utilized to classify the remaining users after filtering into modularity classes. The result is a colorful network where the classes can be distinguished by the color usage.

3.3.2 Visualizations

Visualization is a useful manner to get insights into data. With the chosen data sets, different ways of visualizing the data must be chosen to get the most useful information. In the section below, each visualization is substantiated. For all visualizations, the same settings for ForceAtlas 2 were used, meaning that each visualization was scaled with 2.0, and overlap of nodes is prevented.

Hashtags

- **Global overview of data**
  Before applying filtering, it is useful to get a global overview of the data. This may give interesting insights that can only be retrieved without filtering. **The main goal** of this visualization is to analyze the positions of all nodes and additionally, to see whether a certain pattern in position can be identified. This visualization is made after categorization, meaning that all the nodes with a weighted degree > 300, are categorized and filled with a color. This allows analyzing the position of these topics within the larger network in specific. As the focus is on the positions, rather than the relations between topics, the edges are not shown in the figure.

- **Global overview of data with outliers filtered out**
  As elaborated in section 3.3.1, and visible in Figure 4.3, applying a relatively small filter on the weighted degree of nodes already excludes a large number of data. In this visualization a minimum weighted degree filter of 20 will therefore be applied. **The main goal** of this visualization is to identify what the difference with the previous visualization is. Questions that can be answered by applying this filter, are for instance, where the infrequently used hashtags were located and what are the positions of the remaining nodes. To get more insights in the reasons why some nodes pass this filter but are still located further from the center, the edges are visible in this visualization.

- **Commonly used topics; position overview**
  After the global overviews are made, the main topics are analyzed. As argued in section 3.3.1, to identify the main used topics, a filter of a minimum weighted degree of three hundred is used. Since the remaining nodes are all categorized, this will give the first proper overview of the topics that were used within the time frame of the data set. Additionally, the size of the nodes is based on the weighted degree, which allows for more insights in the frequency of use within the remaining data set. **The main goal** of this visualization is to identify the topics that were used the most, and additionally analyze the positions of the topics.

- **Mainly used topics; relationships overview**
  After the position of the mainly used topics have been identified, the relationships between the topics can be studied. Therefore, the edges are added to the previous visualization. Since the edge weights differ widely, these weights are rescaled with a minimum of 1 and a maximum of 50. **The main goal** of this visualization is to identify the global relationships between the topics. Force Atlas 2 uses a mix of the colors of the two nodes as the color of the edge in between them, which ensures that the strong relationships are clearly visible.

- **Tight relationships between topics overview**
  While the previous visualization may give global insights in the relationships between the main topics, it also includes relationships that have been used once. For this visualization an edge filter will therefore be applied to outline relations between topics that have been used more than or equal to ten times. **The main goal** of this visualization is therefore to identify the strong relationships and get an answer to the question which hashtags are used often together.

- **Positions of hashtags referring to media**
  One of the goals of this study is to identify the role of media within protests and protests motivation. Since all tweets in the data set are retrieved with the filtering word "demonstratie" (protest), it is interesting to see which hashtags referring to media are used within this social network. **The main goal** of this visualization is therefore to identify the position of the hashtags referring to media.

- **Relationships of hashtags referring to media**
  In addition to the positions of hashtags referring to media, the relationships with topics of other categories is relevant as well. By outlining the relationships between the media and other topics, differences may come to light. **The main goal** of this visualization is therefore to identify the relationships of the hashtags referring to media.

- **Commonly used topics overview between February 2021 and February 2022**
  Since Twitter conversations are continuous events, adding a time element can be very useful. Protests are planned at specific times and places, and by adding a time frame the Twitter activity can be analyzed. Before specific target dates can be identified, the global activity during 2021 is visualized. **The main goal** of this visualization is to identify the topics that caused notable Twitter activity within this time span.

- **Commonly used topics overview between February 2022 and February 2023**
  As the previous visualization represents the Twitter activity in 2021, this visualization will represent the activity the year after. **The main goal** of this visualization is therefore the same as the previous one; identifying topics that caused notable Twitter activity within this time span. In addition to identifying the activity, this second visualization offers the opportunity to compare both years.

Mentions

For all visualizations of the mentions, the same settings for ForceAtlas 2 were used, meaning that each visualization was scaled with 50.0, and overlap of nodes is prevented.

- **Global overview of social networks** When the weakest relationships between users have been filtered out, a first visualization can be realized in which the largest social networks are visual. By using a unique color for each modularity class, insights can be retrieved regarding the structure of the network. **The main goal** of this visualization is therefore to analyze the structure of the network and compare the different social networks that it consists of.

Hashtag usage per modularity class

- **Most used hashtags per modularity class between Feb 21 and Feb 22**
  When the largest modularity classes have been identified, it is useful to analyze the hashtags that are used among these classes. This can provide insights in the topics that are discussed within or between certain networks. In order to get useful visualization, the data set is split into two years, where the split is set on February 2022. **The main goal** of this visualization is therefore to identify topic differences or similarities between the identified modularity classes within the time span of Feb 21 and Feb22 .

- **Most used hashtags per modularity class between Feb 22 and Feb 23**
  After the visualization for 2021 is made, the same is done for the year after. This will result in a similar structured network, where differences can be explored. **The main goal** of this visualization is therefore to identify topic differences or similarities between the identified modularity classes within the time span of Feb 22 and Feb23, and to provide reference material for comparison with the year before.

- **Social network interactions before an event based on hashtags.**
  The topics that are discussed in social networks can vary, for this study however, we are interested in topics related to protests. Therefore an analysis is done on the use of hashtags related to a specific protest among social networks. We will use a dynamic visualization to indicate dynamic differences over time. **The main goal** of this visualization is therefore to identify differences in hashtag use related to a specific protest.

3.4 Validity threats

This study seeks to gather evidence that social networks can aid in understanding and predicting protests. In order to achieve results that contribute to this objective, it is crucial to consider the validity threats. By identifying these threats prior to conducting the study, weaknesses in the research design and methodology can be pinpointed, thereby safeguarding against erroneous conclusions drawn from inadequately employed approaches.

3.4.1 Construct validity

Within this study, a distinction is made between three aspects that van be analyzed to gain insights into protest motivation. The first aspect considered is the use of hashtags. Relationships are established between hashtags that are frequently used together within the context of protests. Since hashtags used within the same tweets often share a thematic connection, this approach proves suitable for measuring and analyzing protest dynamics. As a result, it does not pose a threat to construct validity.

The second factor is group dynamics, which is measured by analyzing a social network that is made based on mentions on Twitter. The concepts that are intended to be measure are the centrality, tightness, and hashtag usage of clusters. A Social Network Analysis is a suitable method to do so, as it does not only visually illustrates the differences, but additionally provides quantitative data that can be used for statistical tests. This method is therefore identified as not posing construct validity threats.

The third factor that is analyzed, is the role of the media. This is done by identifying hashtags that refer to media, after which their position based on the relationships can be analyzed. As the focus for this study is on the role of the media in protest motivation within protest-related social networks. This is considered as the appropriate manner to do so. All data about the media is coming from the users within the network itself, allowing to analyze what sources are mostly referred to, and to which topics they are connected.

3.4.2 Content validity

Three possible threats to content validity can be identified when analyzing the used data. The first threat in content validity is about the data used to analyze group dynamics. This is measured by analyzing a social network that is made based on mentions on Twitter. The mention data used for this purpose, however, is coming from multiple users that are mentioned in a single tweet, which is different compared to a single user mentioning another user in a tweet. Relationships that arise, are therefore based on a user that mentions two or more users in a tweet, which does not necessarily mean that the mentioned users have an interaction with each other. This may result in the threat that relations are created between users how do not directly interact with each other.

The second threat in content validity regards the role of the media. The previous section describes how the used method supports the aim that of this study. However, it is noteworthy that caution should be used in drawing conclusions, as the used data is coming from Twitter users and no direct impact of the media is measured.

The third threat is about the modularity classes that are identified with the program Gephi 0.10.1. Although this is an efficient manner to determine classes, in the method used for this study, it is purely based on mentions within the data set. While the assumption is made that each user can be categorized in exactly one class, this may in reality not always be the case.

3.4.3 Internal validity

The internal validity of a study requires a reliable research method to mitigate potential factors that could unintentionally influence the outcomes. In this study, two large data sets are utilized, retrieved using the same method. It is important to note that there is a difference in queries between the two sets. The original data set consists of a single query word, while the evaluation data set is obtained using over seventy query words. However, this difference can be considered minimal since both sets aimed to obtain protest-related data, which aligns with the study's objectives.

To address potential threats of external influencing factors in the statistical analyses, precautions are taken by using portions of the total instead of absolute values. This approach helps mitigate the impact of factors such as category or class size, thereby preserving the internal validity of the study.

3.4.4 External validity

The evaluation phase of this study contributes to the external validity of the research, which pertains to the extent to which the results can be generalize. This study has a well-defined scope, focusing on Dutch protests and using Twitter as the sole data source. However, this can also pose a threat as one data source may not be enough for generalizable results. However, this approach also presents a potential threat, as relying on a single data source may limit the generalizability of the findings. Nevertheless, previous research has demonstrated that Twitter data provides a meaningful reflection of society, thereby supporting its suitability as a data source for this study.

Furthermore, it is important to note that the data used in this study spans a two-year period. This extended time frame enhances the generalizability of the findings, as it is unlikely that different time periods would yield significantly different outcomes.

Chapter 4

Results

The results section is organized into three main phases. First, the findings of the exploratory phase are presented, providing initial insights into the data. Next, the developed method based on these findings is described in detail. Lastly, the method is evaluated using a new data set.

4.1 Distribution of data

Figure 4.1 displays the distribution of all tweets gathered for this study using the filter word "demonstratie". During the period of January 2021 and February 2023, a total of 685,521 tweets were retrieved. The graph reveals that there are four months in which no data was retrieved (July 2021, September 2022, October 2022). Additionally, it is worth noting that the y-axis has a logarithmic scale, which is chosen because of the large difference in tweet retrieval in March 2021. During this month, a total of 92,624 tweets was obtained, while during the months February 2021, April 2021, June, 2021, and December 2022, less than a thousand tweets were retrieved.

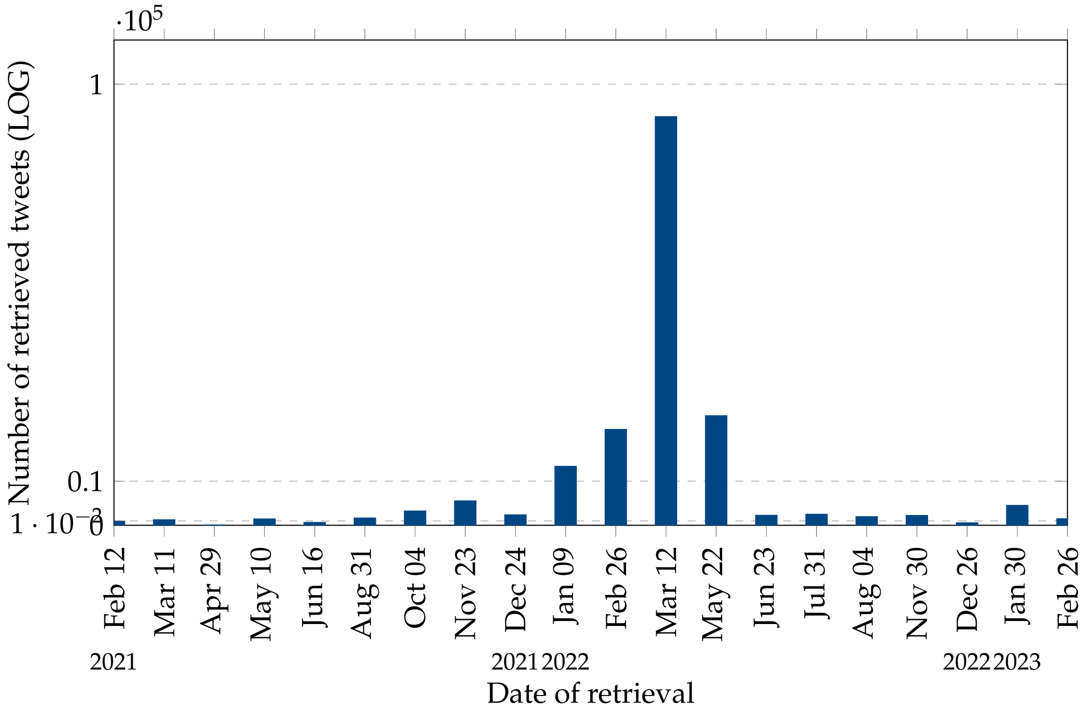

FIGURE 4.1: Distribution of retrieved tweets per month with filter word "demonstratie"

For this study, the objective is to examine the relationships between topics and users. To achieve this, hashtags and mentions used in the tweets are collected to construct social networks. As explained in section 3.3.1, these relationships are established when hashtags have been used in the same tweet. The data set of the hashtags consists of 41,965 tweets, which together consists of 24,319 different hashtags with 157,492 unique edges (relationships) in between them. Taking into account instances where relationships are formed multiple times, such as when the same hashtags are used together more than once, a total of 352,397 relationships are identified. The distribution of this number is depicted in Figure 4.2, which demonstrates the distribution over time and highlights the peak in March 2021. As mentioned in the previous section, certain months were excluded due to the absence of significant protest activities. The figure reveals a decline in relationships during the summer months, suggesting a potential decrease in protest-related activities during that period.

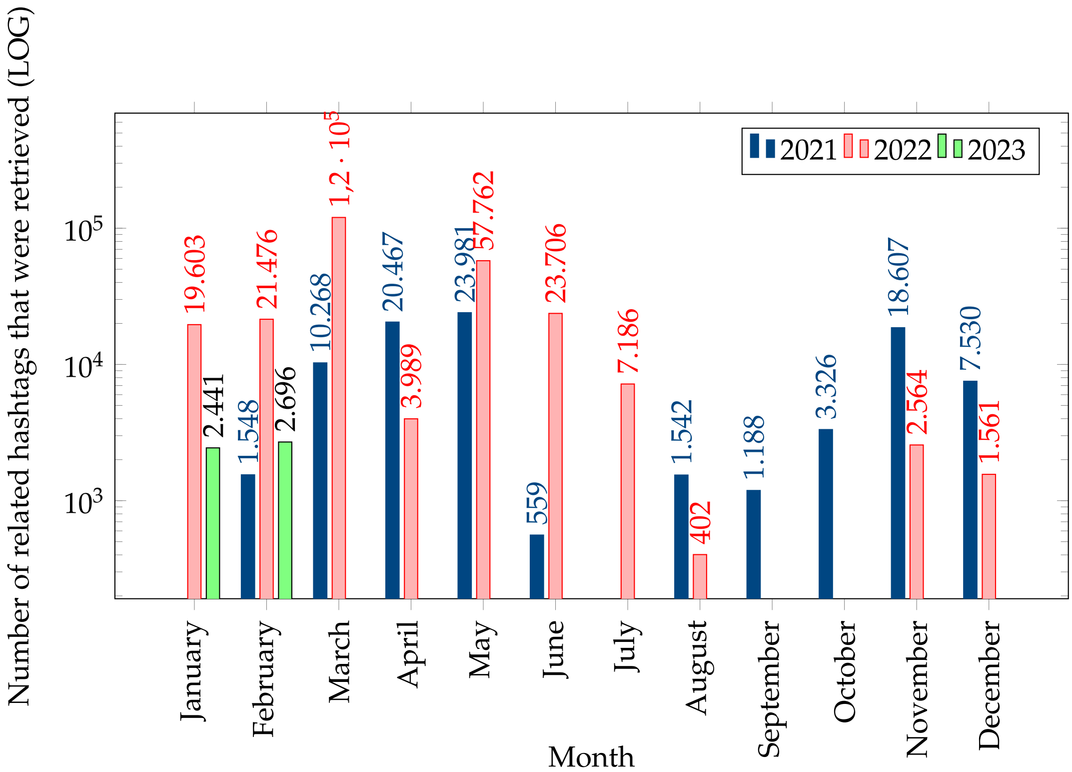

FIGURE 4.2: Distribution of gathered relationships between hashtags between February 2021 and February 2023. Different colors represent the different years.

4.2 Filtering and categorization

4.2.1 Filtering

In Figure 4.3, an overview of the remaining nodes is presented, considering different weighted degree filters with a maximum of one hundred. Without applying any filters, the total number of nodes (24,319) is displayed. However, when a filter is applied to remove nodes with a weighted degree less than 20 ($WD \geq 20$), the number of remaining nodes is only 3,171. This represents a substantial decrease of 21,148 nodes (87%), each corresponding to a unique hashtag. The significant reduction in nodes suggests that a large number of hashtags have minimal connections with other topics, being mentioned less than 20 times.

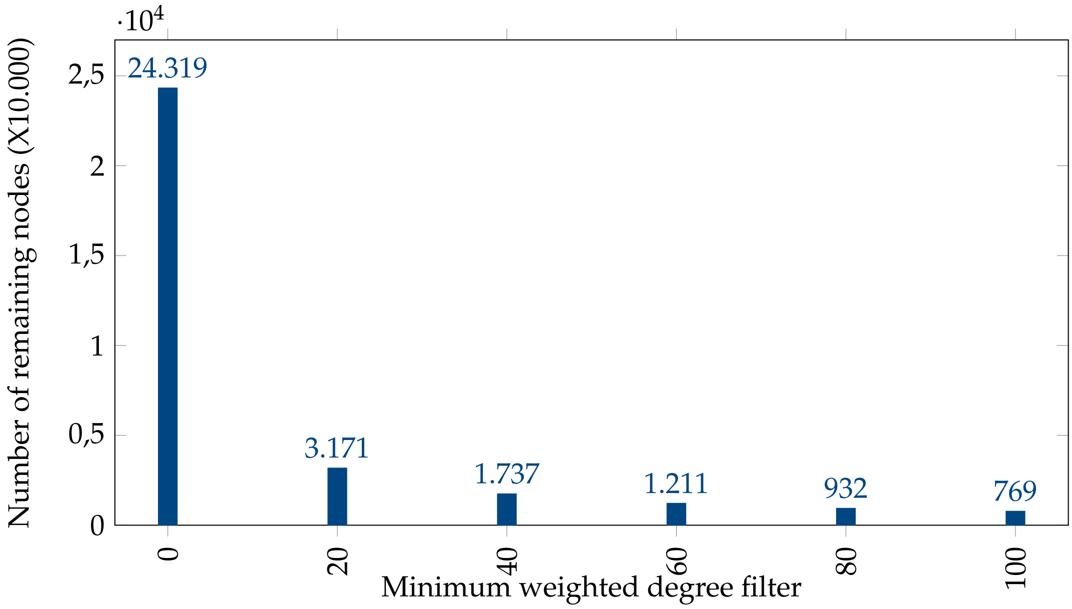

FIGURE 4.3: Number of remaining nodes when applying a minimum weighted degree filter for the values 0, 20, 40, 60, 80, 100

As visible in Figure 4.3, the scale of the y-axis is greatly influenced by the 24,319 nodes when no filtering is applied. Therefore, Figure 4.4 represents the remaining nodes for filtering values equal or above one hundred. Note that the weighted degrees of the hashtags have a long tail distribution, eg., a substantial proportion of the hashtags exhibit infrequent usage, whereas a minute fraction of the hashtags are used extensively.

To define a proper filter threshold, we analyze when the number of nodes that get filtered out stagnated. Looking at Figure 4.4, the number of remaining nodes is still declining sharply for the filters one hundred, two hundred, and three hundred. The difference in remaining nodes for the filters two hundred and three hundred is 150, while the difference in remaining nodes for the filters three hundred and four hundred is only seventy five. When using filter intervals of hundred, this is the first interval where the remaining nodes did not decrease with hundred. The filter that is used for the visualization is therefore set on a minimum of three hundred. Applying this filter, 284 (1.2%) nodes remain which is an appropriate amount to work with.

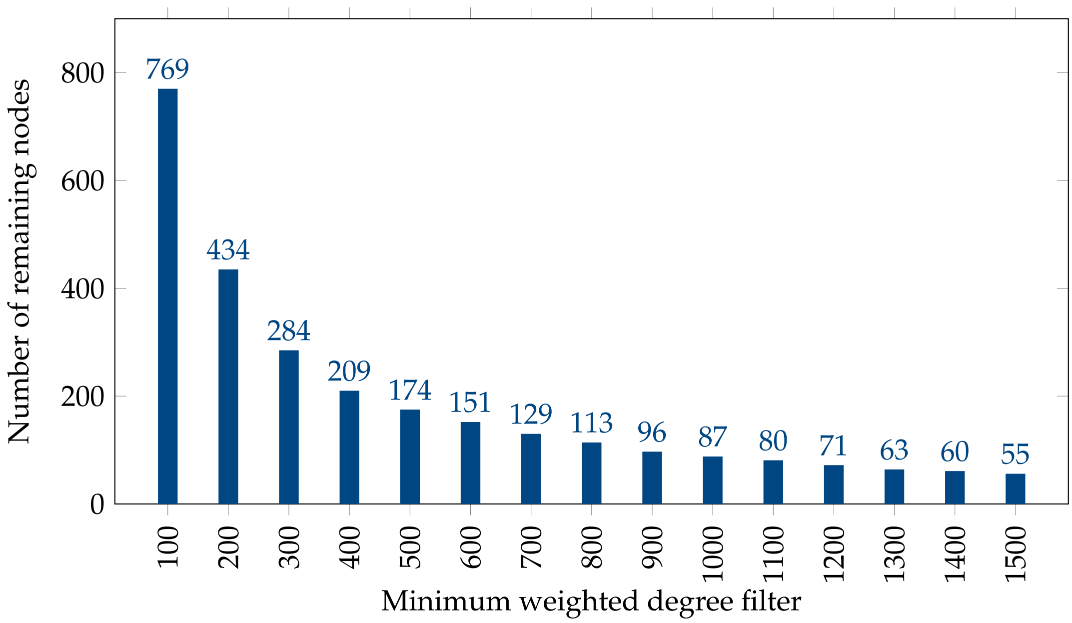

FIGURE 4.4: Number of remaining nodes when applying a minimum weighted degree filter for values equal or above one hundred. Values having a weighted degree of less than one hundred left out for scaling purposes.

4.2.2 Categorization

Hashtags

Following the limitation of twelve colors as defined by RockContent (2022), the remaining nodes after filtering are classified into a maximum of twelve different topics. Additionally, a dark-grey color is assigned to the "other" category, and a light-grey color is assigned to the "uncategorized" category. It is important to note that the "uncategorized" category is only used in the visualization when no filters are applied, and it is not extensively analyzed. Figure 4.5 provides an overview of the defined categories, their respective colors, and the percentage of nodes belonging to each category.

Uncategorized	(98,83%)
War	(0,22%)
Covid	(0,18%)
Location	(0,12%)
Political person / party	(0,11%)
Media	(0,1%)
Politics	(0,09%)
Other	(0,08%)
Police / Protests	(0,08%)
Economy	(0,08%)
Eurovision 2022	(0,05%)
Energy	(0,04%)
Farmers	(0,02%)
Climate	(0,01%)

FIGURE 4.5: Defined categories and their assigned color based on predefined color values by ColorBrewer (2021).

Mentions

Following the limitation of twelve colors, the modularity classes are assigned colors. For this network the ten largest classes are chosen assign a color to, as the classes outside the top contain not more than ten persons. The ten largest modularity classes contain 98.31% of the users in the social network. In Figure 4.6, the assigned colors per modularity classes are visible. It should be noted that the number assigned to the modularity classes is randomly chosen. The number therefore is <u>not</u> referring to data such as size or density.

1	(32,57%)
4	(30,85%)
7	(12,68%)
10	(9,7%)
2	(3,28%)
0	(2,53%)
14	(2,5%)
16	(2,43%)
6	(1,49%)
19	(0,28%)

FIGURE 4.6: Modularity classes and their assigned color based on predefined color values by ColorBrewer (2021).

4.3 Hashtags

Visualizations that contribute to answering the exploratory research questions of this study were devised in the research method. An overview of the visualizations, together with the applied filters and remaining nodes and edges can be found in Table 4.1. In order to answer the research questions related to the Twitter data that was used for this study, visualization are made using the program Gephi together with the algorithm Force Atlas 2. The results are shown in the section below, along with the main findings. The larger versions of all visualisations can be found in Appendix A.

Visualization	**Goal**	**Applied filters**	**Remaining nodes/edges**
4.7 Global overview of data	analyze the positions of all nodes without filtering	-	24,319N
4.8 Global overview of data with with outliers filtered out	identify difference with the previous visualization	$WD \geq 20$	3,312N
4.9 Commonly used topics; position overview	identify most used hashtags	$WD \geq 300$	284N
4.10 Commonly used topics; relationships overview	identify the global relationships between the topics	$WD \geq 300$	284N 10,465E
A.5 Tight relationships between topics overview	identify the strong relationships between the topics	$EW \geq 200$ & $D \geq 1$	46N 150E
4.12 Positions of hashtags referring to media	analyze the position of the hashtags referring to media	$WD \geq 300$	284N
4.14 Relationships of hashtags referring to media	identify the relationships of the hashtags referring to media	$WD \geq 300$ & $D \geq 1$ & ES OR ET = Media	239N 1,745E
4.15 Tight relationships of hashtags referring to media	identify what topics are often associate to which media hashtags	$WD \geq 300$ & $EW \geq 11$ & $D \geq 1$	51N 150E
4.17 Differences in relationships between media hashtag with and without capital letter	analyzing differences in relationships	ES OR ET = (J)(j)inek	117N 162E
4.18 Commonly used hashtags Feb '21 - Jan '22	identify most used hashtags of within this time span	$WD \geq 300$	82N 1,732E
4.20 Commonly used hashtags Feb '22 - Feb '23	identify most used hashtags of within this time span	$WD \geq 300$	211N 6,214E

TABLE 4.1: Overview of visualisations for hashtags together with the applied filters on the weighted degree (WD), edge weight (EW), degree (D), edge source (ES), and edge target (ET).

4.3.1 Global overview of data

Figure 4.7 represents the entire social network of hashtags after categorization without any filters within the time span of February 2021 to February 2023. The network consists of 24,319 nodes which all represent one unique hashtag. Immediately notable is the circle with nodes around the tight network in the middle, which is due to the lack of connections with the center of the network. In addition to this outer ring of the network, it is noticeable that the colored nodes (nodes with a weighted degree of more than three hundred) are predominantly positioned in the center of the network. While all nodes with a weighted degree less than three hundred are depicted in black in this figure, there appears to be a perceptual variation in darkness on the right side above the center compared to the rest of the network. However, it is important to note that this distinction is the result of the high density present in that specific region of the network.

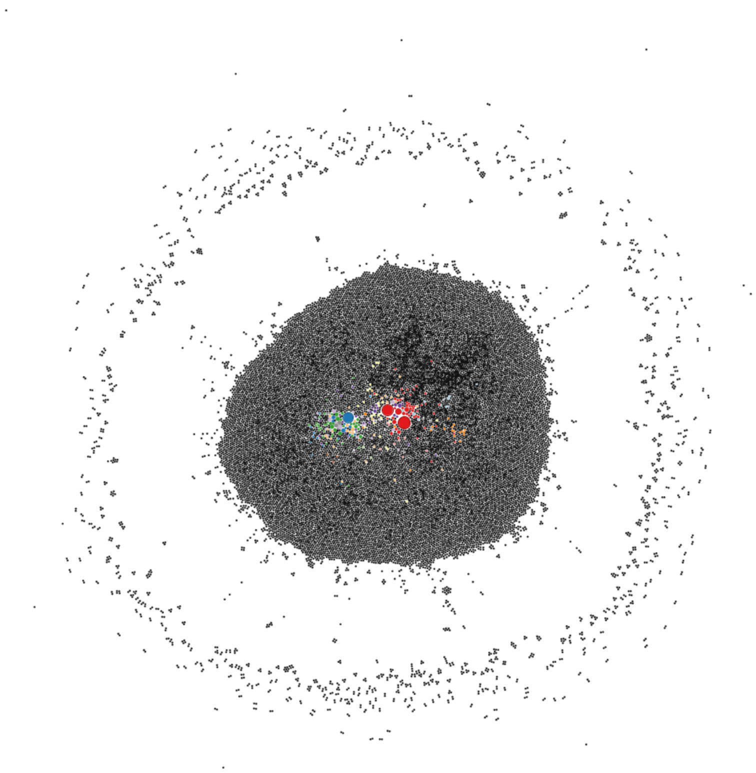

FIGURE 4.7: Categorized social network of hashtags. Filter(s): -. Remaining nodes/edges: 24,319N (100%) Coloring used as defined in Figure 4.5.

4.3.2 Global overview of data with with outliers filtered out

Figure 4.8 depicts the nodes that have a minimum weighted degree of twenty. A notable difference with the previous figure is the significant reduction of the circle surrounding the network, although some remnants are still present. This observation is intriguing as it suggests that despite lacking close connections with the central nodes (based on their position), these nodes still achieve a minimum weighted degree of twenty due to the applied filter.

Regarding the nodes that no longer appear in the central section of the network, it is evident that most of the gaps emerge along the network's periphery. However, a considerable gap also emerges around the larger red nodes representing the "war" category. This finding indicates that even though these hashtags did not meet the filter criterion of $WD \geq 20$, they exhibited a stronger association with the hashtags related to the topic of war.

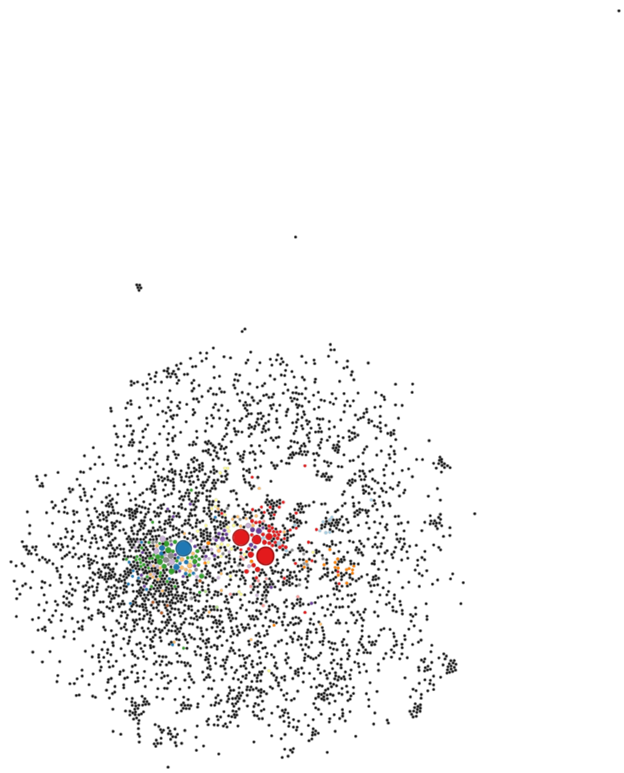

FIGURE 4.8: Categorized social network of hashtags. Filter(s): $WD \geq 20$. Remaining nodes/edges: 3,312N (13.6%). Coloring used as defined in Figure 4.5.

4.3.3 Commonly used topics; position overview

Figure 4.9 illustrates the social network of hashtags that remains after applying the minimum weighted degree filter of three hundred. This visualization enables a thorough examination of the most frequently used hashtags in the data set. The most notable discovery in this network is the presence of two main clusters. One cluster revolves around the frequently used hashtag "demonstratie" (protest), while the other cluster centers around the hashtags "Oekraine" and "Oekraïne" (Ukraine).

The position of a node is based on the relations that it has with other nodes. Clusters therefore emerge when a group of nodes is more closely related to each other compared to their relations with other nodes in the network. It is important to base conclusions regarding the positions of nodes solely on the distances between other nodes. The presence of two clusters in this figure, on the left and right sides, does not imply that their positions would remain the same if the algorithm were rerun. The two clusters could be positioned differently, such as being mirrored or stacked. However, the distances between the nodes would remain unchanged. Therefore, the position of nodes merely represents their relative spatial arrangement. For explanatory purposes however, the two clusters are referred to as the "left" and "right" clusters. The following list presents the categories along with the corresponding number of nodes within the currently depicted network.

- War (54)
- Covid (43)
- Location (30)
- Political person / party (26)
- Media (25)
- Politics (21)
- Other (20)
- Police / Protests (20)
- Economy (19)
- Eurovision 2022 (11)
- Energy (9)
- Farmers (4)
- Climate (2)

Analyzing both clusters, the categories Climate, Covid, Farmers, and Police / Protests are more frequently present in the left cluster. On the other hand, the categories Economy, Energy, Eurovision 2022, and War are more frequently present in the right cluster. The remaining categories Location, Media, Political person / party, and Politics are located in both clusters. Although it is likely that these three are not topic specific, meaning they are linked with different other categories, the location is different.

The category "media" differs from all other categories by being located in the middle of the network. When envisioning two vertical lines; one going straight through the middle of the large blue node "demonstratie" and another one positioned in between the two large red nodes referring to Ukraine, twenty three of the twenty five nodes (92%) categorized as media are located between the two lines. In section 4.3.6, the position of hashtags categorized as media is further analyzed.

Analyzing the category "location", the nodes are located at all positions along and within the two main clusters. However, studying the middle part of the network, the five hashtags located closest to the vertical middle part of network are from top to bottom: "China", "Nederland", "NL", "Groningen", and "groningen".

When doing the same for the category "Political person / party", seven nodes close to the middle are visible. These hashtags are from top to bottom: "D66", "VVD", "Kaag", "Rutte", "Baudet", "rutte", and "kaag".

In addition to the two main clusters, the nodes belonging to the categories Economy and Eurovision 2022 are mainly clustered together with nodes of the same categories, but further from the center of the network and more to the right side.

Furthermore notable, is the hashtags that refer to the same topic, but are written differently. We can distinguish four writing differences; 1. the use of capital letters (Rutte - rutte), 2. the use of diaeresis (Oekraine - Oekraïne), 3. the use of abbreviations (Nederland - NL), and 4. different spelling (Putin - Poetin). The most prominent example of these writing differences can be seen within the category "Eurovision 2022", which consists of 11 hashtags, all written differently.

To conclude this section, a distinction between two clusters can be made. The left cluster is concentrated on topics that directly affect the Netherlands, while the right cluster is more concentrated on the war in Ukraine which affected the Netherlands indirectly.

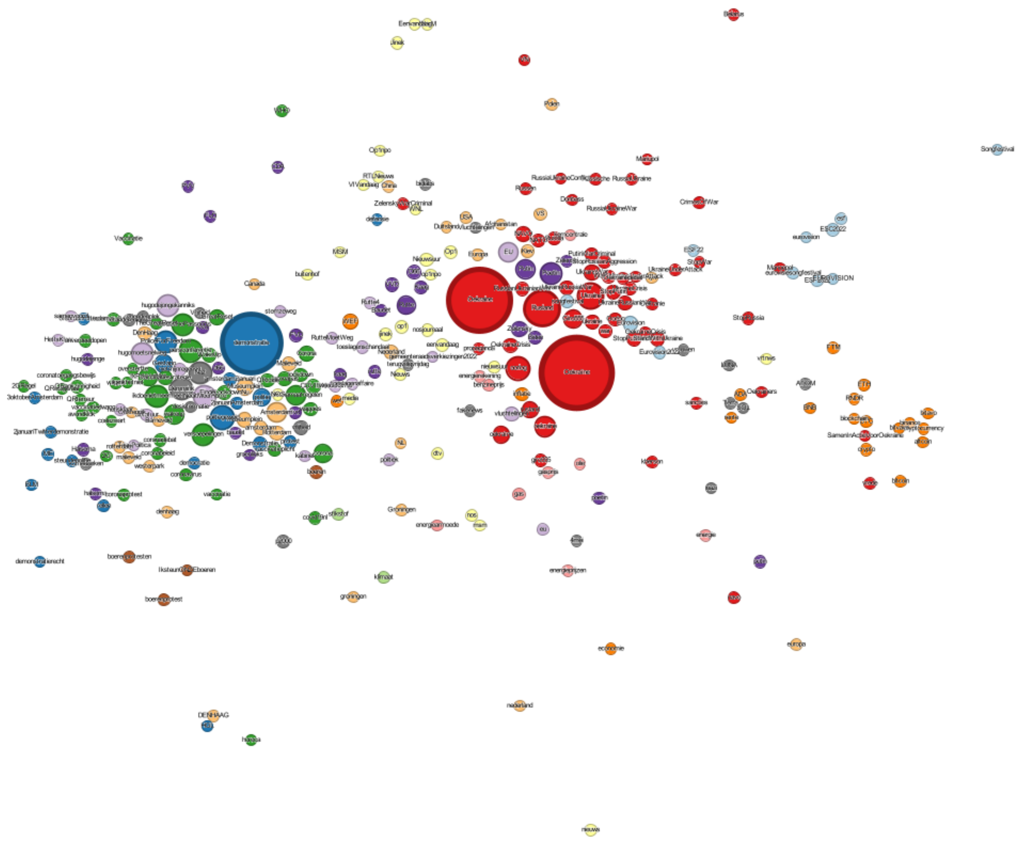

FIGURE 4.9: Categorized social network of hashtags. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N (1.2%). Coloring used as defined in Figure 4.5.

4.3.4 Commonly used topics; relationships overview

The previous section is focused on the positions of the hashtags. In this section, the relationships between the hashtags are described, supported by Figure 4.10. Each line in this network represents the relation between two hashtags, with the thickness of the edge indicating the weight (strength) of the relationship. Since the edge weights in this network fluctuate from 1 to 1,674, the edge weights in this network are rescaled to a range of 1 to 50. This ensures no edges are disturbingly thick, but at the same time still makes it able to distinguish strong relations from weak relations. Furthermore, the color of each edge is a transition from the colors of the connected nodes.

Analyzing the network, it is evident that despite the identification of two distinct clusters in the previous section, there are still numerous connections between the nodes on the right and left sides of the network. The abundance of edges makes it challenging to differentiate between them. Therefore, the upcoming section (4.3.5) focuses on the close relationships within the network.

The edges of the nodes located at the outside of the network are identifiable, and clarify the location of these nodes. Looking at the cluster of the category "Economy", it is visible that most relations are between two nodes from the same category (edges completely orange). Comparing this to the middle of the network, edges of all colors are visible.

Examining the bottom left of the network reveals a significant relationship among the hashtags "DENHAAG," "HGL," and "p2000". Interestingly, despite the relatively thick edge connecting them, the two associated nodes are not positioned in close proximity to the "p2000" hashtag. This observation can be attributed to the influence of other nodes within the network. Specifically, the "p2000" hashtag demonstrates numerous connections with nodes located in the central region, whereas the other two hashtags lack such connections.

It is noteworthy that nodes without direct relationships tend to exert a repelling effect on each other, contributing to the spatial distribution of the network. In this case, the presence of additional connections of the "p2000" hashtag with central nodes results in its displacement from the related nodes "DENHAAG" and "HGL."

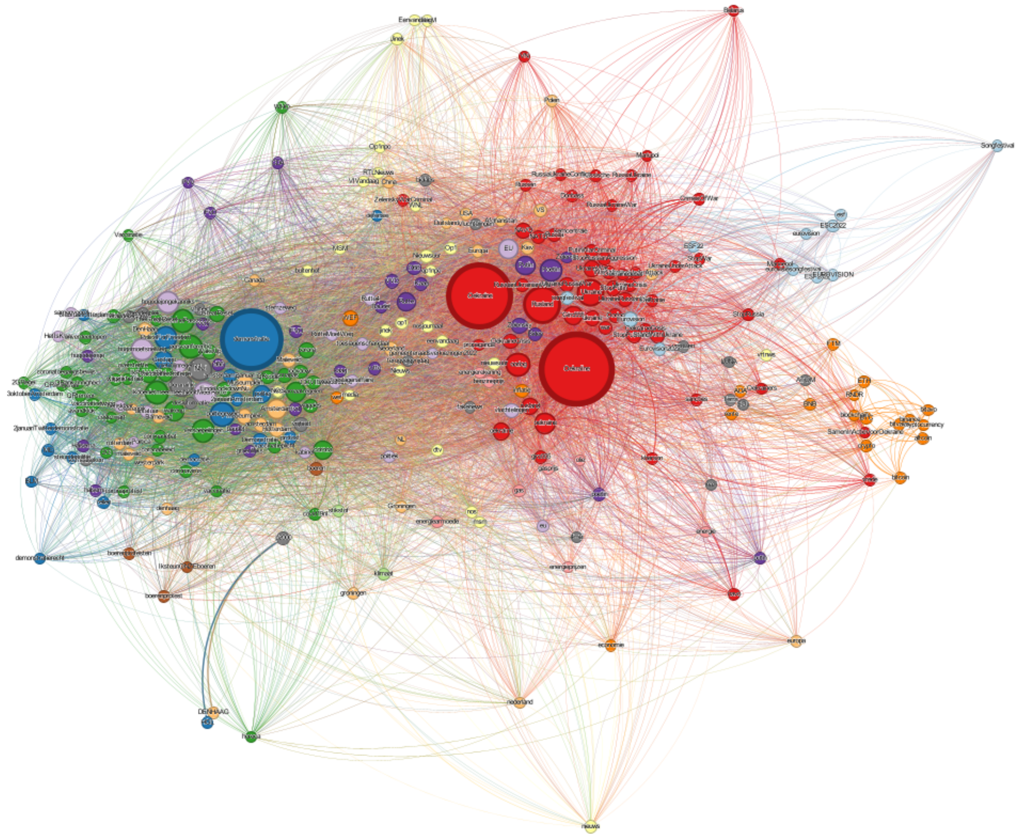

FIGURE 4.10: Categorized social network of hashtags. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N(1.2%) - 10,465E (6.6%)). Coloring used as defined in Figure 4.5.

4.3.5 Tight relationships between topics overview

To identify which hashtags are most used in combination with each other, Figure 4.3.5 is created. Filtering is used to remain the 150 most used relations, which result in an edge weight filter of two hundred. Visible are two main clusters and a third smaller cluster which only consist of three nodes.

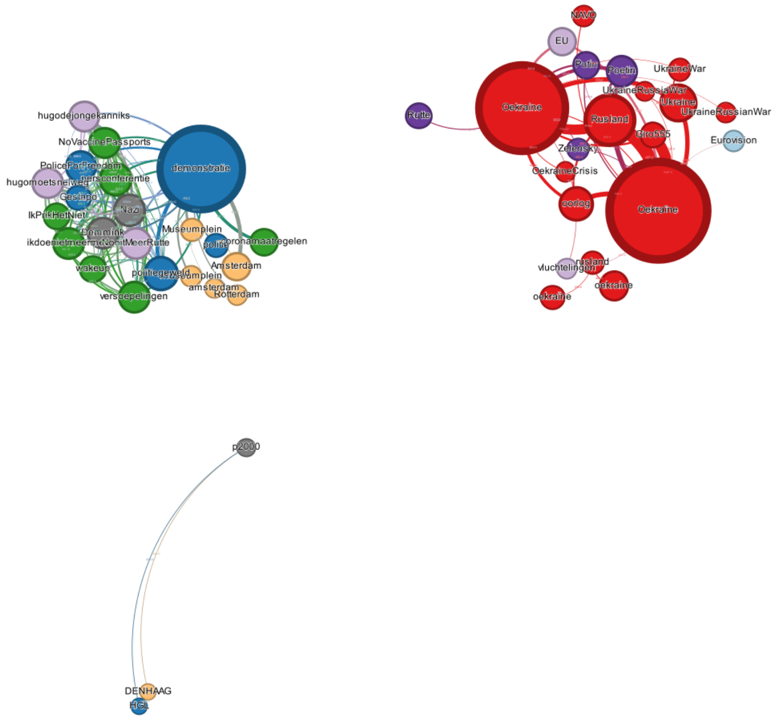

FIGURE 4.11: Categorized social network of strongest relationships between hashtags. Filter(s): $EW \geq 200$. Remaining nodes/edges: 46N(0.2%) - 150E(0.09%). Coloring used as defined in Figure 4.5

The strongest relationships are located in the right cluster, visualized with the thicker edges. The strongest relationship is "Rusland" - "Oekraïne" (EW 1,674), followed by the relation "Giro555" - "Oekraïne" (EW 1,421). Analyzing the categories, it is visible that eight are remaining. Eurovision 2022 only has one node remaining, which is related 205 times to the hashtag "Oekraïne". The category war is present the most with fourteen nodes. Another result that should be outlined is visible when looking at the nodes categorized as political person / party. The four remaining nodes all refer to presidents of the countries; Russia (two times), Ukraine, and the Netherlands. The hashtag referring to the Dutch president, is positioned on the left side of the rights clusters, indicating stronger relations with the cluster on the left, compared to the other presidents. These relations however, do not have weights higher than two hundred, as they are not visible after filtering.

The left cluster consists of five categories, covid being the most present (33.3%). The large number of edges between the green nodes indicates strong relations within this category. Another notable aspect of this network is the presence of three nodes located at the bottom. These nodes demonstrate relatively equal and strong connections with each other. This pattern can be attributed to automated tweets that track police reports, which likely generate similar content and establish mutual relationships among these nodes.

The category "Location" is represented by five nodes in the network, all of which exhibit connections with the prominent blue node "demonstratie". Additionally, among these nodes, "Museumplein" and "Amsterdam" are also linked to each other, suggesting a relationship between them. As this suggest strong relations between the category Location and Police / Protests, an ANOVA test is executed to examine the average edge weight between the category Location and the other categories. For this test, the assumption is made that the category location is connected to all other groups equally (H0). The results of the test are visible in Table 4.2.

Category	Police/Protests	Covid	Political pers/part	Media	Politics	Other	Economy	Eurovision	Energy	Farmers	Climate
μ	27.73	6.10	4.07	2.13	4.51	7.01	1.82	1.65	1.87	5.44	2.72
War	**20.64**	0.99	3.02	4.96	2.59	0.08	5.27	5.45	5.22	1.65	4.37
Police/Protests	0	**21.63**	**23.66**	**25.6**	**23.23**	**20.72**	**25.91**	**26.09**	**25.86**	**22.3**	**25.01**
Covid	**21.63**	0	2.04	3.97	1.6	0.91	4.28	4.46	4.23	0.67	3.38
Polit. pers/part	**23.66**	2.04	0	1.94	0.44	2.94	2.25	2.42	2.2	1.37	1.35
Media	**25.6**	3.97	1.94	0	2.37	4.88	0.31	0.49	0.26	3.31	0.59
Politics	**23.23**	1.6	0.44	2.37	0	2.5	2.68	2.86	2.63	0.93	1.78
Other	**20.72**	0.91	2.94	4.88	2.5	0	5.19	5.37	5.14	1.57	4.29
Economy	**25.91**	4.28	2.25	0.31	2.68	5.19	0	0.18	0.049	3.61	0.9
Eurovision	**26.09**	4.46	2.42	0.49	2.86	5.37	0.18	0	0.23	3.79	1.08
Energy	**25.86**	4.23	2.2	0.26	2.63	5.14	0.049	0.23	0	3.56	0.85
Farmers	**22.3**	0.67	1.37	3.31	0.93	1.57	3.61	3.79	3.56	0	2.72

TABLE 4.2: Absolute difference between means of edge weights with category location. Values in bold indicating significant differences with $p < .01$.

The outcome of the One-Way ANOVA test is a p-value of $< .01$, indicating that some of the category averages are considered to be unequal. Consequently, the null hypothesis (H0) is rejected. Table 4.2 demonstrates that the significant p-value is the primary result of the mean differences with the category Police / Protests. This finding aligns with the network analysis. Based on the results, it can be concluded that within the analyzed protest-related social network, hashtags referring to locations exhibit a significant association with hashtags related to Police / Protests compared to their relationships with other categories.

4.3.6 Positions of hashtags referring to media

As described in section 4.3.3, hashtags categorized as media are located closely to the vertical center of the network. To verify whether this observation is statistically significant, a Chi-Squared test is applied. This test a non-parametric statistical analyzing method often used in experimental work where the data consist in frequencies or "counts" (Zibran, 2007). To use this test, the central region of the network is determined. This allows to compare the expected occurrences of a category in the central region to the actual occurrences of the category. Since two main clusters are identified, the central region of the network is in this case the space right from the middle of the center of node "demonstratie", to the horizontal position between the two large nodes referring to Ukraine. The central region is depicted in Figure 4.12.

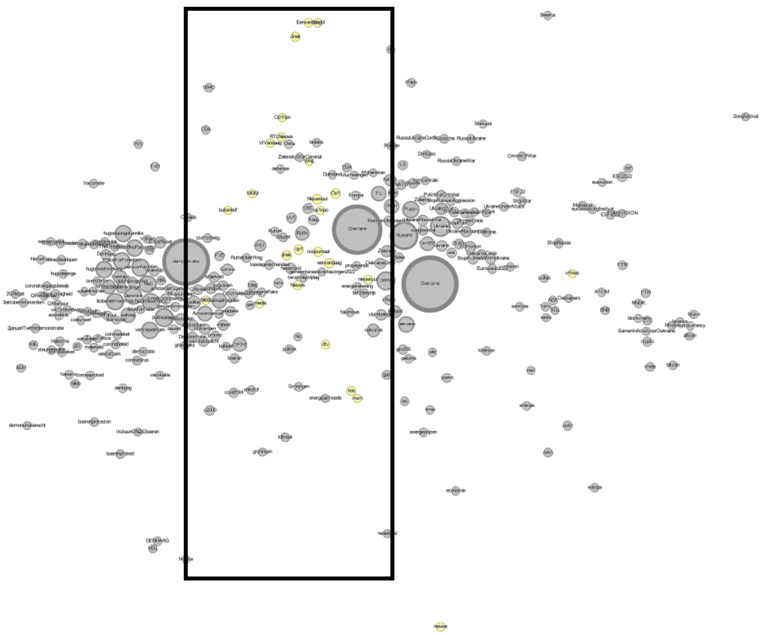

FIGURE 4.12: Position of hashtags categorized as media. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N(1.2%)

To perform the Chi-Squared test, a rule of thumb is used which states that the expected frequencies for all categories being tested should be above five (Cochran, 1954). However, newer calculation tools have shown that this test can also be executed when at least 80% of the expected frequencies are above five (Kingdom, 2023). For this study, the ten categories with the highest number of nodes categorized in them were selected, resulting in one expected frequency of four, which meets the criterion by having 90% above five. The ten used categories together have 92 out of the 269 nodes (34.2%) within the defined center. The hypotheses are defined as follows:

- **Null Hypothesis (H0):** There is no significant association between node category and presence in the central region of the network.

- **Alternative Hypothesis (HA):** There is a significant association between node category and presence in the central region of the network.

H0 suggests that there is no association between the node category and the presence in the central region of the network. The expected frequencies are calculated based on the assumption that each category has an equal proportion (34.2%) of its nodes positioned within the center. The expected frequencies can be obtained by multiplying the total number of nodes in the category by 0.342. The comparison between expected and observed frequencies is shown in Figure 4.13. This figure highlights the difference for the media category compared to the other categories.

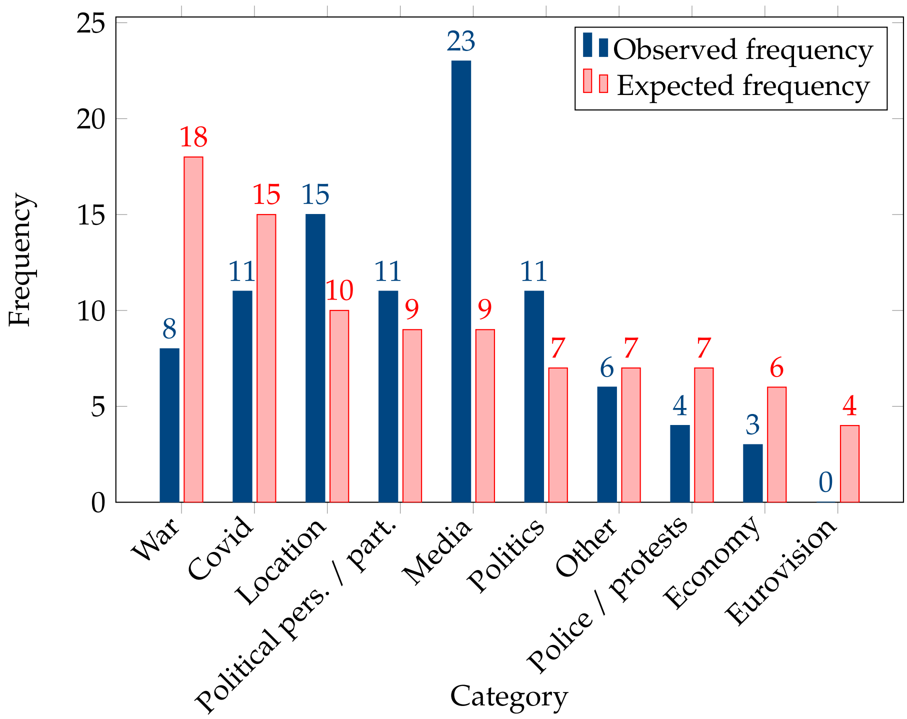

FIGURE 4.13: Category Frequencies (used for Chi-squared test, p < 0.01)

The outcome of the Chi-Squared test is a p-value of $< .01$, indicating that the null hypothesis can be rejected. It can therefore be concluded that there are associations between the node category and the presence in the central region of the network. The media category contributes the most to this outcome with a difference of fourteen, indicating that hashtags categorized as media have a significantly higher chance of being positioned in this central region.

The horizontal position of hashtags related to media can be explained by their connections to hashtags in both clusters. This positioning is likely influenced by news items published by the media that focus on either protests or the war. However, this explanation does not account for the observed differences on the vertical axis.

As shown in Figure 4.12, variations in height are noticeable. The hashtags "Jinek," "Eenvandaag," and "DitisM" are positioned at the top of the network, while the hashtag "nieuws" is located at the bottom. Possible reasons for these differences can be identified by examining the relationships of these hashtags, which will be discussed in the following section.

4.3.7 Relationships of hashtags referring to media

Figure 4.14 illustrates the relationships among hashtags related to media. It is noteworthy that the majority of these relationships are concentrated in the middle of the network. When examining the two clusters, it appears that media hashtags are more closely associated with the right cluster, as indicated by their proximity and connections. However, there is one notable exception with the hashtag "NOS," which is located close to the left cluster. Conversely, the hashtag "vrtnws" is positioned on the right side of the network, possibly due to its reference to a Belgian news channel, which distinguishes it from other media hashtags.

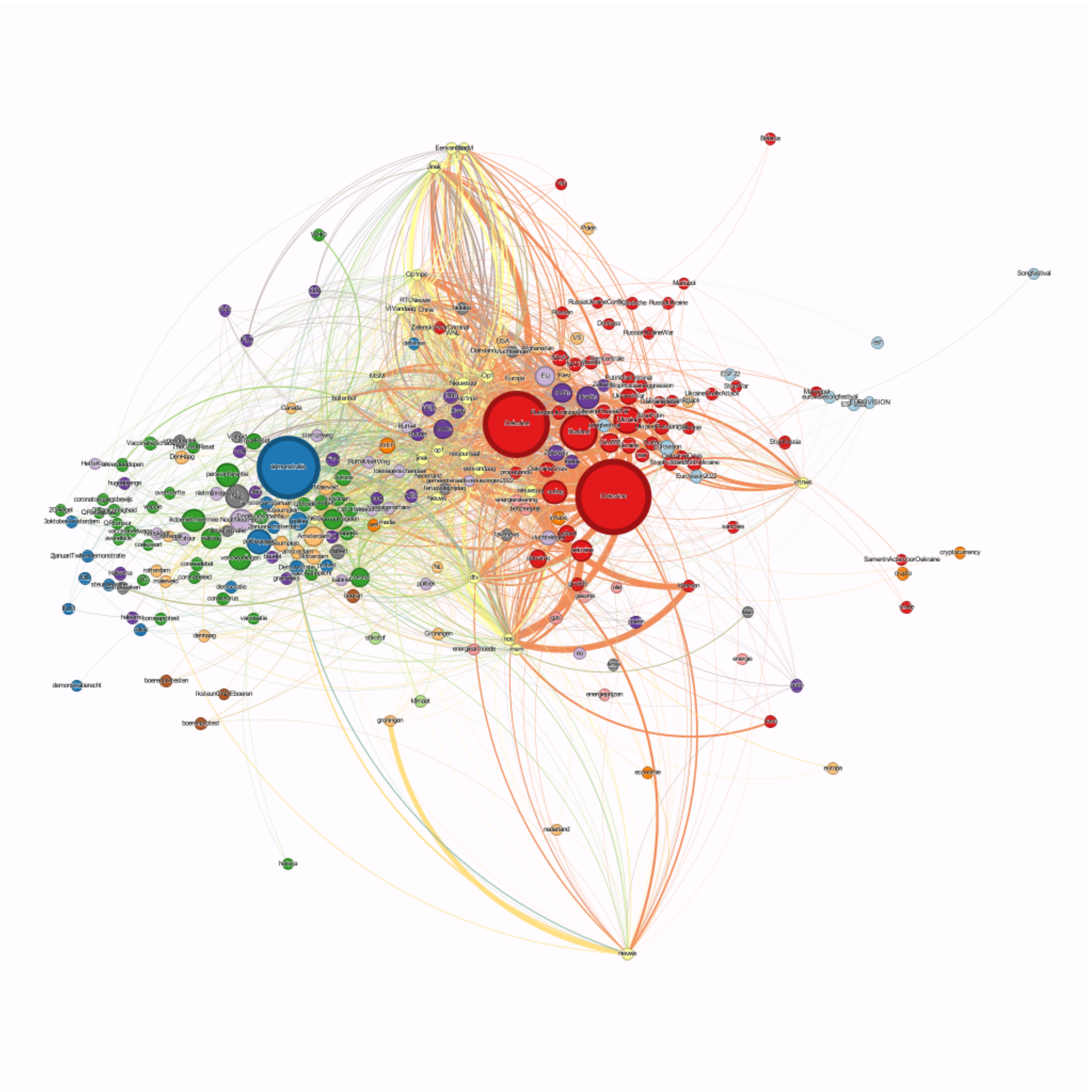

FIGURE 4.14: Relationships of hashtags categorized as media. Filter(s): $WD \geq 300$ AND & ES OR ET = Media AND $D \geq 1$. Remaining nodes/edges: 239N(1.0%) - 1,745E(1.1%). Coloring used as defined in figure 4.5

While the previous figure provides insights into the main relationships among media-related hashtags, its complexity makes it challenging to identify strong connections. Figure 4.15 on the other hand, focuses on relationships that occur at least eleven times ($EW \geq 11$). This filter was applied to highlight the top 150 most frequently occurring relations involving media-related hashtags. The figure reveals that a larger number of strong relationships are concentrated on the right side of the network. In contrast, the left network, originally dominated by the "Covid" category, now only retains a single node representing the "corona" hashtag, which is exclusively related to the hashtag "media".

In section 4.3.7, the visually different position of the hashtag "NOS" was mentioned. The figure below shows that this hashtag is the only Media hashtag that is linked to both the hashtags "demonstratie" and "politie". Given these relationships, this further explains the position of the hashtag "NOS". The position of the hashtag "vrtnws" is explained by the strong relations with the nodes representing war, while not having relations with the left cluster of the network.

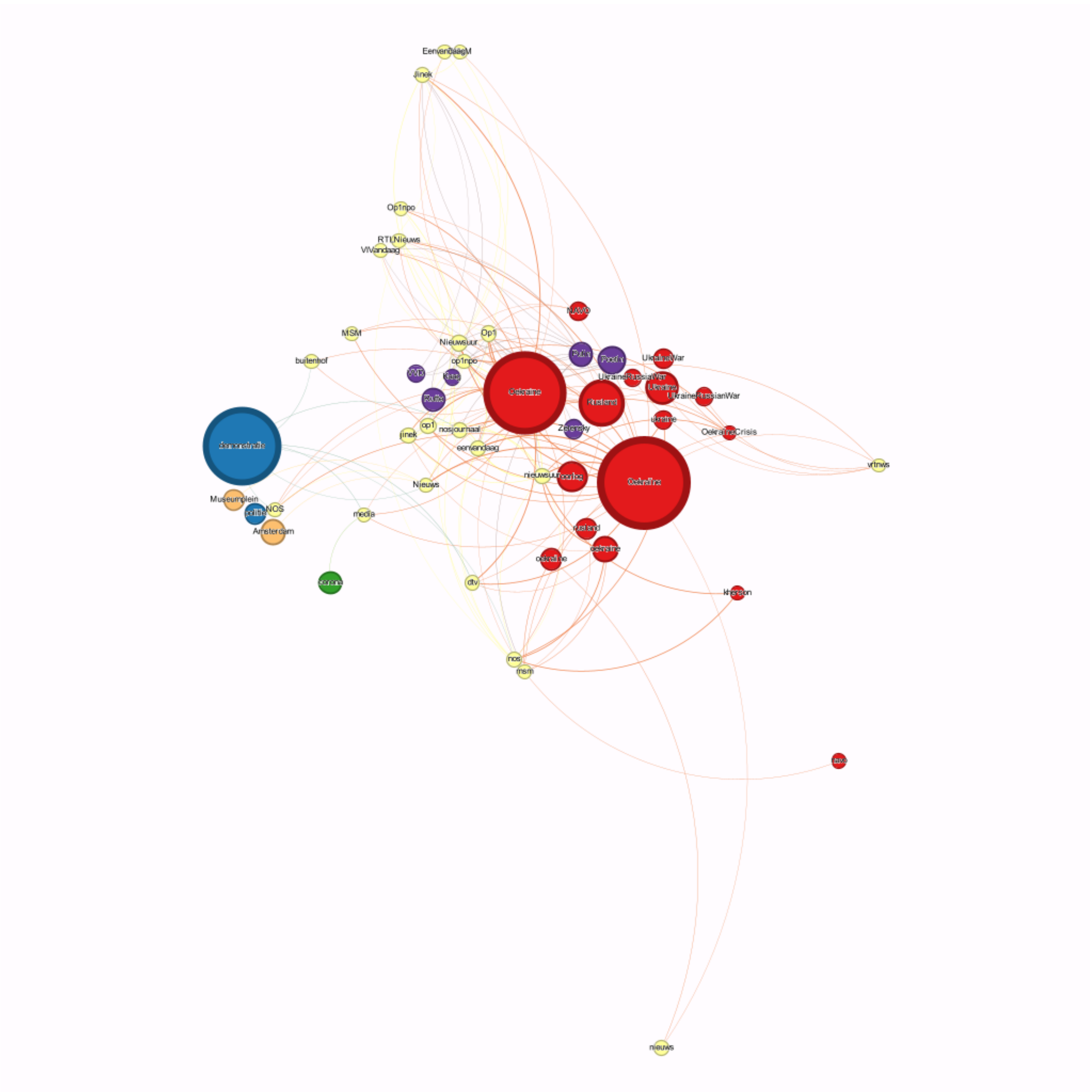

FIGURE 4.15: Tights relationships of hashtags categorized as media. Filter(s): $WD \geq 300$ $EW \geq 11$. Remaining nodes/edges: 51N(0.2%) - 150E(0.09%). Coloring used as defined in Figure 4.5

To test whether the category media does have equally distributed relations with other categories, another Chi-Squared test is performed. The total number of relationships outside the own category is computed for each category (inter-category connections). The assumption is done that all inter-category connections of the media (6,247) are equally distributed compared to the share in the total number of inter-category connections. The expected frequencies can therefore be derived by multiplying the % of the total by 6,247, which is the total number of inter-category connections for the category media. An overview is visible in Table 4.3.

Category	Inter-category connections	% of Total	Expected	Observed
Media	6247	-	-	-
War	25943	16.1	1008.368	3306
Covid	32238	20.1	1253.046	395
Location	14571	9.1	566.354	435
Pol person / part	15466	9.6	601.142	1209
Politics	21572	13.4	838.473	318
Other	14760	9.2	573.701	143
Police / Protests	27651	17.2	1074.756	236
Economy	3088	1.9	120.026	39
Eurovision	1668	1.0	64.833	46
Energy	2186	1.4	84.967	76
Farmers	936	0.6	36.381	23
Climate	642	0.4	24.954	21
SUM	160721	100	6247	6247

TABLE 4.3: Relations with Media - Category Analysis

The hypothesis that are used for the Chi-Squared test are:

- **Null Hypothesis (H0):** There is no significant association between node category and the number of relations with the category media.
- **Alternative Hypothesis (HA):** There is a significant association between node category and the number of relations with the category media.

The outcome of the Chi-Squared test reveals a p-value of $< .01$, indicating the rejection of the null hypothesis (H0) and demonstrating significant associations between the node category and the number of relationships with the Media category. Figure 4.16 visually represents the disparities between observed and expected frequencies. It is noteworthy that the "War" category exhibits three times more relationships with the "Media" category than anticipated. Conversely, the "Police / Protests" and Covid categories display observed relationships that are significantly lower than expected. Therefore, based on the data set employed in this study, it can be inferred that hashtags referring to media are predominantly employed in conjunction with hashtags related to war. Additionally, it can be concluded that hashtags referring to media are infrequently used alongside hashtags referring to Police / Protests and Covid.

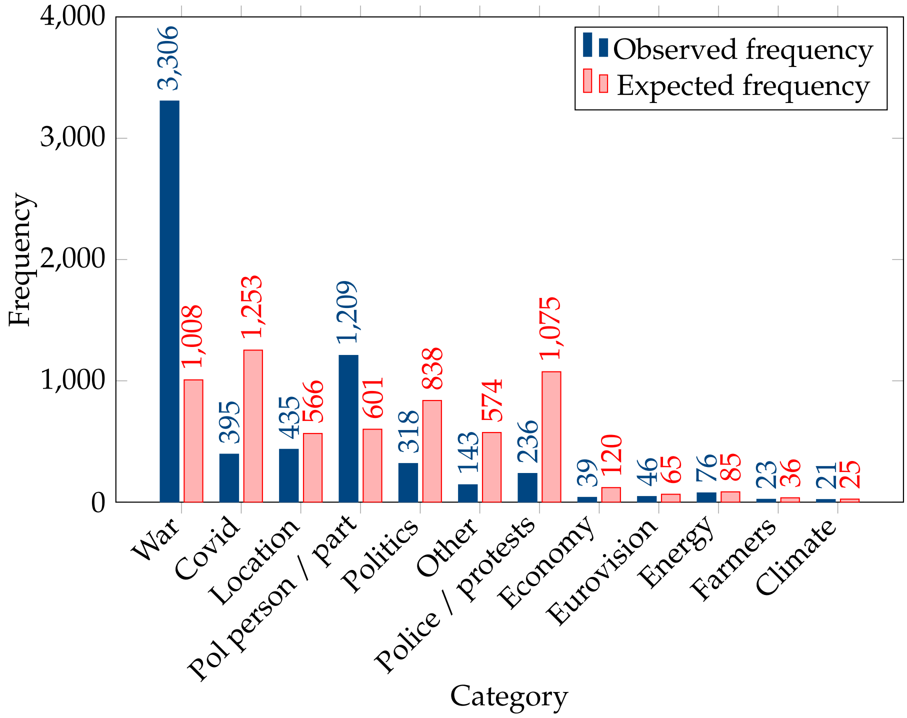

FIGURE 4.16: Category Frequencies (used for Chi-squared test, $p < 0.01$)

To analyze the usage of capital letters within the social network, the example of the hashtag Jinek is examined. This example is chosen due to its significant difference in position within the network. Figure 4.17 depicts the relationships between the hashtag "Jinek" (capitalized) and the hashtag "jinek" (not capitalized). The edges connected to "Jinek" are visualized in blue, while the edges connected to "jinek" are visualized in red.

The hashtag "Jinek" is connected to sixty-six other hashtags, while the hashtag "jinek" is connected to ninety-seven other hashtags. Focusing on relationships that occur twenty times or more ($EW \geq 20$), seven edges remain. Among these, three edges are related to "Jinek" (43%), while the other five edges are related to "jinek" (57%).

The most frequently used relationship is "jinek - op1" (55 times), which explains their close proximity in the network. The second, third, and fourth most used relationships are "Jinek - Oekraine" (54 times), "jinek - Oekraine" (38 times), and "jinek -Oekraïne" (31 times), respectively. Notably, in all three of these relationships, the hashtag "Oekra(i)(ï)ne" is consistently written with a capital letter.

Examining the fifth and sixth most used relationships, the use of capital letters is the same for both sides of the relationships. The fifth most used relationships is "Nieuwsuur - Jinek" (28 times) and the sixth "nieuwsuur - jinek" (27 times). Comparing these relations to a possible mix of capital letter usage, the data shows that the relationship "Jinek - nieuwsuur′ does not exist, "while jinek - Nieuwsuur" is used 10 times. Within this network, twenty-one hashtags exist that are present with both a capital starting letter and a non capital starting letter (e.g., "nieuwsuur", "Nieuwsuur", "rutte", "Rutte"). When seeking for an answer to the question if the capitalized hashtag "Jinek" is connected more often to a capitalized version of one of the 21 hashtags (vice versa for the non capitaliezed hashtag "jinek"), the relations of all

42 hashtags are analyzed. An extensive table with all edge weights is visible in Appendix B.1. When comparing both two writing types of these 21 hashtags, in 76% of the cases a non capitalized hashtag is connected more often to the non capitalized hashtag "jinek" and a capitalized hashtag is connected more often to the capitalized hashtag "Jinek".

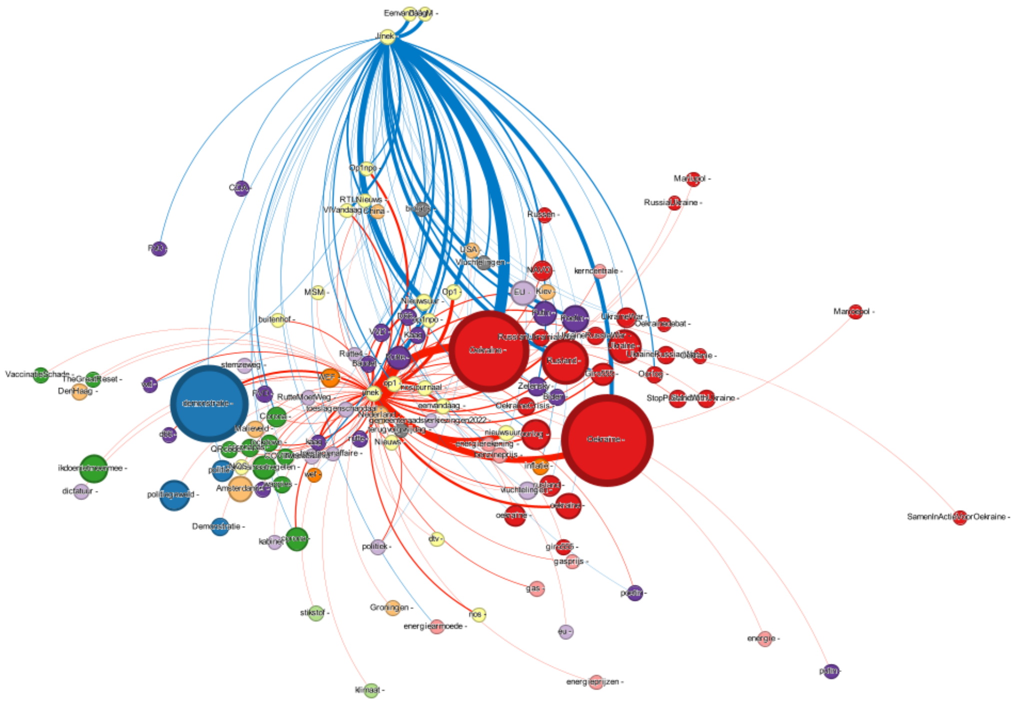

FIGURE 4.17: Relationship differences between the hashtags Jinek and jinek. Edges connected to Jinek visualized in blue. Edges connected to jinek visualized in red. Thickness of edges indicating the edge weight (rescaled to a minimum of 1 and a maximum of 50). Coloring used as defined in Figure 4.5

To test whether significant differences can be found within capital letter usage, four groups are created. The four groups are:

- Non capitalized jinek related to non capitalized other hashtags
- Non capitalized jinek related to capitalized other hashtags
- Capitalized Jinek related to non capitalized other hashtags
- Capitalized Jinek related to capitalized other hashtags

These groups consist of the counts of each combination possible within the "Jinek"/"jinek" hashtag. Both the capitalized and non-capitalized versions of the hashtag "Jinek" were considered, and the number of times they appeared in combination with other hashtags was counted. The null hypothesis (H0) is defined as follows:

- H0: There is no significant association between the use of a capital letter in the hashtag "jinek" and a capitalized other hashtag.

The Chi-Squared test is employed to assess whether the null hypothesis can be accepted because this test is well suited for analysis in which counts of values are used. Among the analyzed relationships between the hashtags "jinek" and "Jinek", it was found that 65.9% of the related hashtags appeared in capitalized form, while 34.1% appeared in non-capitalized form. To calculate the expected occurrences, the total occurrences are multiplied by 0.659 for the capitalized hashtags and 0.341 for the non-capitalized relations. The test results in a p-value of $< .01$, leading to the rejection of H0 and indicating a significant association between the use of a capital letter in the hashtag "jinek" and the use of a capital letter in the related hashtag. The values used in the test are presented in Table 4.4.4.4.

Groups	Observed frequency	Expected frequency
jinek related to non capital related hashtag	170	108 (43.1%)
jinek related to capital related hashtag	145	207 (65.9%)
SUM	315	315
Jinek related to non capital related hashtag	27	89 (43.1%)
Jinek related to capitalized related hashtag	235	173 (65.9%)
SUM	262	262

TABLE 4.4: Observed and expected frequencies for the hashtag Jinek and jinek related to capitalized or non capitalized other hashtags.

4.3.8 Commonly used topics overview between February 2021 and February 2022

Figure 4.18 visualizes the social network of hashtags that have a weighted degree of more than three hundred within the time span of Feb ’21 and Feb ’22. Comparing this network to the original network without date filtering (4.9), this seems to be a zoomed in version of the left cluster. Although the positions of the nodes are unchanged in comparison to the previous networks, the size is. Visual is the large node "demonstratie", which is still the hashtag that is linked to other hashtags most times within this time span. Since this visualization is a result of filter, it is besides interesting to see what was filtered out. In contradiction to the original network, the categories, "War", "Media", "Economy", "Eurovision 2022", "Energy", "Farmers", and "Climate" did not reach the minimum weighted degree of three hundred within this time span. Although these categories are not represented in this network. The category "Corona" is the most representative category within this network. 34 Out of the 43 that were in the original network (79%) have reached the 300 weighted degree within this time span. The category "Police / Protest" is represented with 17 nodes in this network. Compared to the original network with 20 nodes with this category, it can be concluded that almost all hashtags within this category (85%) were predominately used in this time span.

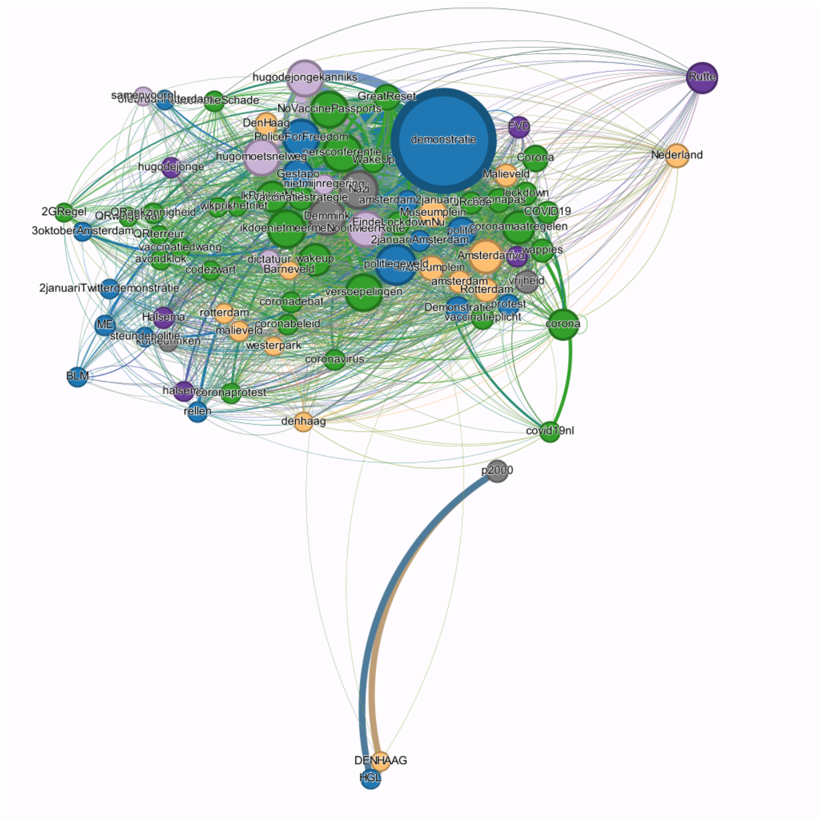

FIGURE 4.18: Categorized social network of hashtags with a minimum weighted degree of 300 between Feb 2021 and Feb 2022. Remaining nodes/edges: 82N(0.3%) - 1,732E(1.1%). Coloring used as defined in Figure 4.5.

Since Figure 4.18 depicts that the hashtag "demonstratie" is the most used hashtag within this time frame, and additionally the query word, a statistical test is ran to check whether this hashtag is used significantly more in certain months compared to other months. This is an important test as this might indicate that large protest were held during this period. As described in the method (Section 3), the test is ran to determine the differences per month. Results are shown in Table 4.5.

Month	**Mar21**	**Apr21**	**May21**	**Jun21**	**Aug21**	**Sep21**	**Oct21**	**Nov21**	**Dec21**	**Jan22**
Feb21	1.27	0.33	5.42	5.76	5.59	1.26	0.46	0.16	3.87	5.58
Mar21	0	1.6	4.16	4.49	4.32	0.0098	1.73	1.43	5.13	6.84
Apr21	1.6	0	5.76	6.09	5.92	1.59	0.13	0.17	3.54	5.24
May21	4.16	5.76	0	0.33	0.17	4.17	5.89	5.58	**9.29**	**11**
Jun21	4.49	6.09	0.33	0	0.17	4.5	6.22	5.92	**9.63**	**11.33**
Aug21	4.32	5.92	0.17	0.17	0	4.33	6.06	5.75	**9.46**	**11.17**
Sep21	0.0098	1.59	4.17	4.5	4.33	0	1.72	1.42	5.13	6.83
Oct21	1.73	0.13	5.89	6.22	6.06	1.72	0	0.31	3.4	5.11
Nov21	1.43	0.17	5.58	5.92	5.75	1.42	0.31	0	3.71	5.42
Dec21	5.13	3.54	**9.29**	**9.63**	**9.46**	5.13	3.4	3.71	0	1.71

TABLE 4.5: Absolute difference between means in daily use of hashtag "demonstratie". Bold values indicate statistical significance at a significance level of 0.05

Additionally to the comparison between months, another test is ran to compare the means of the singular months to the average of this time frame the total amount of hashtags that were used. Results are shown in 4.6 and depicts statistical differences for the months December 2021 and January 2022.

Month	μ	Feb21	Mar21	Apr21	May21	Jun21	Aug21	Sep21	Oct21	Nov21	Dec21	Jan22
AVG	0.15	0.17	0.17	0.19	0.15	0.13	0.15	0.17	0.19	0.17	0.21	0.24
P	-	0.39	0.23	0.08	0.52	0.48	0.83	0.55	0.23	0.33	**0.001**	**0.0001**

TABLE 4.6: Monthly average amount per day of percentage hashtag "demonstratie" compared to all used hashtags bewteen Feb21-Feb22, and Mann-Whitney U Test p-values. Bold values indicate statistical significance at a significance level of 0.01.

The network analysis reveals a strong association between the category "corona" (covid) and the hashtag "demonstratie" (protest). Additionally, specific protest-related hashtags can be identified, such as "2januariTwitterdemonstratie", "2januariamsterdam", and "amsterdam2januari". These hashtags refer to a protest against the corona measures that took place on January 2nd. Figure 4.19 illustrates the most frequently used hashtags in conjunction with other hashtags. Notably, the activity surrounding the protest begins to rise three days prior to the event. In particular, the hashtags "demonstratie" and "Amsterdam" experience a significant surge in usage compared to preceding days. Furthermore, the hashtag "2januariAmsterdam" emerges on the day itself.

On the day of the protest, both "demonstratie" and "Amsterdam" exhibit a remarkable increase of approximately 200% in usage compared to December 30th. In contrast, the hashtag "2januariAmsterdam" experiences a more modest 35% increase in usage. Notably, the hashtag "Amsterdam2januari" is used only twice on December 31st but rises to 382 times on the day of the protest, suggesting its emergence specifically on the protest day.

Another noteworthy observation is the disparity in usage between the generic hashtags "demonstratie" and "protest." The former appears more frequently in the data set, with a total of 5,571 associations with other hashtags, while the latter appears in combination with other hashtags only 752 times. However, no significant differences are found between the two samples.

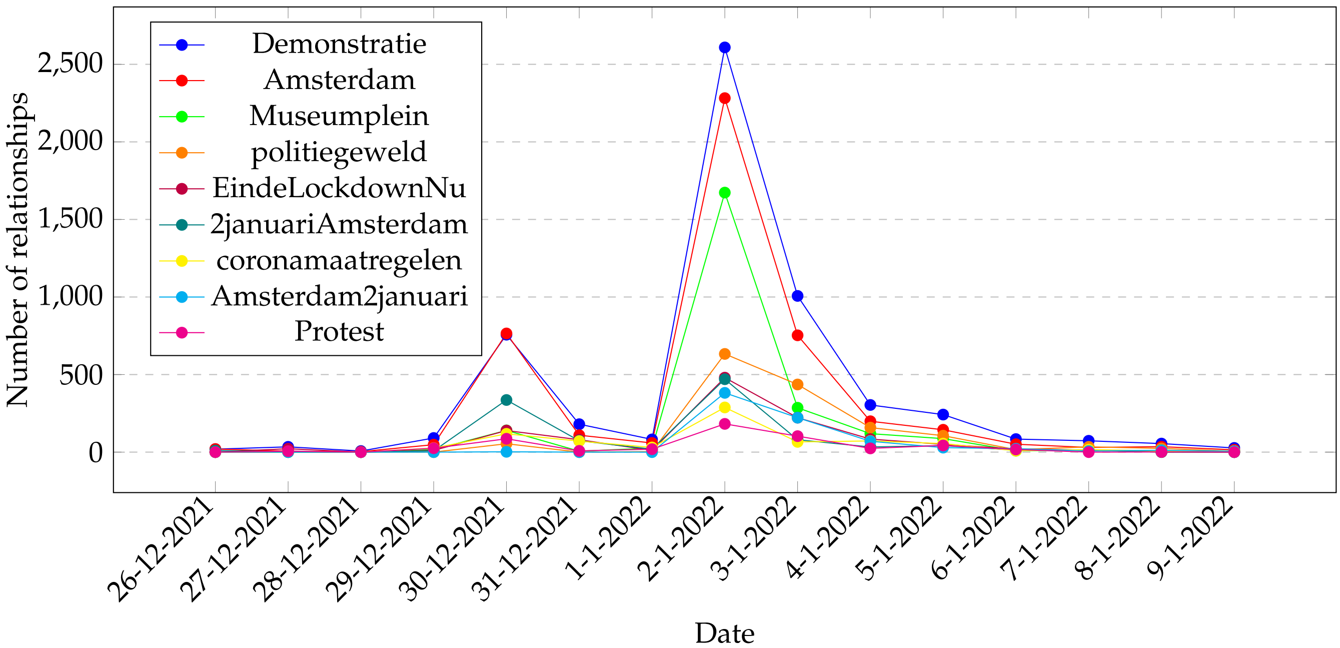

FIGURE 4.19: Frequency of hashtag usage from 26-12-2021 to 9-1-2022

4.3.9 Commonly used topics overview between February 2022 and February 2023

Figure 4.20 visualizes the social network of hashtags that have a weighted degree of more than three hundred within the time span of Feb '22 and Feb '22. Comparing this network to the original network and to the network that was discussed in the previous section (4.3.8), multiple things can be noticed. While the network does not look significantly different from the original network, the left side of the network has changed the most. The blue node "demonstratie", is visualized smaller than in the previous networks, which means it has been related to other hashtags less than the year before. Furthermore, compared the network of the previous year, the topic "war" is positioned as the main category of this year, meaning that all tweets there were retrieved with the filtering word "demonstratie", were most often linked to the war in Ukraine. Comparing this network to the network of the year before, the category "Corona" has been represented with less nodes (15), compared to 34 in the previous year.

Similar to the previous section, a Mann Whitney U test is conducted to identify differences between the mean and the average per month. Results are shown in Table 4.7. While some values are below 0.01, they are not significantly different, which is due to the lacking amount of days for which data was available. Yet, there are significant differences visible for the months May and December.

Month	μ	Feb22	Mar22	Apr22	May22	Jun22	Jul22	Aug22	Nov22	Dec22	Jan23
AVG	0.12	0.09	0.001	0.00	8.68e-05	0.06	0.12	0.17	0.15	0.25	0.14
P	-	0.35	0.03	0.09	**4.165e-8**	0.02	0.47	0.23	0.03	**5.146e-7**	0.28

TABLE 4.7: Monthly average amount per day of percentage hashtag "demonstratie" between Feb22-Feb23 compared to all used hashtags, and Mann-Whitney U Test p-values. Bold values indicate statistical significance at a significance level of 0.01.

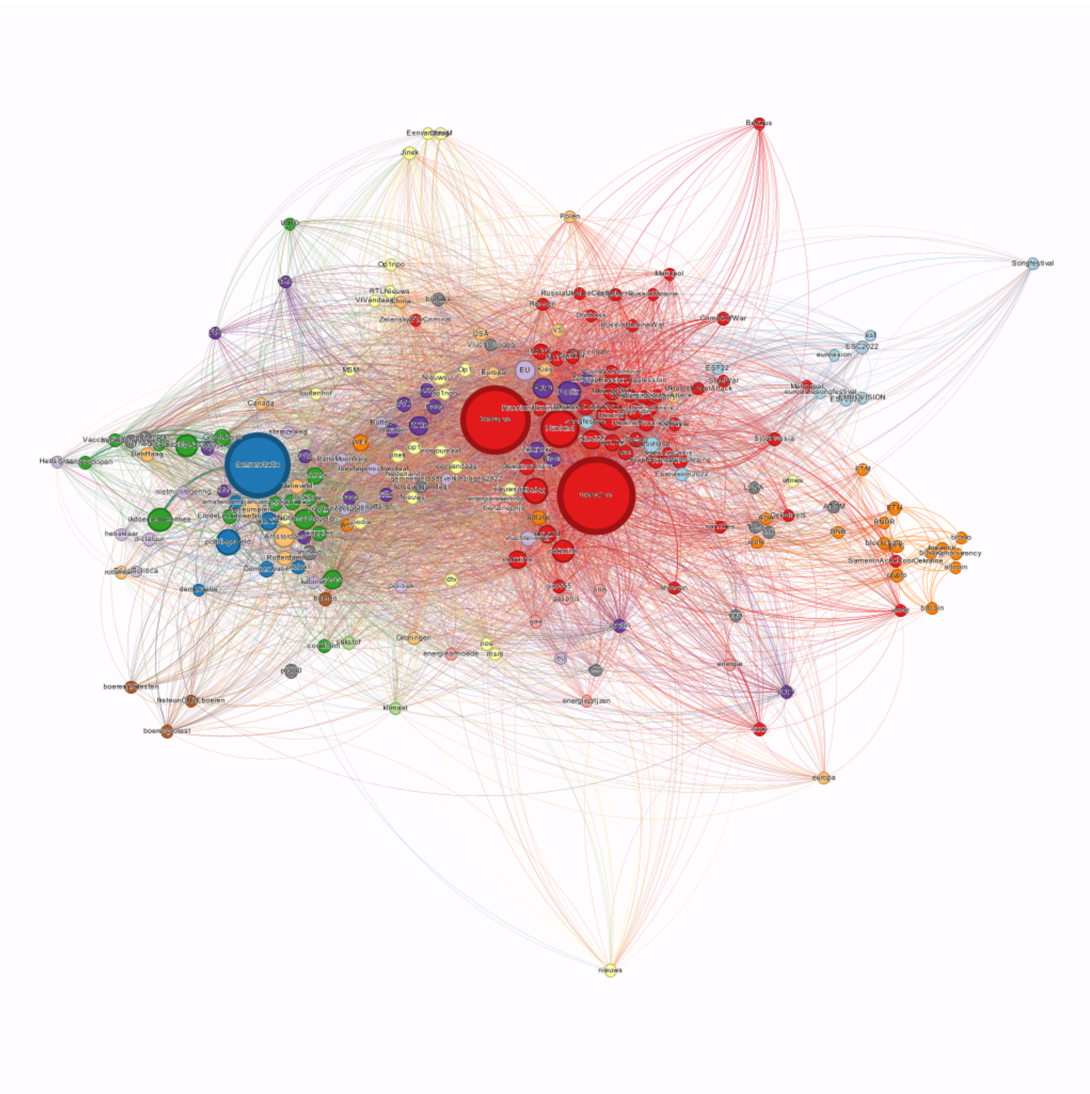

FIGURE 4.20: Categorized social network of hashtags with a minimum weighted degree of 300 between Feb 2022 and Feb 2023. Remaining nodes/edges: 211N(0.9%) - 6,214E(3.9%). Coloring used as defined in Figure 4.5.

4.4 Mentions

4.4.1 network

In Figure 4.21, the social network is shown in which a clustering is made. The clustering is made based on 167,064 tweets that contain mentions. The complete data set consists of 40,832 different users and 187,876 edges. Each node represents a unique twitter user and Gephi created modularity classes (clusters) based on the mentions. The more two persons are mentioned in a single tweet, the higher the chance that they are classified in the same modularity class. When analyzing this network, it is visible that a difference in clusters is visible. Not only does the size of the clusters vary, the clusters at the edge of the network do not contain users of other clusters inside it, while the clusters located more in the center are more intertwined with other modularity classes.

To see if statistical differences can be identified, the centrality of the clusters is compared to each other. A normal distribution within each class was identified which made it possible to apply the one way ANOVA test is to compare the centrality of each class to all other classes. The outcome of this is visible in Table 4.8.

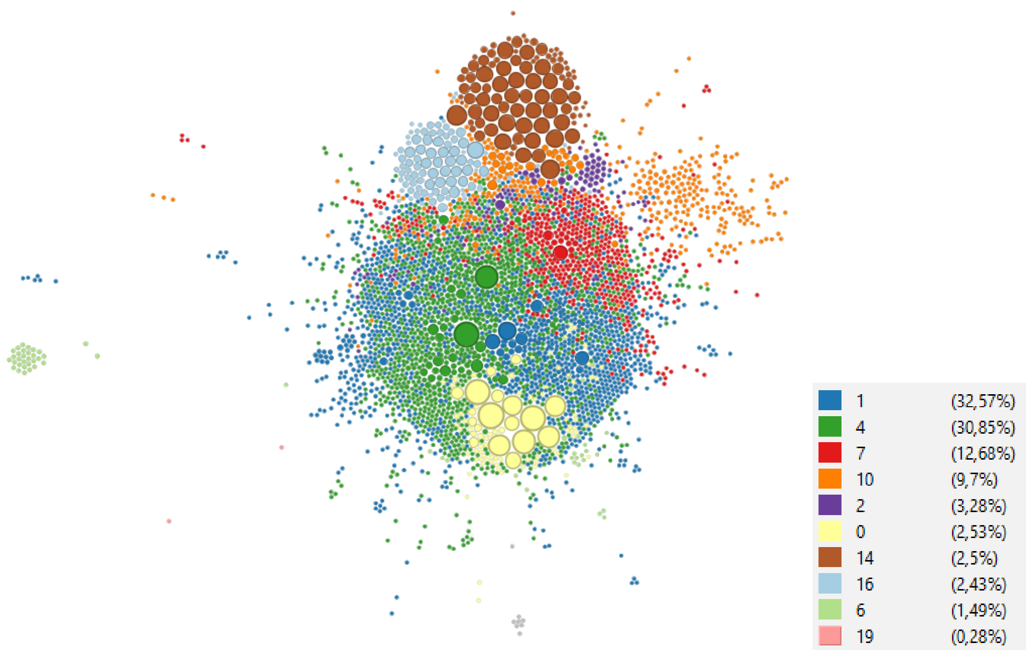

FIGURE 4.21: Social network of users classified in modularity classes based on mentions on Twitter. Filter(s): $W \geq 2$. Legend shown on the right.

The provided table reveals differences in centrality means per class. Values in bold represent significant differences ($p<0.05$). MC14 and MC16 exhibit significant distinction from all other classes. On the other hand, MC2 displays a comparatively lower degree of differentiation, showing significant differences only in relation to MC14 and MC16. Examining the mean centrality values, the largest absolute differences are observed between MC7 and MC14, as well as between MC10 and MC14. It is worth emphasizing that the magnitude of centrality differences can vary depending on the position compared to the center. It is possible for classes positioned closer to each other to exhibit greater differences, even when compared to classes positioned further away from each other. This occurs due to the relative positioning

Class	MC4	MC7	MC10	MC2	MC0	MC14	MC16
MC1	**0.016**	0.007	0.0057	0.012	**0.03**	**0.22**	**0.18**
MC4	0	**0.023**	**0.021**	0.0036	0.014	**0.21**	**0.17**
MC7	**0.023**	0	0.0013	0.019	**0.037**	**0.23**	**0.19**
MC10	**0.021**	0.0013	0	0.018	**0.036**	**0.23**	**0.19**
MC2	0.0036	0.019	0.018	0	0.018	**0.21**	**0.17**
MC0	0.014	**0.037**	**0.036**	0.018	0	**0.19**	**0.15**
MC14	**0.21**	**0.23**	**0.23**	**0.21**	**0.19**	0	**0.039**

TABLE 4.8: Outcomes of the Tukey Kramer test for differences in centrality among classes within the social network. Values are the absolute differences between the means of two classes. Significant values are marked in bold with a p-value < 0.05

in relation to the center, resulting in comparable values for classes equidistant from the center but positioned differently. As a supplementary analysis to the one-way ANOVA test, a two-tailed t-test is conducted to examine potential disparities between the mean centrality of the network and the respective classes. The outcomes of this test for each class are provided in the following list.

MC1 Results of the one-sample t-test indicated that there is a significant small difference between MC1 (M = 0.03 ,SD = 0.06) and the population mean (M = 0.04), t(1288) = 7.9, p < .001, Cohen's d = 0.2.

MC4 Results of the one-sample t-test indicated that there is a non significant very small difference between MC4 (M = 0.04 ,SD = 0.09) and the population mean (M = 0.04), t(1220) = 0.8, p = .432, Cohen's d = 0.02.

MC7 Results of the one-sample t-test indicated that there is a significant medium difference between MC7 (M = 0.02 ,SD = 0.04) and the population mean (M = 0.04), t(501) = 11.8, p < .001, Cohen's d = 0.5.

MC10 Results of the one-sample t-test indicated that there is a significant small difference between MC10 (M = 0.02 ,SD = 0.04) and the population mean (M = 0.04), t(383) = 8.6, p < .001, Cohen's d = 0.4.

MC2 Results of the one-sample t-test indicated that there is a non significant very small difference between MC2 (M = 0.04 ,SD = 0.06) and the population mean (M = 0.04), t(129) = 0.3, p = .766, Cohen's d = 0.03.

MC0 Results of the one-sample t-test indicated that there is a significant small difference between MC0 (M = 0.06 ,SD = 0.08) and the population mean (M = 0.04), t(99) = 2, p = .047, Cohen's d = 0.2.

MC14 Results of the one-sample t-test indicated that there is a significant large difference between MC14 (M = 0.3 ,SD = 0.1) and the population mean (M = 0.04), t(98) = 19.8, p < .001, Cohen's d = 2.

MC16 Results of the one-sample t-test indicated that there is a significant large difference between MC16 (M = 0.2 ,SD = 0.1) and the population mean (M = 0.04), t(95) = 13.4, p < .001, Cohen's d = 1.4.

4.4.2 Group Dynamics

To get insights in the users of the modularity clusters, the users that were mentioned the most are analyzed. Table 4.16 provides insights in the classes. The classes are positioned based on their size which results in the largest class (MC1) being on top and the smallest class (MC19) being on the bottom. The variables M (mentions) and D (degree) provide insights into user interactions within tweets. The M variable indicates how many times a user is mentioned together with another user in a single tweet. The D variable represents the number of different users a user is referenced with in a single tweet. The M/D ratio gives an average measure of the number of times a user is mentioned with the same other user in a single tweet. This ratio can be used as an indicator of the strength or frequency of interactions between specific users. Higher M/D ratios suggest more consistent and repeated mentions between specific users. In this study, individuals who are not publicly known or affiliated with a political party or media channel have been anonymized using an asterisk (*) to ensure privacy and confidentiality.

Analyzing the modularity classes, It is visible that the publicly known persons are classified together in MC1 and MC4. For MC1, the the top 5 users are all related to media and the NOS (position 1) is mentioned almost twice as much as wierdduk (position 5). However, their M/D ration does not differ much, meaning that the average edge strength is roughly equivalent. When analyzing MC4, The usernames indicate that this class contains users that are linked to politics. The two most mentioned users in this class are substantially larger than the other 3 top users out this class. The degree of user "thierrybaudet" is besides the largest, meaning that this user has been present with other users in a tweet most often.

When comparing MC1 and MC4 to the other classes, modularity class 4 can be considered approximately the same size as class 1 in terms of magnitude. When comparing the number of mentions of both top fives, differences are visible. MinPres (user name of Mark Rutte) is mentioned almost 1.5 times the top user of class 1. Furthermore, the difference in M within class 4 is relatively large compared to other classes.

In additional to the large classes, the smaller class MC0, is also notable and exhibits unique characteristics. The main unique characteristic is the M/D ratio, which is close or above one hundred for all top 5 users. This ratio is a result of a large number of mentions (M), in combination with the small degree (D). Since the users within this class all have similar statistics, this indicates that the users within this class are tightly related to each other. This observation holds true for class 14 and class 16 as well, and is supported by Figure 4.21 in which persons of these classes are positioned closely together.

Mod. Class	User 1				User 2				User 3			
	U	M	D	μEW	U	M	D	μEW	U	M	D	μEW
MC1 (1289)	NOS	5098	524	9.7	telegraaf	3856	417	9.2	robdewijk	3456	312	11.1
MC4 (1221)	MinPres	**7459**	448	16.6	thierrybaudet	6860	**677**	10.1	WBHoekstra	3004	208	14.4
MC7 (502)	**************	3782	344	11.0	************	2043	242	8.4	*******	751	132	5.7
MC10 (384)	*********	2004	167	12	***************	1676	132	12.7	************	1429	88	16.2
MC2 (130)	************	1934	221	8.8	************	1115	70	15.9	*********	804	79	10.2
MC0 (100)	*********	**7810**	81	96.4	*******	7588	76	99.8	***********	7456	74	100.8
MC14 (99)	***********	5530	168	32.9	************	5294	127	41.7	*********	4826	121	39.9
MC16 (96)	*******	4389	184	23.9	*******	2349	93	25.3	*******	2317	86	26.9
MC6 (59)	***********	376	75	5.0	**************	376	31	12.1	*******	191	32	6.0
MC19 (11)	*************	65	8	8.1	**********	64	8	8.0	***************	59	8	7.4
Mod. Class	User 4				User 5							
	U	M	D	μEW	U	M	D	μEW				
MC1 (1289)	RTLnieuws	2910	349	8.3	wierdduk	2871	337	8.5				
MC4 (1221)	D66	2406	201	12.0	lientje1967	2402	237	10.1				
MC7 (502)	**********	735	86	8.5	******	709	35	20.3				
MC10 (384)	*****	1419	95	14.9	********	1328	68	19.6				
MC2 (130)	*************	708	12	59	*********	677	12	56.4				
MC0 (100)	****	6992	67	104.4	******	6280	50	**125.6**				
MC14 (99)	********	4529	125	36.2	*******	4473	120	37.3				
MC16 (96)	************	2317	86	26.9	*********	2316	86	86				
MC6 (59)	*************	188	30	6.3	************	187	30	6.2				
MC19 (11)	**********	58	9	6.4	***********	53	8	6.6				

TABLE 4.9: Top 5 most mentioned users per Modularity Class (MC). U representing the username (* used to anonymized accounts that are not publicly linked to media or politics). M representing the amount of times a person is mentioned, D the number of different persons an individual is linked to, and μEW the average Edge Weight (amount of times a person is mentioned with the same other person) which can be retrieved by dividing M by D.

As evident from the preceding table, there are notable variations in the average edge weights among different classes. To assess the significance of these differences, a one-way ANOVA test was conducted, comparing all classes. It is important to note that this test is based on the means of each entire cluster and not exclusively on the data presented in Table 4.16. The results, presented in Table 4.10, indicate the absolute differences in average edge weights. Values in bold indicate statistically significant differences at a significance level of .05. The chosen significance level takes into account the large period of data collection and the substantial size of the groups exhibiting significant differences in characteristics. This ensures that conclusions are not drawn based on chance alone

As shown in the table, groups 2, 10, 14, and 16 show significant differences. Especially the last three mentioned classes have a significant higher average edge weight compared to the other classes. MC0 has the highest average edge weight, but does not significantly differ from class 14 and 16.

Class	MC4	MC7	MC10	MC2	MC0	MC14	MC16
MC1	0.057	0.81	0.95	**2.22**	**38.84**	**18.59**	**9.11**
MC4	0	0.76	0.89	**2.16**	**38.79**	**18.53**	**9.05**
MC7	0.76	0	0.13	1.4	**38.03**	**17.78**	**8.29**
MC10	0.89	0.13	0	1.27	**37.9**	**17.64**	**8.16**
MC2	**2.16**	1.4	1.27	0	**36.63**	**16.37**	**6.89**
MC0	**38.79**	**38.03**	**37.9**	**36.63**	0	20.25	29.74
MC14	**18.53**	**17.78**	**17.64**	**16.37**	**20.25**	0	**9.48**

TABLE 4.10: Absolute differences between the means of average edge weights for different classes in the social network. Values in bold indicate significant differences at a significance level of 0.05.

To compare each individual class with the population mean, one-sample t-tests are used. Each class is compared to the mean of the entire network. Results show similar results as obtained through the ANOVA test; modularity classes fourteen and sixteen show significant large differences compared to the mean. This indicates that these classes are consistent in the high average weight among all participants within the cluster. While, for MC0, for instance, users could also be identified further outside the cluster when analyzing the network.

MC1 Results of the one-sample t-test indicated that there is a significant small difference between MC1 ($M = 5.6$,$SD = 2.9$) and the population mean ($M = 7$), $t(1288) = 17.4$, $p < .001$, Cohen's $d = 0.5$.

MC4 Results of the one-sample t-test indicated that there is a significant small difference between MC1 ($M = 5.7$,$SD = 3.8$) and the population mean ($M = 7$), $t(1220) = 12.4$, $p < .001$, Cohen's $d = 0.4$.

MC7 Results of the one-sample t-test indicated that there is a significant very small difference between MC7 ($M = 6.4$,$SD = 4.5$) and the population mean ($M = 7$), $t(501) = 2.9$, $p = .004$, Cohen's $d = 0.1$.

MC10 Results of the one-sample t-test indicated that there is a non significant very small difference between MC10 ($M = 6.6$,$SD = 5.5$) and the population mean ($M = 7$), $t(383) = 1.6$, $p = .108$, Cohen's $d = 0.08$.

MC2 Results of the one-sample t-test indicated that there is a non significant very small difference between MC2 ($M = 7.8$,$SD = 9.4$) and the population mean ($M = 7$), $t(129) = 1$, $p = .319$, Cohen's $d = 0.09$.

MC0 Results of the one-sample t-test indicated that there is a significant medium difference between MC0 ($M = 44.5$,$SD = 47.3$) and the population mean ($M = 7$), $t(32) = 4.5$, $p < .001$, Cohen's $d = 0.8$.

MC14 Results of the one-sample t-test indicated that there is a **significant large difference** between MC14 ($M = 24.2$,$SD = 15.6$) and the population mean ($M = 7$), $t(98) = 11$, $p < .001$, Cohen's $d = 1.1$.

MC16 Results of the one-sample t-test indicated that there is a **significant large difference** between MC16 ($M = 14.7$,$SD = 8.5$) and the population mean ($M = 7$), $t(95) = 8.9$, $p < .001$, Cohen's $d = 0.9$.

4.5 Hashtags and mentions combined

4.5.1 Visualizations

To get valuable insights into the topics discussed within the modularity classes, the incorporation of hashtags can serve as a useful tool, without the need to analyze the actual tweets. In section 4.3, the main hashtags have already been identified. This section establishes the link between the identified modularity classes, facilitating a comprehensive understanding of the classes and their respective subjects of interest. Note that a larger version of each visualization can be found in the appendix (A).

To get a global overview of the hashtag usage within the social network, Figure 4.22 is created. This network is a representation of all the hashtags that are used by the modularity classes. Although the coloring usage is equal to the coloring used in section 4.4, the figure looks very different. In this network all users that are in the same class, are merged into one node. Each colored node therefore represents one cluster and can be identified by the color. All grey nodes represent hashtags and give information by analyzing the position of the hashtag compared to the classes. The size of the colored nodes represents the size of the cluster, while the size of the grey nodes represents the number this hashtag is used. By using a minimum WD of 2, hashtags that are used once have been filtered out. This ensures that typos and non-important hashtags are out filtered. It is identifiable that each class has unique hashtags that are used only in this class. However, the size of these nodes shows that hashtags that are only used by one class, are mostly nodes not used in large amounts.
In the center of the cluster however, the positioned hashtags have been used by multiple classes, and the size shows that these are the hashtags that have been used most frequent. The ten most used hashtags within this network are named as visible in the figure. Nine of them are positioned in the actual center of the network, while the hashtag "NieuwsInPerspectief" is located on the right of the network. This hashtag is only used by MC0 which is unique for a hashtag that is in the top 10 most used hashtags of this network.

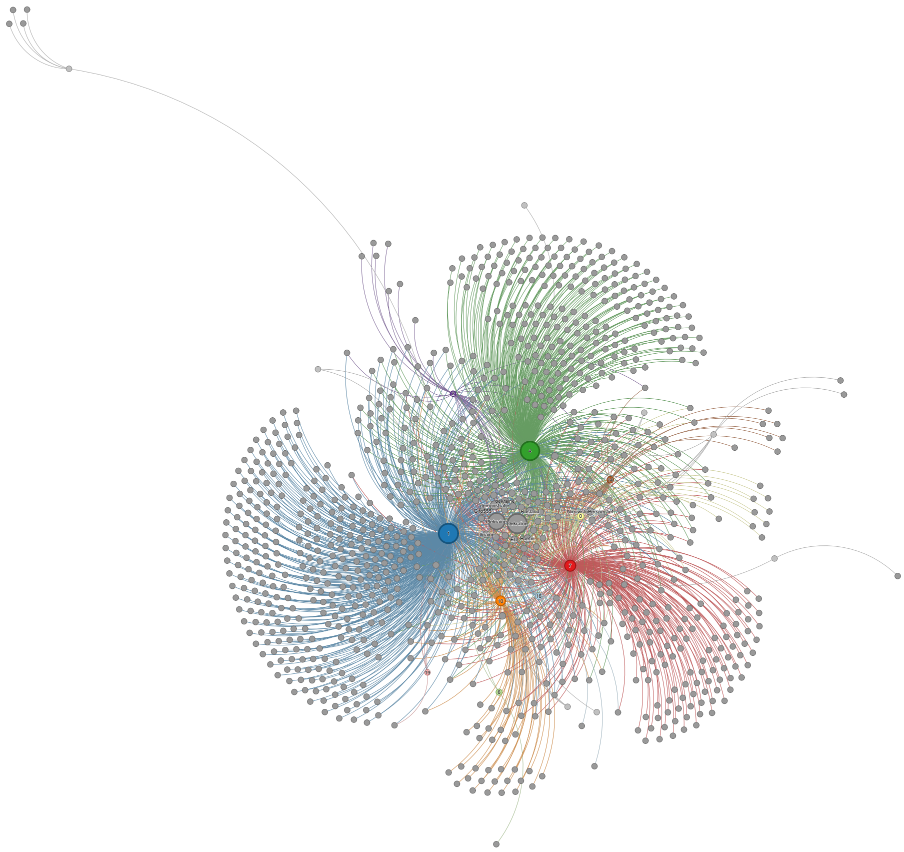

FIGURE 4.22: Social network of hashtags used by the defined modularity classes between 1 Feb 2021 and 31 Jan 2023. Filter(s): $EW \geq 2$. Remaining nodes/edges: 1,030N(24.7%) - 1,685E(28.8%). all edges rescaled to min 0.1 and max 1 for visibility purposes, network scaled X1000.

After analyzing the overall positioning of the hashtags, the element of time is now incorporated to identify changes in topics over time. Figure 4.23 depicts all hashtags that have been used at least five times by one of the classes in 2021. It is evident that a reduced number of hashtags remains, representing the most frequently used ones, with the thickness of the edges indicating the level of usage by a particular class. The smaller classes exhibit fewer remaining connections, indicating a limited number of hashtags that were used more than five times within this time frame. Conversely, the larger classes still retain some edges, and while most connections are associated with the central hashtag, MC1, MC4, and MC7 show notable relationships with hashtags that are not connected to other classes. This indicates that these topics have only been spoken about frequently within these specific classes.

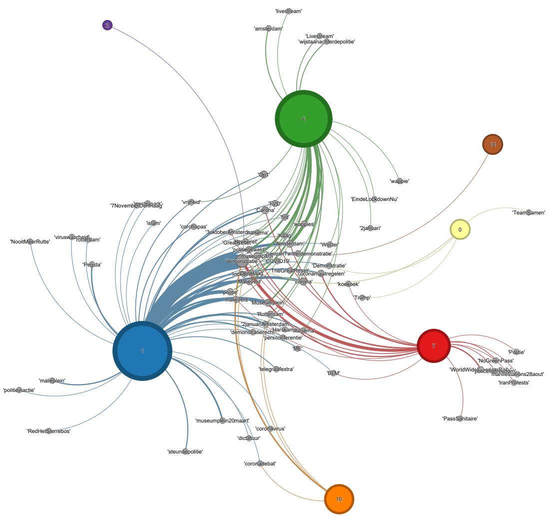

FIGURE 4.23: Social network of hashtags used by the defined modularity classes between 1 Feb 2021 and 31 Jan 2022. Filter(s): $EW \geq 5$ AND $D \geq 1$. Remaining nodes/edges: 77N(1.84%) - 109E(1.86%). Re-scaled edges visible, Network scaled X50.

Now that the network of 2021 has been analyzed, a comparison with 2022 can be made. Figure 4.24 depicts the hashtag usage over the course of 2022. It is immediately evident that the activity has increased, with more nodes and edges visible, indicating that the classes have tweeted more or, at the very least, used more hashtags.Compared to the network of 2021, the number of hashtags has tripled, and the number of edges has almost quadrupled. This demonstrates an increase in both the variety of hashtags used and the frequency of a single hashtag's usage by a particular class. Upon closer analysis of these differences, it is notable that the smaller classes now exhibit a significantly higher number of connections compared to the previous year. Remarkably, when examining hashtags that are specific to a single class, those connected to the smaller classes are positioned closer to the classes themselves, while hashtags exclusively connected to, for example, class 1 and 4, are situated further outside the network. The distance between these hashtags, the respective cluster, and the center provides insights into the active tweeting about these 'unique' topics within these smaller clusters. Although the larger classes also have hashtags that are exclusively used by them, these hashtags are positioned at an equal distance from the class as the cluster is positioned from the center. This indicates that these are not topics that are discussed particularly more than the topics in the center.

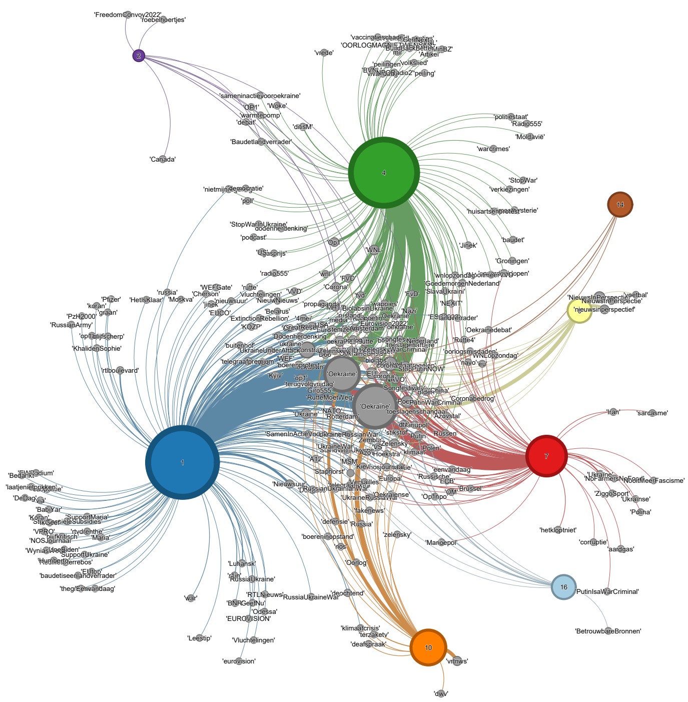

FIGURE 4.24: Social network of hashtags used by the defined modularity classes between 1 Feb 2022 and 31 Jan 2023. Filter(s): $EW \geq 5$. 277N(6.63%) - 433E(7.4%). scaled edges visible, network scaled X50.

4.5.2 Hashtag usage per class

To provide a general overview of the activity per modularity class, Figure 4.25 illustrates the average amount of tweets per user and the average amount of hashtags per hundred tweets for each class. The classes are positioned in order of size, resulting in the largest class MC1 being on the outer left and the smallest class MC16 being on the outer right. The horizontal lines are used to illustrate the averages per category.

Upon analyzing the graph, the results shows large differences per class. The size of the class (shown below class number) does not seem to have an effect on one of the two categories as now trend in one the bars can be found.

Analyzing the tweets per user, the largest difference can be identified between MC4 (minimum) and MC0 (maximum) as the users in the last named class send 2.6 times more tweets per users.

Although MC4 has the lowest value for the tweets per users, it shows the maximum value (34.9) for the hashtags per hundred tweets.

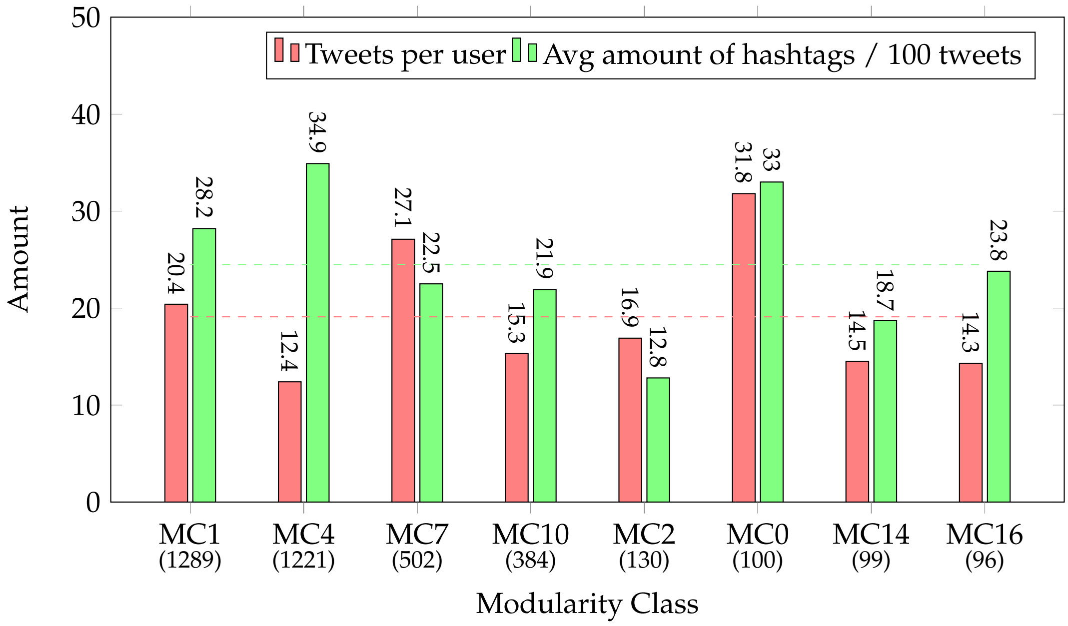

FIGURE 4.25: Frequency of hashtag usage per modularity class. Classes divided along x-axis based on size where MC1 is the largest class, MC4 the second largest class, etc.

To get deeper insights into the topics discussed within the identified modularity classes, an analysis of the most used hashtags is conducted. Consistent with earlier sections of the study, a distinction is made between tweets sent between Feb'21-Feb'22 (2021) and Feb'22-Feb'23 (2022). Table 4.11 presents the six most frequently used hashtags per modularity class, highlighting the variations across classes. A more extensive table in which the top ten hashtags is visible can be found in C.

A noteworthy observation arises when comparing the primary hashtags used in 2021 with those in 2022. A significant shift in subjects can be identified, primarily attributed to the commencement of the war in Ukraine. This shift not only corresponds to a change in the predominant topic of discussion but also results in increased activity, evident from the higher usage of hashtags. This trend is observed across all classes when comparing the usage amounts between 2021 and 2022.

Analyzing the hashtags used in 2021 reveals that the hashtag "demonstratie" is the most commonly employed hashtag in five out of the six analyzed classes, which can be explained by the use of this word as the query to retrieve this data. It is worth noting that for the class in which "demonstratie" is not the top hashtag (MC4), this particular hashtag does not appear among the top ten at all. However, for this class, the hashtag "demonstratie is present in the top ten of 2022", while this is not the case for all other classes.

Additionally, a variety of location-specific hashtags are observed, including city names such as Amsterdam and Rotterdam, names of protest sites within cities such as Museumplein and Malieveld, and combinations of city names with protest dates, for instance, "3oktoberAmsterdam" and "2januariAmsterdam.

When analyzing the hashtags used in 2022, it is evident that two variations of the word 'Oekraine' dominate the top two positions. Interestingly, Modularity Class 0 (MC0) stands out with a different most frequently used hashtag, namely 'NieuwsInPerspectief.' Notably, this hashtag does not appear in the top ten for any other class. Furthermore, it is worth noting that two other variations of the same hashtag are among the top ten hashtags in MC0. This observation, coupled with the findings presented in section 4.4.2, establishes the uniqueness of this class in terms of tie strength and topics.

Class	Hashtag 1				Hashtag 2				Hashtag 3			
	2021		2022		2021		2022		2021		2022	
	#	A	#	**A**	#	**A**	#	**A**	#	**A**	#	**A**
MC1	demonstratie	246	Oekraine	738	politie	57	Oekraine	612	Museumplein	55	Rusland	125
MC4	Amsterdam	50	Oekraïne	536	museumplein	38	Oekraine	438	politiegeweld	27	VWNL	151
MC7	demonstratie	45	Oekraïne	362	WordWideFreedomRally	20	Oekraine	308	coronamaatregelen	19	Poetin	112
MC10	demonstratie	23	Oekraïne	179	Rotterdam	8	vrtnws	134	2januar iAmsterdam	5	Oekraine	111
MC2	demonstratie	8	Oekraine	15	3oktoberAmsterdam	4	Oekraïne	13	politiegeweld	5	FreedomConvoy2022	11
MC0	demonstratie	12	NieuwsInPerspectief	222	Demonstratie	9	Oekraïne	144	TeamSamen-	5	nieuwsinperspectief	117

Class	Hashtag 4				Hashtag 5				Hashtag 6			
	2021		2022		2021		2022		2021		2022	
	#	A	#	**A**	#	**A**	#	**A**	#	**A**	#	**A**
MC1	museumplein	52	Ukraine	101	corona	43	Poetin	90	Amsterdam	43	Giro555	85
MC4	Malieveld	26	demonstratie	143	wappies	26	Rusland	116	Museumplein	22	Op1	100
MC7	Museumplein	19	Rusland	98	museumplein	15	EU	57	Malieveld	13	Ukraine	44
MC10	coronadebat	5	deochtend	28	lockdown	5	EU	18	KoffieDrinken	4	terzaketv	17
MC2	politiegeweld	2	Canada	11	2januariTwitterdemonstratie	2	Corona	11			Rusland	8
MC0	Trump	5	Oekraine	44	museumplein	4	NieuwsInPerspectie	31	Wijster	3	Rusland	26

TABLE 4.11: Top 6 Hashtags per Modularity Class (MC) divided over 2021 and 2022. Values that are left out have not been used more than two times.

4.26 visualizes the variation in activity before during and after the protest. The first snapshot represents the Twitter activity on March 18 and 19, the two days preceding the protest. As observed in this figure, there is minimal activity, with only Cluster 1 (Media) and Cluster 4 (Politics) being active during this time. However, these clusters do not specifically use hashtags related to the protest.

The second snapshot captures the Twitter activity on the day of the protest (March 20) until noon. Again, there is limited activity, although the hashtags "demonstratie" and "museumplein" appear for the first time within the analyzed period. Additionally, no changes are observed in the active modularity classes.

The third snapshot represents the remaining hours of the protest day (March 20, 12:01-23:59), during which a visible increase in activity is noted. Not only have the already active clusters used more hashtags, the other clusters started participating in protests related conversations.

Lastly, the fourth snapshot depicts the Twitter activity in the days following the protest (March 21-23). As shown in the figure, the activity continues to increase, particularly in modularity class 1.

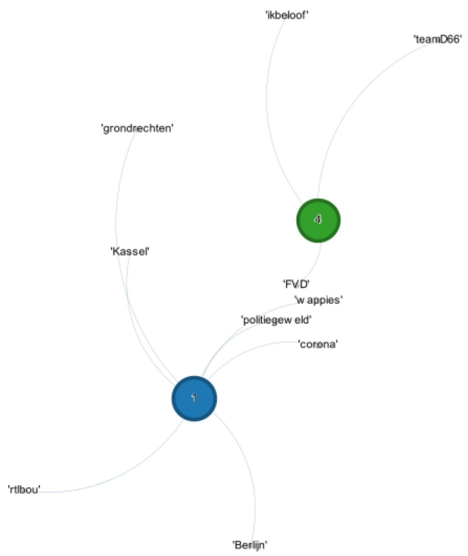

FIGURE 4.26: GIF visualizing the activity among modularity classes before - during - after the protest on 10 March 2021. This protest was located in Amsterdam. Four images, each showing different time intervals. (GIF only supported by certain PDF programs, adobe acrobat reader works)

4.6 PReSNA - A Systematic Method for Analyzing Protest-Related Social Networks on Twitter

The analysis of the original data set has provided valuable insights into protest-related networks. Building upon these findings, a systematic method is developed to analyze protest-related social networks comprehensively, both before and after the occurrence of the protest. The method is named PReSNA, which stands for **P**rotest-**Re**lated **S**ocial **N**etwork **A**nalysis, and it aims to facilitate the examination of protest-related social networks on Twitter.

PReSNA: A Systematic Method for Analyzing Protest-Related Social Networks on Twitter.

The PReSNA method consists of six distinct phases, which are enumerated below:

1. Scope Definition
2. Query Formulation
3. Tweet Retrieval
4. Data Transformation
5. Network Creation
6. Quantitative Analysis

4.6.1 1. Scope Definition

The first phase of the method involves defining the scope of the analysis. This phase consists of two key steps.

The first step is to identify the specific protest(s) or time frame that will be analyzed. This step establishes the boundaries and focus of the analysis, which can vary depending on whether the emphasis is on a particular protest, protests within a specific time frame, or protests in general. To gain insights into the protest, studying media sources can be an effective approach as they provide information about the related groups, topics, and goals of the protest. In addition to these protest-related attributes, it is crucial to consider the likelihood and impact of the protest. Protests with significant impact require more attention from the police, as mitigating their effects demands greater effort. However, it is also possible to analyze protests with lower impact but higher likelihood. To assess the likelihood and impact of the protest under analysis, the following question needs to be addressed.

1. What are the main motivations and demands of the protest?
2. Does the protest relate to a controversial or highly-discussed topic
3. What is the level of public support or opposition to the protest?
4. What is the potential for escalation or violence during the protest?
5. What is the historical context and precedent for similar protests?
6. What is the level of organization and coordination among protest organizers?
7. What is the potential influence of external actors or interest groups on the protest?
8. What is the expected media coverage and public attention for the protest?

9. What are the potential legal and regulatory implications of the protest?

The second step of this phase involves determining the specific aim of the analysis. There are several potential aims that can be identified, including:

- Analyzing the topics and discussions related to a specific protest.
- Investigating the Twitter activity leading up to a protest.
- Identifying and examining groups and communities within protest-related conversations.

4.6.2 2. Query Formulation

The second phase involves formulating the query that will be used to retrieve the data. The specific content of the query will depend on the chosen scope of the analysis. The effectiveness of the query can be enhanced by using keywords that are already identified as being used in mainstream and social media. It is therefore useful to extract these keywords from the media analyzed in Phase 1.

In this study, differences in hashtag usage per protest are identified, highlighting the importance of carefully selecting the keywords. For instance, when analyzing a specific protest, the query can be tailored to retrieve tweets specifically related to that protest, using specific keywords such as "Amsterdam2januari" and "ikdoenietmeermee". In addition to these specific keywords, it is also important to include generic keywords such as "demonstratie" and "protest" in the query. However, it should be noted that including generic keywords may result in retrieving data that is not solely related to the specific protest under analysis, as it may include tweets related to other protests during the same period.

While the inclusion of generic keywords may introduce bias when analyzing a specific protest, they are crucial when analyzing protests within a certain time frame or protests in general. These generic keywords can help identify differences in protest-related social networks and provide insights into the broader landscape of protests.

It is essential to carefully consider the selection of both generic and specific keywords during query formulation to ensure the relevance and accuracy of the retrieved data for the analysis. To provide further guidance in this phase, additional categories as defined by Müter et al. (2023) and discussed in the related work chapter (Chapter 2) are integrated into this method. The resulting list of categories, along with corresponding questions, are presented below.

- **Generic:** 'demonstration', 'protest', 'riots', 'demonstrators', 'unrest', 'social unrest', 'danger', 'occupation', 'on the streets', 'dissatisfaction', 'uprising', 'police', 'outbreak', 'hooligans', 'rioters', 'activists', 'arrested activists', 'demonstrat
- **Specific:** What specific words are used in relation to this protest?
- **Related Groups:** Which groups are associated with or connected to this protest?
- **Counter Groups:** Which groups oppose or are against the goal of the protest?
- **Related Topics:** Which topics are closely related to this protest?
- **Supporting Parties:** Which political parties are providing support for the protest?
- **Opposed Parties:** Which political parties are opposed to the protest?

4.6.3 3. Tweet Retrieval

When the query is formulated, the next step is to retrieve the tweets based on the formulated query. Apart from the query itself, other parameters such as language and date range need to be determined. The tweet retrieval process for this study is explained in the method section (Section 3).

In this particular study, a decision was made not to trigger the Twitter API during periods when protest-related Twitter activity was expected to be low. However, it is important to note that this approach may result in missing data, and should be avoided whenever possible, especially when analyzing specific time frames. While the same risk exists for protest-specific analyses, it is typically feasible to retrieve the required data in one or two attempts, depending on the time period that can be retrieved through the API. However, for larger time frames, more frequent API triggers and careful monitoring are necessary to ensure a complete data collection.

- What specific data do I want to retrieve? (e.g., user profiles, trending topics, hashtags, mentions)
- What time period do I want to cover? (e.g., past 7 days, past month, past year)
- What is the geographical scope of the data I want to retrieve? (e.g., worldwide, specific countries, specific cities)
- Do I want to filter the data based on specific keywords, hashtags, or users? If yes, which keywords, hashtags, or users are relevant to my analysis?
- Do I want to limit the data to specific languages? If yes, which languages are relevant to my analysis?
- Should I consider retweets or only original tweets?
- How much data do I want to retrieve? (e.g., maximum number of tweets per API call, total number of tweets to retrieve)
- Do I need to retrieve real-time data or historical data?
- Do I have access to the required API permissions and limitations?
- Should I use the streaming API to receive real-time data, or is it sufficient to use the RESTful API?
- Is personal data relevant for my analysis? (e.g., user, user name, location, user description)

4.6.4 4. Data Transformation

Once the API is successfully used to retrieve tweets posted within the defined period that match the query, the next phase is the data transformation phase. The specific transformation process will depend on the defined aim of the study, as different aims require different transformations. The primary objective of this phase is to establish connections between attributes, such as hashtags and/or mentions, within the retrieved tweets. The choice of connections to be established will be determined by the specific aim of the study.

Table 4.12 provides insights into which attributes should be used for specific goals within the study. This table serves as a guide to determine which attributes are most relevant for achieving each specific goal.

Analysis	**Goal**
Hashtag analysis	Analyzing the usage of hashtags to identify popular topics, trends, patterns, and relationships between topics.
Mention analysis	Examining the interactions between users by analyzing mentions to identify key actors, network structures, and protest-related online groups (modularity classes)
Hashtag-mention analysis	Studying the relationship between topics and groups by linking the hashtags to the groups that use it.

TABLE 4.12: Goals and analysis methods for Social Network Analysis of tweets

Before transforming the data, an important consideration is whether the created network should be directed or undirected. The choice between directed and undirected networks depends on the specific analysis being conducted.

If the analysis focuses on interactions between or within groups, it is recommended to create a directed network. In this case, the tweeter would be identified as the source node, and the mentioned users or groups would be the target nodes. This allows for the examination of the direction and strength of relationships between different entities in the network.

On the other hand, if the analysis aims to identify relationships between attributes, such as the co-occurrence of hashtags within the same tweet, the direction of the links is not relevant. In this scenario, an undirected network can be created, where the focus is on identifying connections or associations between attributes without considering their directionality.

The decision to use either a directed or undirected network should align with the research objectives and the specific research questions being investigated. By carefully considering the nature of the relationships and the goals of the analysis, researchers can determine the appropriate type of network to construct during the data transformation phase.

4.6.5 5. Network Creation

Once the data is transformed, the network can be created using visualization software such as Gephi, which is recommended for use with the PReSNA method. Gephi offers various functionalities and algorithms, including the ForceAtlas2 algorithm, which can be used to visualize and analyze the network.

To learn how to use Gephi and explore its features, you can refer to the official Gephi website (https://gephi.org/users/). The website provides comprehensive documentation and resources that can guide you in using the program effectively.

In the method, particular attention is given to categorization and filtering techniques, which can provide valuable insights into the protest-related data. Categorization allows for the application of colors within the network visualization, enhancing the amount of information that can be displayed. Different categorization techniques can be employed based on specific goals and objectives. Table 4.13 presents a list of categorization techniques along with their respective goals, providing a framework for organizing and interpreting the network data.

Goal	**Categorization**
Discovering of hidden patterns	By categorizing hashtags in topics and users in groups, insights into how different topics, groups, or users are connected can be gained.
Sentiment analysis	By categorizing hashtags or mentions as positive, negative, or neutral, you can gain insight into the prevailing sentiment around certain topics or groups.
Identifying influential users	By categorizing users based on their involvement with specific hashtags or groups, you can identify influential users. These users may have a large number of followers, generate significant interactions, or play a significant role in certain communities.

TABLE 4.13: Goals and categorization techniques for Social Network Analysis of tweets

4.6.6 6. Quantitative Analysis

The final phase of the method involves conducting quantitative analysis based on the insights gained from the constructed networks. The networks provide valuable information about relationships and frequently used hashtags or mentioned users. The choice of quantitative analysis techniques will depend on the specific objectives of the analysis.

One commonly used statistical test is the t-test, which can be employed to compare samples with the overall mean or with each other. The t-test is particularly useful for identifying significant differences between groups or samples. Other quantitative analyses that can be conducted may include correlation analysis, regression analysis, or network centrality measures, depending on the research questions and objectives.

Table 4.14 presents a list of potential quantitative analyses that can be performed within the PReSNA method. These analyses serve to deepen the understanding of the protest-related social networks and uncover patterns or relationships within the data.

Goal	**Analysis**
Network Centrality Analysis	Degree centrality, Betweenness centrality, and Eigenvector centrality can be used to identify the most important users or topics within a network.
Community Detection	Modularity classes can be used to analyze different groups of people or hashtag within the network and compare their attributes (tightness, in-degree, out-degree) with each other.
Sentiment Analysis	When using categorization based on sentiment, quantitative analysis can be used to compare the sentiment of groups or identify the overall sentiment that is related to a topic.
Temporal Analysis	By adding a time element to the network, all earlier mentioned analyses can be done to identify temporal patterns or differences.

TABLE 4.14: Quantitative analyses for Strengthening Network Analysis

4.7 Evaluation

Now that the development phase of the method has been completed, it is necessary to evaluate its effectiveness using new data. This evaluation is crucial for assessing the generalizability of the method and validating its results.

By analyzing this new data set using the developed method, a further evaluation of the effectiveness and applicability of the method in capturing and analyzing protest-related social networks can be done. The findings from this evaluation will provide valuable insights into the performance and utility of the method, enhancing its credibility and relevance for future research and practical applications.

In the first phase of the method, the scope and aim are defined. The scope for the evaluation it defined on the time frame of February 12, 2023, to March 10, 2023. This period precedes the events of March 11, when both the Extinction Rebellion and the farmers' protests took place. Additionally, the provincial elections held on March 15 are considered, as they are relevant to the context of these protests.

The aim of the evaluation is threefold. Firstly, it is to identify and analyze the relationships between topics within the protest-related social network captured during the defined time frame. This analysis aims to uncover the interconnectedness of different topics and their significance within the network. By examining the patterns of topic interactions, valuable insights can be gained into the structure and dynamics of the protest-related social network.

Secondly, the evaluation aims to identify and compare differences in group characteristics within the network. These group characteristics can include factors such as user demographics, engagement levels, or influential roles within the network. By comparing these characteristics to the findings from the exploratory phase of the study, a deeper understanding of the protest-related social network dynamics can be achieved.

Thirdly, the identified groups, along with the associated hashtags, are utilized to create networks that provide insights into the specific topics and discussions taking place within those groups. These networks serve as visual representations of the topical focus and interactions within each group. By examining these networks, a comprehensive understanding of the major themes and conversations happening within different segments of the protest-related social network can be obtained.

In the second phase of the method , the query is defined to retrieve relevant data for analysis. As stated in the PReSNA method description, the query categories that are defined in the paper by Müter et al., 2023 are used. The evaluation data set for this study is the same compared to the data set used in the mentioned study. The used query, along with the categories can therefore be found in Section 2. The paper provides more comprehensive insights into the query formulation process, including the specific keywords and criteria employed. It is important to note that this query formulation aligns with the principles and guidelines outlined in the PReSNA method, as this phase of the method is derived from the referenced study.

In the third phase of the method, the tweets are retrieved. Based on the aim and scope of the study, choices were made for the tweet retrieval. When considering the questions defined in the method section, the following choices were made. The aim of the study is to analyze the hashtags, as well as the mentions to get insights into topics and relations between groups. In addition, the differences over time will be analyzed. Therefore, the data that was necessary to fulfil this aim, is at least the

hashtags, mentions, and creation date of the tweet are retrieved. Because another goal is to identify whether the groups identified in the first phase of the study, are active within this context as well, the hashed usernames are retrieved, so that the same groups can be analyzed. The time period that is covered in this evaluation is twenty-six days, which was the maximum of days of which data could still be retrieved when the decision was made to analyze these protests. The geographical scope of the data is in line with the scope of this study, namely protests in the Netherlands, which is also the argument that only Dutch tweets are retrieved. The keywords that are used for the tweet retrieval are already defined in the second phase and are based on terms used in the mainstream media.
As retweets can be seen as a confirmation of opinion, these are integrated in the retrieval.
For the amount of data that should be retrieved, the decision is made to retrieve as much as possible, as this enable to draw conclusions on a larger data set. Moreover, the time frame of the scope is relatively short, which therefore will not cause in an overload of data.
The analysis is done after all tweets are retrieved and the choice is therefore made to do not use real-time data which also makes it possible to use the RESTful Twitter API.

During the fourth phase of the method, the data is transformed to facilitate the analysis of relationships between topics, groups, and topics within groups. The transformation process differs depending on the specific analysis being conducted.
For the analysis of hashtags, undirected networks are employed since the focus is on identifying relationships between hashtags, irrespective of their direction. Similarly, for the analysis of mentions, where users mentioned in the same tweet are linked together, the direction of the relations is not relevant.
However, in the analysis of topics within groups, the direction of the relationships becomes important. In this case, the groups are considered as the source and the hashtags as the targets. This approach allows for the examination of topics (hashtags) that generate activity within specific groups, providing insights into the dynamics and interests of these groups.

The upcoming sections provide an overview of the decisions made during the creation of the networks, as well as the outcomes of the quantitative analysis, which are crucial for phase five and six of the PReSNA method.

4.7.1 Hashtags

To identify the primary topics discussed on Twitter, an analysis of the hashtags in this new data set is conducted. A total of 21,494 hashtags are retrieved from tweets that contain more than one hashtag. This methodology ensures the establishment of relationships between hashtags, forming a social network. Due to the large number of retrieved hashtags, filtering techniques are employed to examine the most frequently used ones. The objective of analyzing this new data set is to evaluate the developed method, after which the results can be compared to the results of the original data set. To maintain impartiality, similar settings and filtering approaches are applied as for the original data set, to facilitate an optimal comparison.

During the initial data analysis of this study, a Weighted Degree (WD) filter of >300 was utilized, and this same filter is also applied to the new data set. By implementing this filter, a network consisting of 233 hashtags with 9,385 distinct edges is retained. Similar to the previous method, the remaining hashtags are categorized within this network. Figure 4.27 depicts the categorization along with the corresponding percentages, illustrating the distribution of each category within the network. Note that the categories remain consistent with the previous categorization, with the exception of Eurovision 2022 and Energy, which have been replaced by Opinions and Elections, reflecting their greater presence in this data set. Additionally, the category Politics has been modified to Political topics/ politics related to encompass a wider range of hashtags relevant to this theme.

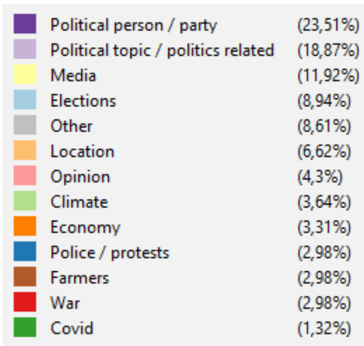

FIGURE 4.27: Defined categories for the evaluation data set and their assigned color based on predefined color values by ColorBrewer (2021).

As previously mentioned, in the evaluation data set, 233 hashtags remaining when applying the same weighted degree filter as used for the original data set. For the original data set, 284 hashtags were remaining with this filter. To test whether the average weighted degrees for both remaining networks are significantly different. An ANOVA test is used to compare the population in the original data set (μ 1,431), with the evaluation data set (μ 1,364). The outcomes of the ANOVA test shows a p-value of 0.769, which is higher than the significance level of 0.05 and results the conclusion that both networks do not significantly differ from each other in weighted degrees. This can be defined as remarkable, since the first data set consists of a time frame of 2 years, while the evaluation data set only consist of data from one month. Given this difference in time frames, the activity within evaluation data set is significantly higher per day, compared to the original data set.

To conduct a more in-depth analysis of the remaining network after filtering (WD > 300), the algorithm was ran incorporating only the relationships between the remaining hashtags. By focusing solely on the relationships between the most used hashtags, it becomes possible to gain a better understanding of the connections among these hashtags. The ForceAtlas 2 algorithm was found to be particularly effective in positioning the hashtags when the network was initially scaled to accommodate the large space. Once the positions were determined, the scaling was adjusted accordingly. However, in the case of this network, a majority of the hashtags exhibited strong relationships with other hashtags, necessitating the retention of a scaling factor of 300 to visualize the differences in positions.

Figure A.16 illustrates the social network of the remaining nodes after filtering, with the algorithm being re-executed. Notably, the categories are now more prominently clustered together. Particularly striking is the central position occupied by the "Elections" category within the network. Additionally, the positioning of media-related hashtags differs from that in the original data set, where they were situated in the middle of the network.

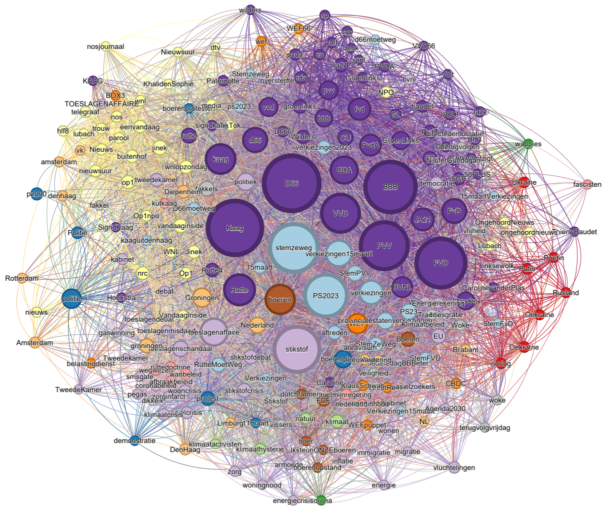

FIGURE 4.28: Social network of hashtags used between February 12 and March 10 2023. Position based on only the edges that are visible. Filter(s): $WD \geq 300$. Remaining nodes/edges: 233N(1.1%). Network scaled X300. Coloring used as defined in 4.27

Now that the main topics have been analyzed, it is important to examine the strongest relationships within the network. Figure 4.29 illustrates the remaining hashtags when only the most significant relationships are highlighted. This network was created using an Edge Weight (EW) filter of >200, focusing on the most robust connections. It is worth noting that, unlike the previous network (Figure A.16), the positions of the hashtags have not been altered as the algorithm was not re-run. The only change made was the application of the filtering criteria.

The majority of these relationships are observed between political parties or their members, reflecting the close associations within the political landscape. Additionally, certain political topics such as "stikstof" (nitrogen), "toeslagenaffaire" (childcare allowance affair), and "toeslagenschandaal" (childcare allowance scandal) emerge as significant themes that have generated considerable activity during the analyzed period. It is worth noting that the hashtag referring to the current Prime Minister, Mark Rutte, is positioned closest to these issues compared to other hashtags within the Political person/party category.

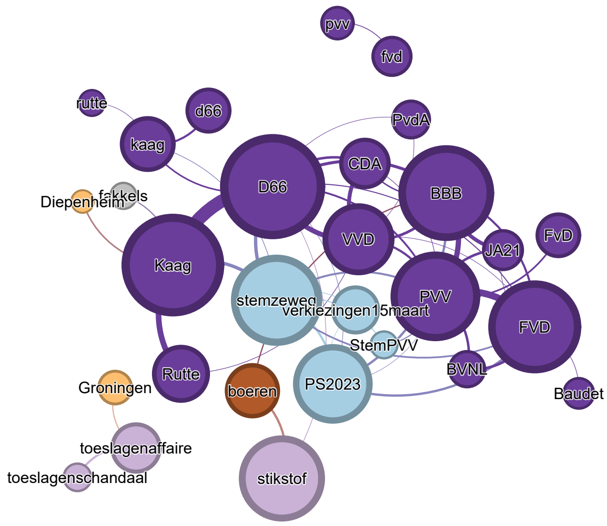

FIGURE 4.29: Social network illustrating the strongest relationships between hashtags used between February 12 and March 10 2023. Filter(s): $EW \geq 200$. Remaining nodes/edges: 29N - 52E. Network scaled X300. Coloring used as defined in 4.27

In the original data set, the category Location showed a significantly association with the category Police / Protests. However, in this network, the last mentioned category is not present when filtering for the strongest relations. Nevertheless, the category Location is represented by two nodes. To evaluate possible differences in the usage of locations in conjunction with other categories, the same ANOVA test was conducted. The results of this test are presented in Table 4.15.

The p-value obtained from the ANOVA test is 0.34, which exceeds the significance level of 0.05, leading to the acceptance of the null hypothesis (H0). Therefore, within the evaluation data set, there is no significant association between the category 'Location' and any specific other category.

Category	Political top/rel	Media	Elections	Other	Economy	Opinion	Police/Protests	Farmers	War	Climate	Covid
μ	9.96	6.03	7.30	12.28	5.38	10.21	8.04	7.39	2.19	3.91	1.25
Polit. pers/part	0.81	3.12	1.85	3.12	3.77	1.06	1.11	1.77	6.96	5.24	7.9
Polit. top/rel	0	3.92	2.66	2.32	4.58	0.25	1.92	2.57	7.77	6.05	8.71
Media	3.92	0	1.26	6.24	0.65	4.18	2	1.35	3.84	2.12	4.78
Elections	2.66	1.26	0	4.98	1.92	2.91	0.74	0.089	5.11	3.39	6.05
Other	2.32	6.24	4.98	0	6.9	2.06	4.24	4.89	10.08	8.37	11.03
Economy	4.58	0.65	1.92	6.9	0	4.83	2.66	2.01	3.19	1.47	4.13
Opinion	0.25	4.18	2.91	2.06	4.83	0	2.17	2.83	8.02	6.3	8.96
Police/Protests	1.92	2	0.74	4.24	2.66	2.17	0	0.65	5.85	4.13	6.79
Farmers	2.57	1.35	0.089	4.89	2.01	2.83	0.65	0	5.19	3.48	6.13
War	7.77	3.84	5.11	10.08	3.19	8.02	5.85	5.19	0	1.72	0.94
Climate	6.05	2.12	3.39	8.37	1.47	6.3	4.13	3.48	1.72	0	2.66

TABLE 4.15: Absolute difference between means of edge weights with category location. No significant differences between the categories identified with the significance level of 0.05.

Having analyzed this network of hashtags, it is now pertinent to incorporate the element of time to gain insights in the evolution of this network and evaluate how these tight relationships emerged. Before delving into the dynamics of the network, further insights are provided into the Twitter activity on a daily basis.

Figure 4.30 depicts the daily count of relationships between hashtags during the specified time period. It is crucial to note that a tweet containing multiple hashtags results in multiple relationships within the network. For instance, a tweet with four hashtags leads to six combinations of relationships. Furthermore, it is worth mentioning that tweet retrieval ceased on March 10 at 11:50 AM, which explains the lower value observed on the last day of data collection. Upon analyzing the line diagram, three distinct peaks can be observed, corresponding to the dates 20-02, 24-02, and 07-03. Additionally, a decline in hashtag use is noticeable towards the end of the depicted time period. These variations in hashtag usage necessitate further analysis, which is elaborated on in the following paragraph.

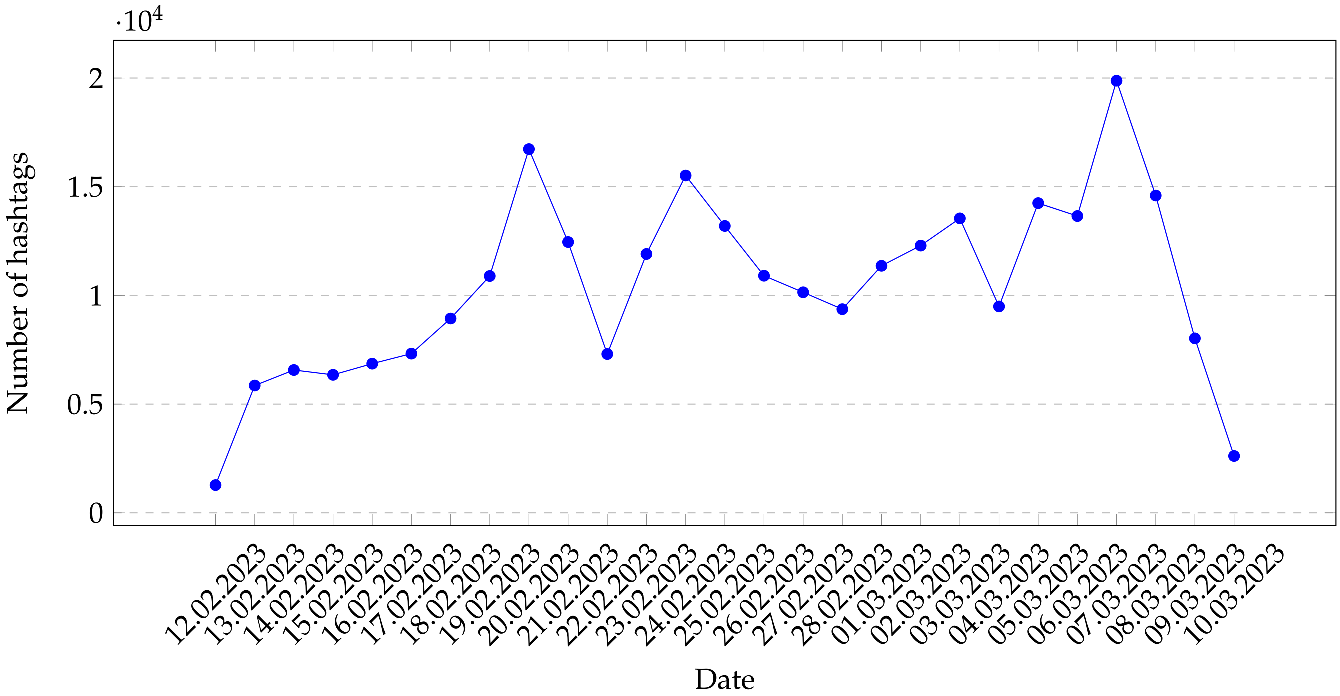

FIGURE 4.30: Daily count of retrieved hashtag relationships.

Figure 4.30 provides insights into the fluctuating use of hashtags over time. To further understand the topics that drove the increased activities on specif days, a dynamic social network is constructed using the same hashtags as in the network without the time element (Figure A.16). Given that this data set spans approximately a month, it is possible to analyze this network on a daily basis.

The edges associated with the remaining hashtags are distributed across the days, with each day representing the day of tweet creation and the subsequent twenty-four hours. This approach enabled the calculation of the Weighted Degree (WD) for each hashtag per day. Subsequently, the size of the nodes is determined based on the WD, with a minimum size of thirty and a maximum size of 200. This resulted in a dynamic network in which the most frequently used hashtags per day are visualized as the largest nodes. The outcome is presented in Figure 4.31, which is a dynamic figure (GIF) that is only supported by certain PDF readers, such as Adobe Acrobat.

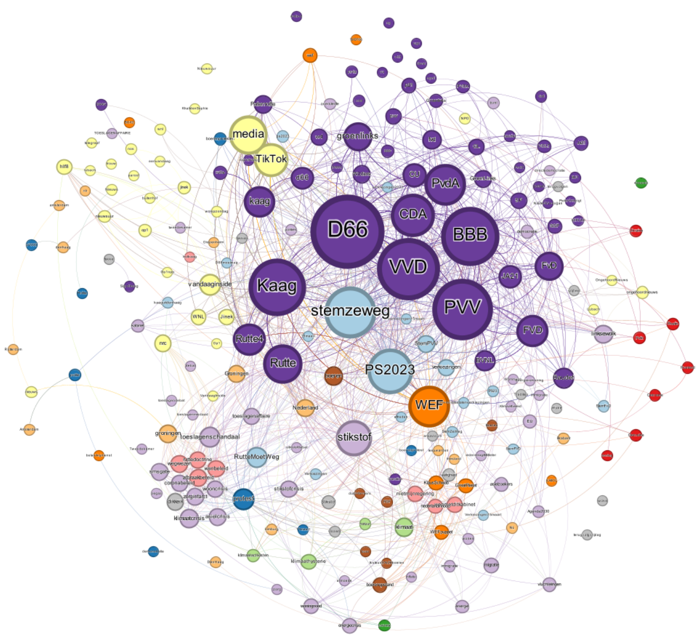

Feb 14

FIGURE 4.31: GIF visualizing the hashtag activity during the period of February 12 and March 10. Coloring used as defined in Figure 4.27. (GIF only supported by certain PDF readers, Adobe Acrobat Reader works)

The three peaks that were identified earlier on the days Feb 20, Feb 24, and Mar 7, can be better explained by this dynamic network.
it is visible that the hashtag "Fakkels" is prominently represented in the network on Feb 20. On the day before (19 Feb), Minister Kaag was awaited by a group of protesters carrying torches in a small village called Diepenheim, which is also mentioned in the network. This event garnered significant attention on Twitter based on this representation.

Furthermore, on Feb 24 the overall activity in the network increased. In addition to the frequently used hashtags in the center of the network throughout the entire time frame, there was surge in activity within the yellow cluster representing the media on the left side, as well as among the political-related topics on the bottom left. This increased activity can be attributed to the presentation of the report "Groningers above gas", which questioned whether Prime Minister Rutte had done enough for the citizens of Groningen (NOS, 2023). Moreover, the preceding day witnessed a parliamentary debate on the issue of nitrogen (stikstof).

On March 7, there was a notable increase in the use of hashtags related to political parties, election related topics, and political issues in general. The childcare allowance affair emerged as the most discussed political topics, followed by the nitrogen crisis. The increased activity on this day can be attributed to a parliamentary debate about the Tax and Customs Administration's handling of the childcare allowance affair.

Among the three events, the incident in Diepenheim is particularly noteworthy due to its nature as an unannounced protest. Further analysis of how this event generated increased activity is warranted. Figure 4.32 depicts the activity per two-hour intervals on February 19, 20, and 21.

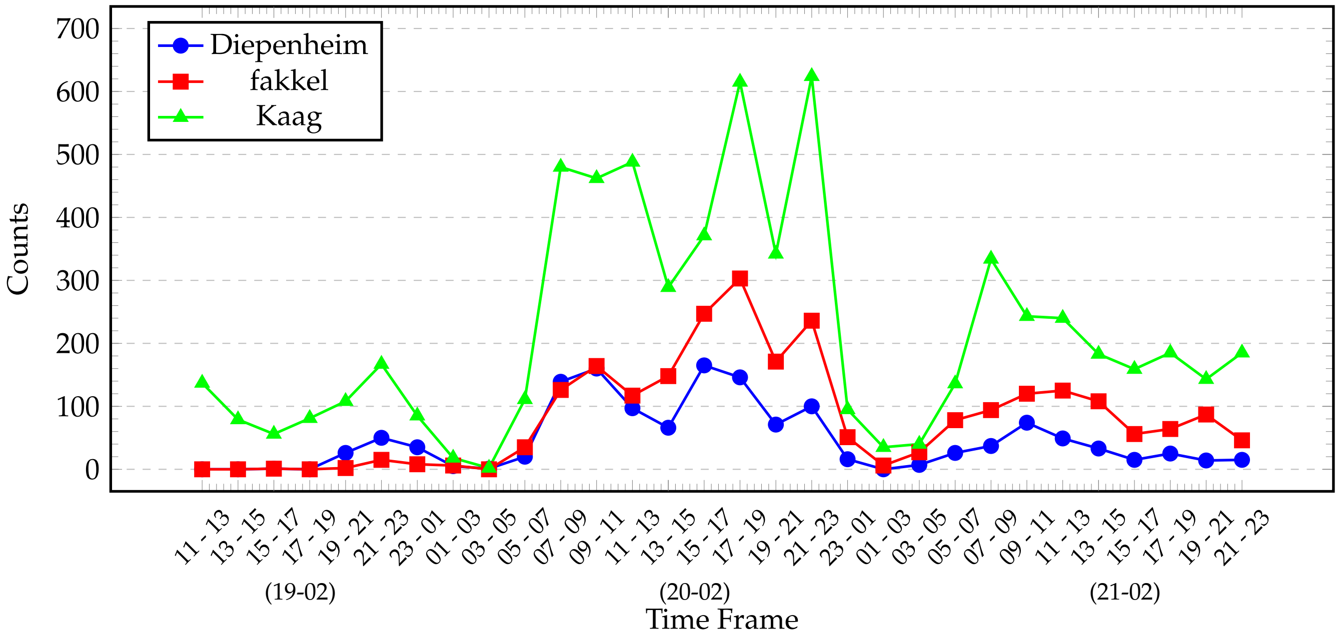

FIGURE 4.32: Counts of relationships with the hashtags: "Diepenheim", "Fakkel", or "Kaag" within time frame 19-02 - 21-02.

While the event occurred on February 19, the significant surge in activity occurred predominantly on the following day. Similar patterns can be observed on this day for the three series, where activity rapidly rises in the morning hours (05-11), followed by a slight decline until 15:00, after which activity starts increasing again until 19:00, when a short dip occurs across all three time frames (19-21), followed by one more increase before nightfall. On February 21, activity related to this topic continues to rise in the morning, after which it starts declining throughout the day.

To gain further insights into the reasons behind the highest activity levels observed on the day after the event, a dynamic network is explained in section (4.7.3) to illustrate the usage of these hashtags by different modularity classes.

Before adding modularity classes into the networks, a more in depth analysis is done on the hashtags related to the two protests within the context of this data set.

With this evaluation data set, the focus is on the announced protests of Extinction Rebellion and the farmers. In the dynamic network (Figure 4.31), it is evident that on February the 23rd, there was an increase in activity among the hashtags related to nitrogen (stikstof). This significant increase is depicted in Figure 4.33 and can be attributed to the debate on this topic in the chamber of representatives on that day. As shown in the graph, the hashtag "stikstof" was not used more frequently in the period leading up to the protest compared to that particular day. Further analysis focuses on the hashtags "protest" and "boerenprotest", as they directly refer to protests, whereas the hashtags "stikstof" and "boeren" can be used in broader contexts.

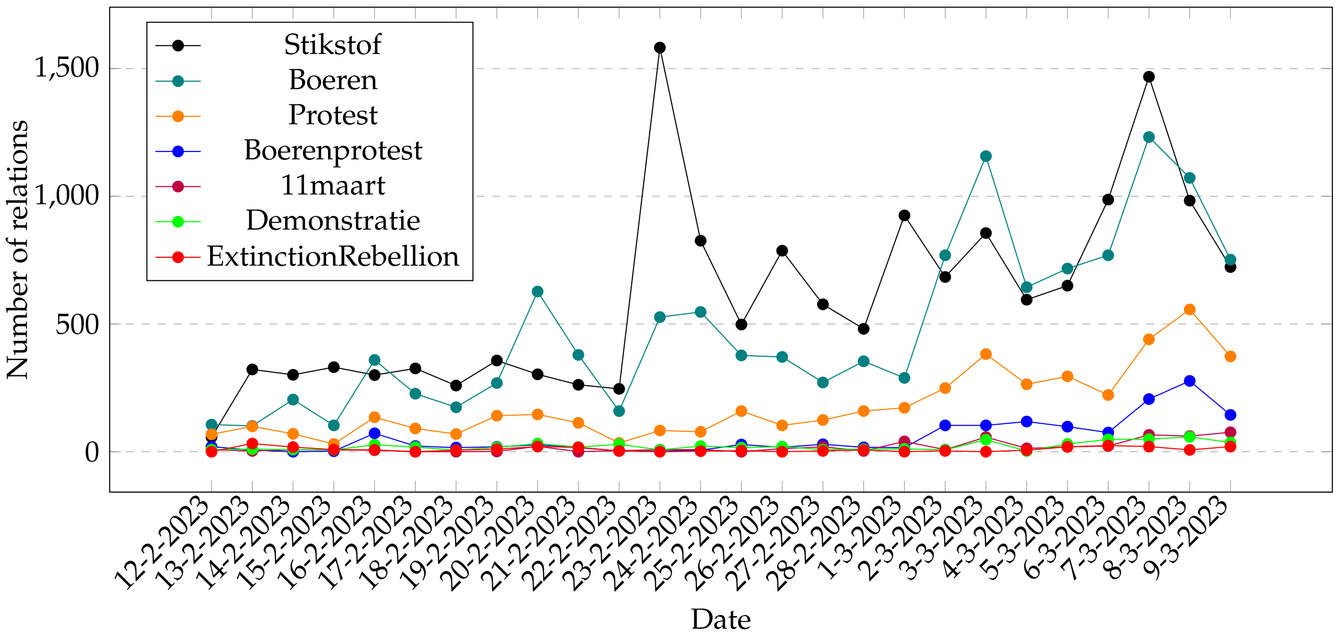

FIGURE 4.33: Frequency of hashtag relationships from 12-2-2023 until 9-3-2023: Pre-Protest Trends.

To detect rising activity prior to the actual peak, which can be determined afterward, a two-sample t-test is conducted. The tests compare samples of seven days with the seven preceding days. By shifting these comparisons one day at a time, the days on which the first significant differences are present can be identified. For the hashtag 'protest', a significant difference is observed when comparing the activity between the samples from 18/02 to 24/02 and 25/12 to 03/01. The result of the two-sample t-test is p-value $< .05$, indicating a statistically significant increase in activity between these weeks. For the hashtag "boerenprotest", this significant difference can first be identified one day later, on March twenty-sixth. On March twenty-seventh, the outcome of the tests for both analyzed hashtags shows a significant increase in usage with a p-value of $< .01$.

In addition to identifying the first significant increase in activity among these hashtags, it is also pertinent to determine the period during which the activity reaches its peak. This information provides valuable insights into online activity preceding protests, which can be applied to future events. The largest difference in activity for the hashtag "protest" when using one-week samples, is observed between 02/03 to 08/03 and 23/02 to 01/03. The difference in averages between these samples amounts to 218.71, representing a substantial increase of 174.4%. Similarly, for the hashtag "boerenprotest", the maximum difference in averages is found within the same date range. Notably, the maximum difference for this hashtag is 123.57, signifying a notable increase of 752.1%. It is worth noting that the percentage increase for the "boerenprotest" hashtag is particularly significant as it predominantly starts gaining widespread usage shortly before the actual protest.

While the hashtag "protest" is used significantly, the hashtag "demonstratie" is used relatively little. The difference in usage between these two hashtags is significant (*p-value* $< .01$). The first mentioned hashtag is used 179 times on average within this data set, whereas the second mentioned hashtag is only used 21 times on average per day. This finding is notable considering that in the original data set, the hashtag "demonstratie' had higher usage compared to the hashtag "protest". It suggest that the choice between the two hashtags is protest dependent.

4.7.2 Mentions

Now that the hashtag networks have been computed and analyzed, the modularity classes are included. To evaluate the outcomes of the previous data set, the same modularity classes are used. To reach this goal, the usernames of individuals categorized in the modularity classes of the previous data set are used to filter mention data in the current data set. This approach enables a comparison of class activity between the two data sets. Figure A.17 illustrated the social network of users in the evaluation data set, with coloring based on the classification defined in section 4.2.2 to indicate users identified in the previous set.

Analyzing this network, notable differences can be observed when comparing this network to the network that was made for the original data set. The original data set exhibited distinct clusters that could be identified as separate classes, such as classes 0, 14, and 16, while the larger classes 1 and 4 were more distributed throughout the network. In this new network, the presence of numerous grey nodes represents users not previously incorporated in the previous data set. Among the colored nodes, classes 14 and 16 are not as easily discernible as in the previous set. However, modularity class 0 remains clearly identifiable, with the larger node sizes denoting the increased frequency of user mentions. Similarly, the largest classes 1 and 4 are spread throughout the network, with some users within these classes receiving frequent mentions, as evidenced by their larger node sizes.

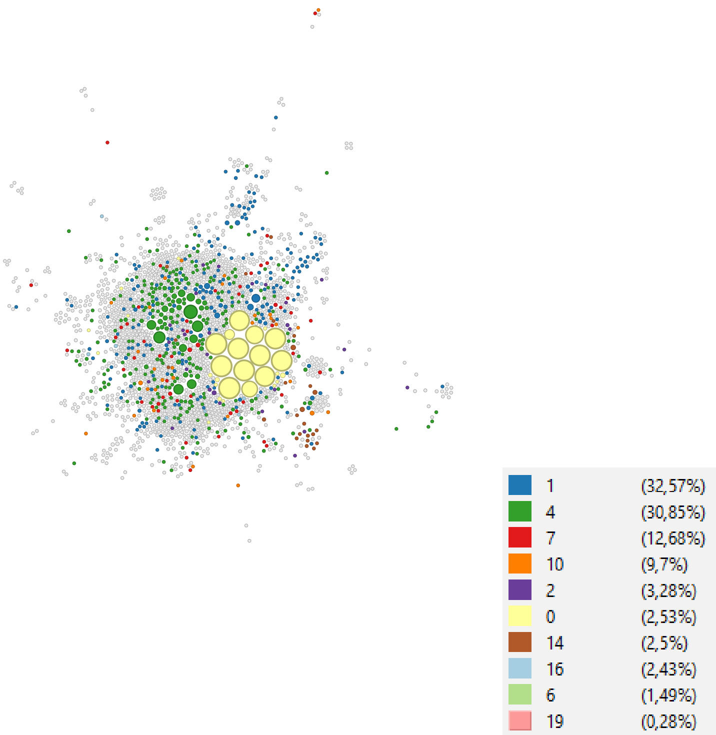

FIGURE 4.34: Social network with the users of the earlier identified classes in which the position is based on the mentions in the evaluation data set. Filter(s): $EW \geq 2$. Legend visible on the right.

Mod. Class	User 1				User 2				User 3			
	U	M	D	μEW	U	M	D	μEW	U	M	D	μEW
MC1 (1289)	Politie	1663	452	3.7	Het_OM	992	209	4.7	NOS	840	374	2.2
MC4 (1221)	D66	3416	776	4.4	lientje1967	2758	**878**	3.1	SigridKaag	2525	670	3.8
MC7 (502)	************	391	69	5.7	*********	362	52	10.0	**********	208	105	2.0
MC10 (384)	****************	380	77	4.9	**********	364	119	3.1	****************	146	87	1.7
MC2 (130)	*************	413	189	2.2	************	152	103	1.5	*********	145	69	2.1
MC0 (100)	************	**5939**	98	60.6	************	5924	72	82.3	**********	5834	83	70.3
MC14 (99)	***********	585	122	4.8	****************	507	92	5.5	**********	429	56	7.7
MC16 (96)	**********	144	82	1.8	*******	94	69	1.4	**************	76	13	5.8
Mod. Class	User 4				User 5							
	U	M	D	μEW	U	M	D	μEW				
MC1 (1289)	telegraaf	759	352	2.2	NLRebellion	557	206	2.7				
MC4 (1221)	thierrybaudet	2200	728	3.0	fvdemocratie	1963	700	2.8				
MC7 (502)	*********	207	65	3.2	***************	174	125	1.4				
MC10 (384)	************	134	64	2.1	*********	111	24	4.6				
MC2 (130)	***********	86	34	2.5	********	67	27	2.5				
MC0 (100)	**************	5760	55	**104.7**	*************	5757	62	92.9				
MC14 (99)	*************	185	54	3.4	************	110	57	1.9				
MC16 (96)	**********	49	21	2.3	**********	47	38	1.2				

TABLE 4.16: Top 5 most mentioned users per Modularity Class (MC). U representing the username (* used to anonymized accounts that are not publicly linked to media or politics). M representing the amount of times a person is mentioned, D the amount of different persons an individual is linked to, and μEW the average Edge Weight (amount of times a person is mentioned with the same other person) which can be retrieved by dividing M by D.

Class	MC4	MC7	MC10	MC2	MC0	MC14	MC16
MC1	0.041	0.13	0.34	0.054	**27.27**	0.12	0.27
MC4	0	0.17	0.38	0.095	**27.23**	0.079	0.31
MC7	0.17	0	0.2	0.079	**27.41**	0.25	0.14
MC10	0.38	0.2	0	0.28	**27.61**	0.45	0.067
MC2	0.095	0.079	0.28	0	**27.33**	0.17	0.21
MC0	**27.23**	**27.41**	**27.61**	**27.33**	0	**27.15**	**27.54**
MC14	0.079	0.25	0.45	0.17	**27.15**	0	0.39

TABLE 4.17: Absolute differences between the means of average edge weights for different classes in the social network of the March 11 data set. Values in bold indicate significant differences at a significance level of 0.01.

4.7.3 Hashtags and mentions combined

In line with the method used for the original data set, an analysis done on the tweet and hashtag usage per classes. 4.35 depicts the average amount of tweets per user per class and the average amount of hashtags per hundred tweets. Analyzing the results, compared to the averages, no class differs largely for the tweets per user. For the hashtag usage however, MC4 and MC6 stand out in hashtag usage with values above forty.

Comparing these results to the results of original data set (4.25), similarities and differences can be identified. Similar to the original data data set, no patterns are visible when comparing the results to the size of the class. Differences are visible when comparing the overages and individual amounts per class.

When comparing the averages, the average amounts of tweets per user is lower for this data set, this however can be explained by the shorter time span. For the hashtags per hundred tweets, it is visible that the minimum, average, and maximum of this set are: 22.8, 34.8, 46.2. This is higher compared to the original data set (12.8, 24.5, 34.9). Comparing the individual classes for both sets, it is visible that MC4 is using the most hashtags per tweets in both sets, while the tweets per user is relatively low for this class. MC0 has the highest tweets per user ratio in both classes, meaning that this class can be defined as the most active class.

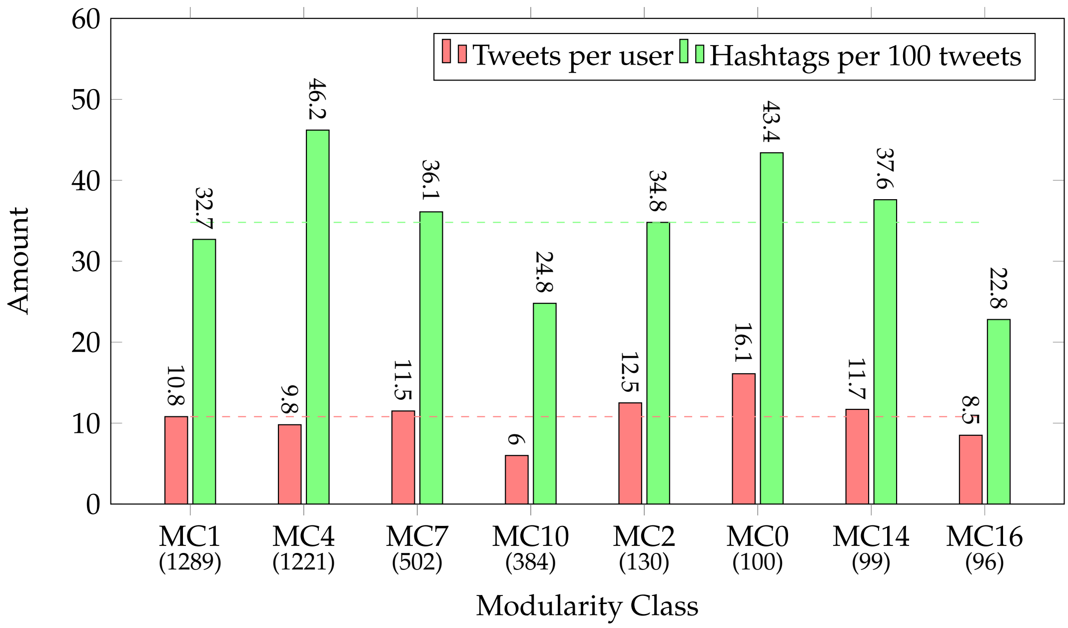

FIGURE 4.35: Frequency of hashtag usage per modularity class based on the evaluation data set. Classes divided along x-axis based on size where MC1 is the largest class, MC4 the second largest class, etc.

To analyze the hashtag usage, a social network was computed and is depicted in Figure A.18. This network provides insights into the distribution of hashtags across different modularity classes. The computation of the network involved applying an edge weight filter of > 2 to exclude typos and singularly used hashtags. Notably, each class possesses its own set of hashtags that are exclusively used within that class. Among the top ten most frequently used hashtags (whose names are added in figure), it is worth mentioning that one hashtag, ("NieuwsInPerspectief") is specific to class 0. This finding corroborates the results obtained from the original data set, indicating that this particular class consistently employs a unique hashtag with a comparable frequency to hashtags used by all the classes combined.

Another noteworthy observation, consistent with the findings from the original data set, is that the remaining hashtags from the top ten are positioned in the center of the network, indicating their link to multiple classes.

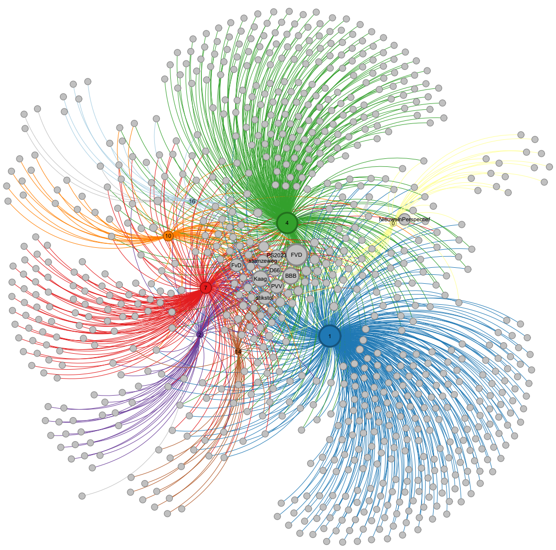

FIGURE 4.36: Social network of hashtags used by the defined modularity classes between 12 Feb 2023 and 10 Mar 2023. Filter(s): $EW \geq 2$. Remaining nodes/edges: 814N(25.44%) - 1,249E(28.76%). Edges rescaled to min 0.1 and max 1 for visibility purposes. Network scaled X1000.

Additional to the global overview of hashtag usage per class, a dynamic social network of the hashtags related to the event on February 19 was computed. By filtering out all other hashtags, the resulting network focuses on interactions involving the three hashtags defined in section 4.7.1. The network is depicted in Figure 4.37. The Force Atlas 2 algorithm was rerun based on the edges within the time span of February 19 at 11:00 to February 21 at 23:00.

The positioning of MC1 and MC4 in the center of the network indicates that these classes have the most interactions with all hashtags, whereas other classes are situated more towards the top right of the network, suggesting fewer interactions with the hashtags on the left. MC0 occupies a distinct position at the bottom of the network, indicating limited or infrequent interaction with the hashtags at the top. Each displayed edge illustrates one or more interactions between the modularity class and the hashtag, allowing for analysis of how the peaks in the previously created line diagram (fig 4.32) originated.

The event took place on February 19 in the afternoon, when protesters awaited Minister Kaag with torches. The network reveals activity primarily among classes MC1, MC4, and MC7, all of which make use of the hashtag "Kaag". The hashtag "Diepenheim" is first used by MC1 at 15:51, which was previously identified as the media class. MC1 employs this hashtag once more at 19:03, before MC4, identified as political persons, starts using it for the first time at 19:52. Prior to any other class using this hashtag, MC1 and MC4 both utilize it three times. Between 22:00 and 22:30, class 10 and class 7 both use the hashtag "Diepenheim" for the first time.

In the early morning of February 20, all classes are actively using the hashtags shown in the network. Each class either uses the hashtag "kaag" or "Kaag". Additionally, between 6:00 and 8:00 the hashtags fakkel (torch) and fakkels (torches) are first used by classes MC7 (6:38) and MC1 (8:05). From 8:28 until 11:28, these hashtags are used eleven more times by various other classes. Activity continues to increase throughout the day and reaches its peak between 17:00 and 21:00, during which MC1 uses the hashtags together 38 times, followed by MC4 with 32 mentions.

In the early morning of February 21 (5:00-9:00), clusters 1, 4, and 7 are mainly active, and later in the day, the other classes also become active later.

To summarize this event, it can be said that classes MC1, MC4, and MC7 show the most activity with the defined hashtags within this time frame. MC1 and MC4 are positioned in the center of the network, indicating frequent interactions with all hashtags. These classes are therefore likely to be more involved in the event. MC0 has a unique position, which suggests that this class has limited interactions with the hashtags used by other classes. The overall activity reaches its peak between 17:00 and 21:00 on February 20, with MC1 using the hashtags most frequently (38 times), followed by MC4 (32 times). This suggests that these classes are highly engaged during this time period.

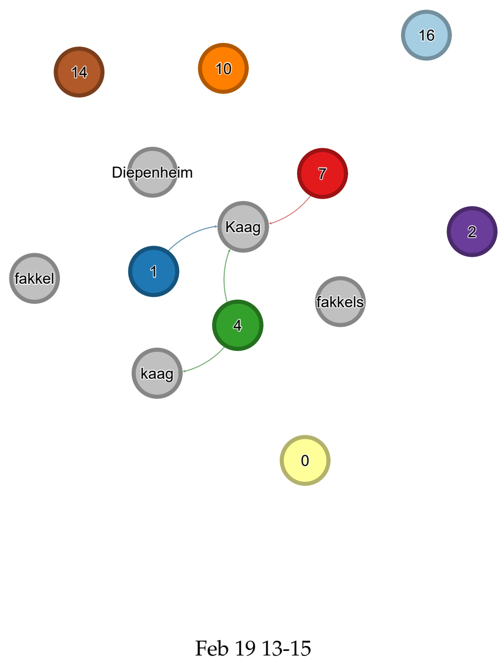

FIGURE 4.37: GIF visualizing the hashtag activity of hashtags "Diepenheim", "Kaag", "Fakkel" during the period of February 19 and February 21. (GIF only supported by certain PDF readers, Adobe Acrobat Reader works)

Chapter 5

Discussion

This study introduces a novel approach to identify protest motivation through the application of Social Network Analysis. The developed method was evaluated by applying it to a new data set, thereby assessing it robustness across different time frames. Notably. the most used topics within the tweets of the new data set differed from those observed in the original data set, making it an invaluable resource for examining the generalizability of findings. However, it is important to acknowledge the limitations associated with this study, which are addressed in this section. The key points for discussion are elaborated upon, and recommendations for future research are provided.

5.1 Limitations

5.1.1 Twitter as the only data source

Although the data sets used in this study consisted of significant amounts of data for analysis and were retrieved over a large time span, it is important to note that the data was sourced solely from Twitter. While previous studies have demonstrated the utility of Twitter as a valuable source for behavioral science, it is also recognized that protest-related conversations take place on other platforms as well. Although Twitter provides researchers with publicly available tweets and facilitates the retrieval of large data sets with relative ease, this very public nature of tweets also poses certain limitations.

It is worth mentioning that the number of protests has been increasing in recent years due to societal unrest. Protesters are increasingly engaging in actions that are not tolerated in the Netherlands, such as highway blockades. Since protesters are aware that such actions are not permitted, it is possible, and yet proven, that more protests are being organized without prior announcement, particularly during the Covid-19 period when regular protests were not permitted at all.

Unannounced protests are likely to be organized through private channels, which result in limited media attention and limited Twitter activity. Consequently. predicting these protests using public data becomes challenging. However, it is important to acknowledge that Twitter conversations related to these events often occur after the fact, enabling their analysis, as demonstrated in this study.

Additionally to the limitation of Twitter being the only data source, the retrieved data sets consisted of gaps within certain time periods. Although these gaps were intentionally created, by only using the Twitter API when protest-related conversations were relevant, this also poses a limitation. The used data for this study could therefore be biased and consists of more activity compared to a data set that also consist of tweets sent within periods without large protests. This bias makes it harder to find statistical differences between months, as the average might not be the exact reflection of the entire time period.

5.1.2 Differences in data set queries

For the exploration of patterns in protest-related social networks, a large data set is used covering two years of protests. This data set enabled to find patterns in hashtag usage and group dynamics. To evaluate the method's effectiveness, usefulness, and usability, as suggested by Peffers et al., 2007, a new data set was used. This evaluation data set consisted of data preceding two major protests and coinciding with upcoming elections. However, it is important to note that this data set was retrieved using a query of more than seventy words, whereas the original data set was obtained using a single keyword, "demonstratie" (protest).

This difference in query might explain why the measured activity per day within the evaluation data set is higher compared to the original data set, as a larger query results in more retrieved tweets per day. Additionally, the statistical tests to examine the relations of the category Location were also different. This discrepancy may be attributed to the difference in queries, combined with the different time frames. It is important to note that not only is there a difference in the length of the time frames, but the evaluation data set only includes data prior to the protests, while the original data set consists of data prior, during, and after the protests. This distinction is important because locations may be predominantly linked to protests when they are occurring or have concluded.
Because it can not be substantiated what directly caused these differences, it is necessary in further research to keep the queries equal for better evaluation results.
Although this evaluation data set consisted of data that could be used to predict protests, it is subjected to bias compared to the original data set. Despite this discrepancy, it offers valuable insights in the use of specific queries to retrieve more precise information about a protest, as queries can be tailored to gather tweets relevant to a given protest.

5.1.3 Bias in relationships data

In addition to the bias in the evaluation data set, there is a potential bias in the method used to establish relationships between users within the protest-related social network. For both the establishment of relationships between hashtags and relationships between users, the same method was used. For the hashtags, the most appropriate method to identify relations was to establish a connection when two hashtags were used in the same tweet. However, the same construction was applied to mentions, resulting in two users being linked in the social network if they were mentioned in the same tweet. This method leads to connected users based on tweets from other users, meaning that the person who sent the tweet and the user mentioned in the tweet do not have a direct connection. These users are only considered connected in the network if another user mentions both of them in the same tweet.

This method was chosen to generate more data, as a tweet with five mentioned users results in ten connections, compared to five connections when the mentioned users are only linked to the tweeter. However, it can be debated whether users who do not directly mention each other should be connected in a social network.

5.1.4 Categorization technique

In this study, the categorization of topics was performed based on the interpretation of the researchers without a specific categorization technique or a strong rationale. As a result, replicating the study with the same categorization as a result may be challenging. Additionally, by choosing to categorize based on topics, the opportunity to incorporate the underlying opinions expressed through hashtags was missed. This could have been a valuable addition to the study, as it would enable sentiment analysis.

However, it is worth noting that the PReSNA method, addresses this limitation by emphasizing the potential for integrating the sentiment analysis of hashtags. This ensures that future researchers are aware of this aspect and can incorporate it into their analyses.

5.1.5 The effect of the war in Ukraine

As depicted in the constructed networks using the original data set, the category labeled as "War" dominated the protest-related conversations from the moment in started. This category contained all topics referring to the war in Ukraine, and made visible impact, as shown in the networks. While it provided valuable insights into the relationships between this topic and others, such as the position of the media, it is important to note that for the purposes of this study, the relevance lies in the occurrence of protests within the Netherlands. With this defined scope, the topic of the war in Ukraine, although influential globally, is not directly pertinent unless it becomes a catalyst for protests within the country.

5.2 Future research

SNA is already proven to be a useful method to provide insights into networks. It has however, not been broadly used in the context of protests as was done in this study. Useful patterns were identified that can aid in predicting future protest motivation. Additional to the findings resulted from this study, the experiences can be used to guide future research. In this section, the directions for future research are proposed.

5.2.1 Automated tweet retrieval

As mentioned in section 5.1.2, two different queries were used to retrieve tweets. While this may not have been the best method for this study, it could be a start for future work as unexpected differences in for instance hashtag use were identified in this study. One key difference is for example that within the original data set, the hashtag "demonstratie" was one of the most used hashtags, as well as the query word. For the evaluation data set, both the query words "demonstratie" and "protest" were used. Results showed that the hashtag "protest", was used significant more and was one of the hashtags that could be used to predict protests. These results suggest that generic hashtags can differ per protest and this provides space for future work. Possible questions that should be answered are for instance whether these generic hashtags differ per protests or subject of protests. Results can provide guidance in identifying the most relevant queries for protest related studies. This can hereafter be used to create a database with the most effective query words per protest, which can then be used for new studies.

5.2.2 The use of more data sources

As mentioned earlier, this study relied solely on Twitter as the data source. While valuable results are obtained from this data, it is worth considering future research that explores other data sources. Examining alternative data sources may reveal differences that provide new insights into protest-related social networks. Furthermore, investigating additional data sources can help strengthen the findings of this study by assessing the consistency and robustness of the results across different platforms, assuming minimal differences between the data sources.

5.2.3 Comparison with protests course

In this study, a particular event (Minister Kaag awaited with tortures) was employed to demonstrate how social networks can illustrate the shift in protest-related conversations from one topic to another when protests occur. However, it should be noted that this event was relatively small in scale compared to protests that are extensively announced in advance. Consequently, this study has limitations in identifying distinct patterns in protest-related conversations that could offer insights in to the potential course of the protest. Further research within this area could be highly relevant for the police, as it can assist in optimizing the allocation of officers based on a better understanding of the course of a protest.

5.2.4 Group characteristics comparison over time

While group characteristics are an important aspect of this study, the comparison is limited to two data sets, and the evaluation of these characteristics over time within a single data set was not conducted. However, exploring group characteristics over time holds significant relevance for future research, as it can provide additional insights into protest motivation.

To address this limitation and gain a more comprehensive understanding of the dynamics of group characteristics, it is recommended to employ the PReSNA method. By utilizing this approach, future studies can analyze the changes and differences in group characteristics over time within a specific protest-related social network. This analysis can offer valuable insights into the evolution of group dynamics, their motivations, and potentially uncover patterns or trends related to protest participation.

By examining group characteristics longitudinally, researchers can gain a deeper understanding of how these factors fluctuate and interact with other variables over time. This can contribute to a more nuanced understanding of protest-related social networks and further enhance our knowledge of the underlying dynamics that drive protest participation.

5.2.5 Sentiment analysis through natural language processing

One of the limitations of this study is the lack of integrating the opinion behind the use of hashtags, which introduces uncertainty regarding whether significant activity increases are driven by protesters or individuals opposed to the protests. This limitation opens up substantial potential for future studies to address this issue.

The integration of sentiment analysis within the data transformation phase of the PReSNA method could greatly enhance the utility of this approach for gaining insights into protest-related networks. By assigning opinion labels to the tweets/hashtags used, analyses can be conducted to examine the overall sentiment associated with specific topics. Additionally, networks can be developed to illustrate how different groups react to events, depicting the flow of tweets as more positive or negative within the networks.

These enhancements would not only provide a deeper understanding of the sentiment surrounding protest-related topics but also offer insights into the dynamics of different groups within these networks.

Chapter 6

Conclusion

The aim of this study is to leverage the insights gained from the pattern identification phase to develop a robust and effective method. Through the application of Social Network Analysis and the analysis of quantitative data derived from social networks, the study seeks to achieve this aim. This chapter serves as the conclusion of the study and provides a reflection on the research questions and findings.

To guide the study and accomplish the main research objective, four research questions were formulated. These research questions were designed to investigate specific aspects related to protest-related social networks and provide a comprehensive understanding of the phenomena under study. By addressing these research questions, the study aims to shed light on the dynamics, characteristics, and relationships within protest-related social networks.

As the main research objective is reached through answering the research questions. The questions are answered first, after which the main research objective is discussed. The first research question was defined to identify the current state of the art within the field of protest motivation and social network analysis. Research question one was formulated as follows:

- **RQ 1** "Which motivating factors to protest can be identified from existing literature?"

To answer this research question, the related literature was studied, which is detailed in chapter 2. The literature revealed that that protest motivation is influenced by various factors. The personality trait openness was found to have a positive effect on protest participation, while agreeableness and consciousness had a negative impact.

In addition to personality traits, four main pathways were identified as motivators: instrumentality, identity, ideology, and group-based anger. Among these pathways, identity, ideology, and group-based anger directly influence the strength of motivation. Moreover, identity enhances instrumental and ideological motives, intensifying anger, and thereby further increasing individuals' motivation to participate in protests.

Furthermore, an individual's social network plays a crucial role in protest motivation and can shape their opinions. Connected users within social networks tend to influence each other, leading to increased polarization within society. The analysis of these social networks provides valuable insights, as studies have found similar group dynamics between online and offline groups. The widespread use of social media further contributes to group dynamics by reducing the barrier to forming groups. Tailored news content contributes to polarization, as online groups are continuously exposed to content that reinforces their existing opinions while missing alternative perspectives. All of these factors ultimately contribute to protest motivation.

To gain more insights in topics that are related to protests, RQ2 was formulated as follows:

- **RQ 2** "What differences in hashtag usage can be used to analyze protest-related social networks?"

To analyze the protest related topics, a social network analysis of hashtags was conducted. The categorized hashtags revealed visually appealing networks that facilitated the immediate identification of category distributions. This analysis demonstrated that hashtags can be used to identify relations between topics, and by incorporating a dynamic component, temporal variations could be observed. In relation to protests, these temporal differences over time highlighted the most prominently discussed topics within protest-related social networks at specific moments in time. By selecting relevant hashtags derived from the networks, the temporal distribution of these hashtags could be depicted using line diagrams, exhibiting significant peaks corresponding to protest events.

The role of the media within this network could also be identified, having a central relation in the original network, illustrating the relations with the two main topics within the network, which were protests in the Netherlands and the war in Ukraine. While this central position could not be identified in the evaluation network due to the absence of distinct clusters, the hashtags that referred to media were still positioned together, suggesting that different media sources are associated with similar topics.

These results that showed interesting results therefore form animportant basis in fulfilling the main research objective.

In addition to the relationships of media-related hashtags, RQ3 was formulated to identify possible differences in media usage between clusters of a protest-related social network. Literature showed that tailor made news items, contributed to the polarization of networks, and it is therefore relevant to assess whether differences in media usage can be identified. RQ3 was formulated as follows:

- **RQ 3** "What differences in media consumption can be identified when comparing clusters of a protest-related social network?"

To find differences in media consumption among the identified modularity clusters, the hashtags were used. Among these hashtags, significant differences are found that show a unique media related hashtag for a particular group. This did not only show that this group seemed to use a unique hashtag to find each other, but the hashtags also gave insights in the opinion of this group on the media.

In addition to disparities in hashtag usage, variations were observed in the activity of the group categorized as "media." During the evaluation phase of the study, the dynamic network displayed the positioning of this group in the center of the network, suggesting high levels of activity. Furthermore, the dynamic network indicated that this media group had an activating effect, as overall activity increased whenever this group was active.

These findings underscore the significance of media-related dynamics within protest-related social networks and highlight the potential for further investigation using the PReSNA method. By delving deeper into the differences in hashtag usage and exploring the interactions and effects between groups, researchers can gain valuable insights into the role of media and its influence on protest-related activities.

The last research question was dedicated to the group dynamics within protest-related social networks. RQ4 was formulated as follows.

- **RQ 4** "What differences in group characteristics can be identified when comparing clusters of a protest-related social network?"

In order to identify groups within protest-related social networks, a function of Gephi 0.10.1, the program used for all networks, was used that was capable of identifying modularity classes based on node connections. The program successfully distinguished tens of classes, of which the eight largest classes, which were all compliant with the minimal number of members of ninety-five members, were extensively analyzed. Various attributes were examined to identify potential differences among these groups.

First, the centrality of the groups was examined, revealing significant differences among the classes. Notably, the two smallest classes (MC14 & MC16) within the analyzed groups were not only significant different from all other classes but also demonstrated significantly large differences compared to the mean.

Secondly, the average edge weight was compared to assess differences in tightness of the clusters. While the same clusters as in the centrality test showed significant differences, another class (MC0) emerged as the tightest class. However, further t-tests indicated that due to a large standard deviation, this class exhibited a significant medium difference compared to the mean, while the other two classes showed significantly large differences again.

Thirdly, differences in hashtag use were identified, focusing on the average amount of hashtags per hundred tweets and the average amount of tweets per user. The class MC4, which had the lowest number of tweets per user, also used the most hashtags per hundred tweets, making it unique with a large difference between these two values. Furthermore, the class MC0 employed a hashtag that was not used by any other class, although the frequency of this hashtag was equal to those used by all classes. This observation suggests that MC0 is the only class significantly using a unique hashtag compared to the average hashtag usage in the network.

The evaluation data set revealed that the two smallest classes from the analyzed groups were not active in this new data set. However, the class MC0, which already exhibited unique characteristics in the original data set, yielded similar results in this data set. This class displayed a significantly higher average edge weight and continued to utilize the same unique hashtag. Furthermore, MC4 again demonstrated the highest number of hashtags per hundred tweets, and while it did not have the lowest number of tweets per user this time, it still remained below average, resulting in the largest difference between these two values among all classes.

- **MRO** Develop a method to analyze protest-related social networks with Twitter data

The insights gained from addressing research questions SQ2, SQ3, and SQ4 have provided valuable contributions towards achieving the main research objective. These findings have revealed patterns within protest-related groups, suggesting the potential for predicting protest motivation. To further investigate this possibility, the PReSNA method is developed as a systematic approach for analyzing protest-related social networks on Twitter. By incorporating the experiences obtained in this study, the method aims to leverage the strengths of the research while addressing its limitations.

In conclusion, the results obtained from applying the PReSNA method have shown promising outcomes in terms of topic analysis through hashtag usage and group analysis through mentions. Further research can build upon this method to explore its full potential in analyzing and understanding protest-related social networks. By utilizing the PReSNA method and incorporating advancements, researchers can delve deeper into the dynamics of these networks, contributing to a comprehensive understanding of protest motivation and behavior by analyzing protest-related social networks.

Bibliography

Audiense (2023). *Audiense intelligence*. URL: https://www.audiense.com.

Back, Mitja D et al. (2010). "Facebook profiles reflect actual personality, not self-idealization". In: *Psychological science* 21.3, pp. 372–374.

Bahrami, Mohsen et al. (May 2018). "Twitter Reveals: Using Twitter Analytics to Predict Public Protests." In: *arXiv: Social and Information Networks*. URL: https://arxiv.org/pdf/1805.00358.

Bakker, Jeroen et al. (2021). "Van scherm naar straat: Hoe sociale media-conversaties protest op straat mobiliseren". In.

Baumann, Fabian et al. (2020). "Modeling echo chambers and polarization dynamics in social networks". In: *Physical Review Letters* 124.4, p. 048301.

Bruns, Axel and Jean Burgess (2011). "The use of Twitter hashtags in the formation of ad hoc publics". In: *Proceedings of the 6th European consortium for political research (ECPR) general conference 2011*. The European Consortium for Political Research (ECPR), pp. 1–9.

Chenoweth, Erica et al. (2022). *Who Protests, What Do They Protest, and Why?* Tech. rep. National Bureau of Economic Research.

Cochran, William G (1954). "Some methods for strengthening the common χ 2 tests". In: *Biometrics* 10.4, pp. 417–451.

ColorBrewer (2021). *ColorBrewer: Color Advice for Maps*. URL: https://colorbrewer2.org/#type=qualitative&scheme=Paired&n=12.

Creswell, J. C. (2003). *Research design: Qualitative, quantitative and mixed methods approaches*. 2nd ed. Thousand Oaks, CA: SAGE Publications.

De Weerd, Marga and Bert Klandermans (1999). "Group identification and political protest: Farmers' protest in the Netherlands". In: *European Journal of Social Psychology* 29.8, pp. 1073–1095.

DeYoung, Colin G, Lena C Quilty, and Jordan B Peterson (2007). "Between facets and domains: 10 aspects of the Big Five." In: *Journal of personality and social psychology* 93.5, p. 880.

Diener, Edward and Richard E Lucas (2019). "Personality traits". In: *General psychology: Required reading* 278.

Donnellan, M Brent et al. (2006). "The mini-IPIP scales: tiny-yet-effective measures of the Big Five factors of personality." In: *Psychological assessment* 18.2, p. 192.

Enthoven, Anniek and Sjors Hofstede (Jan. 2022). "Den Haag heeft grootste tekort aan agenten: 'Dit moet niet te lang duren'". In: URL: https://www.omroepwest.nl/nieuws/4512779/den-haag-heeft-grootste-tekort-aan-agenten-dit-moet-niet-te-lang-duren.

Es, Karin van, Daniela van Geenen, and Thomas Boeschoten (2015). "Re-imagining Television Audience Research: Tracing Viewing Patterns on Twitter". In: *M/C Journal* 18.6.

Forsyth, Donelson R (2018). *Group dynamics*. Cengage Learning.

Golbeck, Jennifer et al. (2011). "Predicting personality from twitter". In: *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*. IEEE, pp. 149–156.

Goldberg, Lewis R et al. (1999). "A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models". In: *Personality psychology in Europe* 7.1, pp. 7–28.
González-Bailón, Sandra et al. (2011). "The dynamics of protest recruitment through an online network". In: *Scientific reports* 1.1, pp. 1–7.
Gosling, Samuel D, Carson J Sandy, and Jeff Potter (2010). "Personalities of self-identified "dog people" and "cat people"". In: *Anthrozoös* 23.3, pp. 213–222.
Groot, Diederik de (Feb. 2022). "Grote toename demonstraties in 2021, gemeenten steeds vaker niet vooraf ingelicht". In: URL: https://www.bnr.nl/nieuws/binnenland/10466698/grote-toename-demonstraties-in-2021-gemeenten-steeds-vaker-niet-vooraf-ingelicht.
Hevner, Alan R et al. (2004). "Design science in information systems research". In: *MIS quarterly*, pp. 75–105.
Hogg, Michael A, Deborah J Terry, and Katherine M White (1995). "A tale of two theories: A critical comparison of identity theory with social identity theory". In: *Social psychology quarterly*, pp. 255–269.
Jenkins, Henry (2006). *Convergence Culture*. Amsterdam, Netherlands: Amsterdam University Press.
John, Oliver P, Laura P Naumann, and Christopher J Soto (2008). "Paradigm shift to the integrative Big Five trait taxonomy: History, measurement, and conceptual issues." In.
Jost, John T, Tessa V West, and Samuel D Gosling (2009). "Personality and ideology as determinants of candidate preferences and "Obama conversion" in the 2008 US presidential election". In: *Du Bois Review: Social Science Research on Race* 6.1, pp. 103–124.
Kingdom, Statistics (2023). *Chi-Squared test calculator*. URL: https://www.statskingdom.com/310GoodnessChi.html.
Klandermans, Bert (1984). "Mobilization and participation: Social-psychological expansisons of resource mobilization theory". In: *American sociological review*, pp. 583–600.
— (2002). "How group identification helps to overcome the dilemma of collective action". In: *American Behavioral Scientist* 45.5, pp. 887–900.
Larson, Jennifer M et al. (2019). "Social networks and protest participation: Evidence from 130 million Twitter users". In: *American Journal of Political Science* 63.3, pp. 690–705.
Lazer, David et al. (2010). "The coevolution of networks and political attitudes". In: *Political Communication* 27.3, pp. 248–274.
Meltwater (2023). *Media Monitoring*. URL: https://explore.meltwater.com/nl-media-monitoring?utm_source=google.
Mitchell, Amy et al. (2014). "Political polarization & media habits". In.
Mummendey, Amelie et al. (1999). "Strategies to cope with negative social identity: predictions by social identity theory and relative deprivation theory." In: *Journal of personality and social psychology* 76.2, p. 229.
Mundt, Marcia, Karen Ross, and Charla M Burnett (2018). "Scaling social movements through social media: The case of Black Lives Matter". In: *Social Media+ Society* 4.4, p. 2056305118807911.
Müter, Laurens, Meagan Loerakker, and Remco Veltkamp (2023). *Information Extraction of Protest Related Tweets: an Exploratory Analysis*.
Nassauer, Anne (2016). "From peaceful marches to violent clashes: A micro-situational analysis". In: *Social Movement Studies* 15.5, pp. 515–530.

Nguyen, Dong et al. (2013). “" How old do you think I am?" A study of language and age in Twitter”. In: *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 7. 1, pp. 439–448.

NOS (Sept. 2021). “Zo schat de politie in hoeveel mensen er op een demonstratie komen”. In: URL: https://nos.nl/artikel/2397182-zo-schat-de-politie-in-hoeveel-mensen-er-op-een-demonstratie-komen.

— (Feb. 2023). “Rutte: conclusies over Groningen zijn hard binnengekomen, ook bij mij persoonlijk ”. In: URL: https://nos.nl/collectie/13902/artikel/2465107-rutte-conclusies-over-groningen-zijn-hard-binnengekomen-ook-bij-mij-persoonlijk.

Ooijevaar, JW and GLM Kraaykamp (2005). “Links in beeld: Een explorerend onderzoek naar de sociale kenmerken van extreem links in Nederland [A focus on the radical left: An exploration of social characteristics of left wing radicals in The Netherlands]”. In.

Opp, Karl-Dieter and Hermann Brandstätter (2010). “Political protest and personality traits: A neglected link”. In: *Mobilization: An International Quarterly* 15.3, pp. 323–346.

Peffers, Ken et al. (2007). “A design science research methodology for information systems research”. In: *Journal of management information systems* 24.3, pp. 45–77.

Pols, Gidi (Feb. 2021). “Van klimaatmars tot coronaprotest: Nederlanders gaan steeds vaker de straat op”. In: URL: https://www.trouw.nl/binnenland/van-klimaatmars-tot-coronaprotest-nederlanders-gaan-steeds-vaker-de-straat-op~b397c93b/?referrer=https%3A%2F%2Fwww.google.com%2F.

Quattrociocchi, Walter, Antonio Scala, and Cass R Sunstein (2016). “Echo chambers on Facebook”. In: *Available at SSRN 2795110*.

Raafat, Ramsey M, Nick Chater, and Chris Frith (2009). “Herding in humans”. In: *Trends in cognitive sciences* 13.10, pp. 420–428.

Rentfrow, Peter J and Samuel D Gosling (2003). “The do re mi’s of everyday life: the structure and personality correlates of music preferences.” In: *Journal of personality and social psychology* 84.6, p. 1236.

RockContent (Mar. 2022). “Best Practices: Maximum Elements For Different Visualization Types”. In: *Rock Content*. URL: https://rockcontent.com/blog/maximum-elements-for-visualization-types/.

Rogers, Richard (2013). *Digital methods*. MIT press.

Selfhout, Maarten et al. (2010). “Emerging late adolescent friendship networks and Big Five personality traits: A social network approach”. In: *Journal of personality* 78.2, pp. 509–538.

Shi, Juan et al. (2017). “Understanding and predicting individual retweeting behavior: Receiver perspectives”. In: *Applied Soft Computing* 60, pp. 844–857.

Shirky, Clay (2009). *Here Comes Everybody*. Amsterdam, Netherlands: Adfo Books.

Sobolev, Anton et al. (June 2020). “News and Geolocated Social Media Accurately Measure Protest Size Variation”. In: *American Political Science Review* 114.4, pp. 1343–1351. DOI: 10.1017/s0003055420000295. URL: http://dx.doi.org/10.1017/s0003055420000295.

Stangor, Charles (2014). *Research methods for the behavioral sciences*. Cengage Learning.

Statistiek, Centraal Bureau voor de (2023). *Onderzoek naar gebruik van sociale media*. URL: https://www.cbs.nl/nl-nl/over-ons/onderzoek-en-innovatie/project/onderzoek-naar-gebruik-van-sociale-media.

Stekelenburg, Jacquelien van, Bert Klandermans, and Wilco W. van Dijk (Jan. 2011). “Combining motivations and emotion: The motivational dynamics of protest participation". In: *Revista de Psicología Social* 26.1, pp. 91–104. DOI: 10.1174/021347411794078426. URL: http://dx.doi.org/10.1174/021347411794078426.

Stroud, Natalie Jomini (2008). "Media use and political predispositions: Revisiting the concept of selective exposure". In: *Political Behavior* 30, pp. 341–366.

Stürmer, Stefan and Bernd Simon (2004). "The role of collective identification in social movement participation: A panel study in the context of the German gay movement". In: *Personality and Social Psychology Bulletin* 30.3, pp. 263–277.

Sunstein, Cass R (1999). "The law of group polarization". In: *University of Chicago Law School, John M. Olin Law & Economics Working Paper* 91.

Van Eck, Levi (Sept. 2017). "Vijf jaar na Project X Haren, hoe kon het zo mis gaan?" In: URL: https://nos.nl/artikel/2194074-vijf-jaar-na-project-x-haren-hoe-kon-het-zo-mis-gaan.

Van Stekelenburg, Jacquelien and Bert Klandermans (2017). "Individuals in movements: A social psychology of contention". In: *Handbook of social movements across disciplines*, pp. 103–139.

Veerbeek, Joris et al. (2022). "Online antisemitisme in 2020". In.

Walgrave, Stefaan et al. (2013). "Why do people protest? Comparing demonstrators' motives across issues and nations". In: *Media, movements and politics*.

Wellman, Barry (1997). "An electronic group is virtually a social network". In: *Culture of the Internet* 4, pp. 179–205.

Welser, Howard T et al. (2007). "Visualizing the signatures of social roles in online discussion groups". In: *Journal of social structure* 8.2, pp. 1–32.

Wieringa, Maranke et al. (2018). "Political topic-communities and their framing practices in the Dutch Twittersphere". In: *Internet Policy Review* 7.2, pp. 1–16.

Wright, Stephen C (2003). "Strategic collective action: Social psychology and social change". In: *Blackwell handbook of social psychology: Intergroup processes*, pp. 409–430.

Zibran, Minhaz Fahim (2007). "Chi-squared test of independence". In: *Department of Computer Science, University of Calgary, Alberta, Canada* 1.1, pp. 1–7.

Appendix A

Appendix - Enlarged visualizations

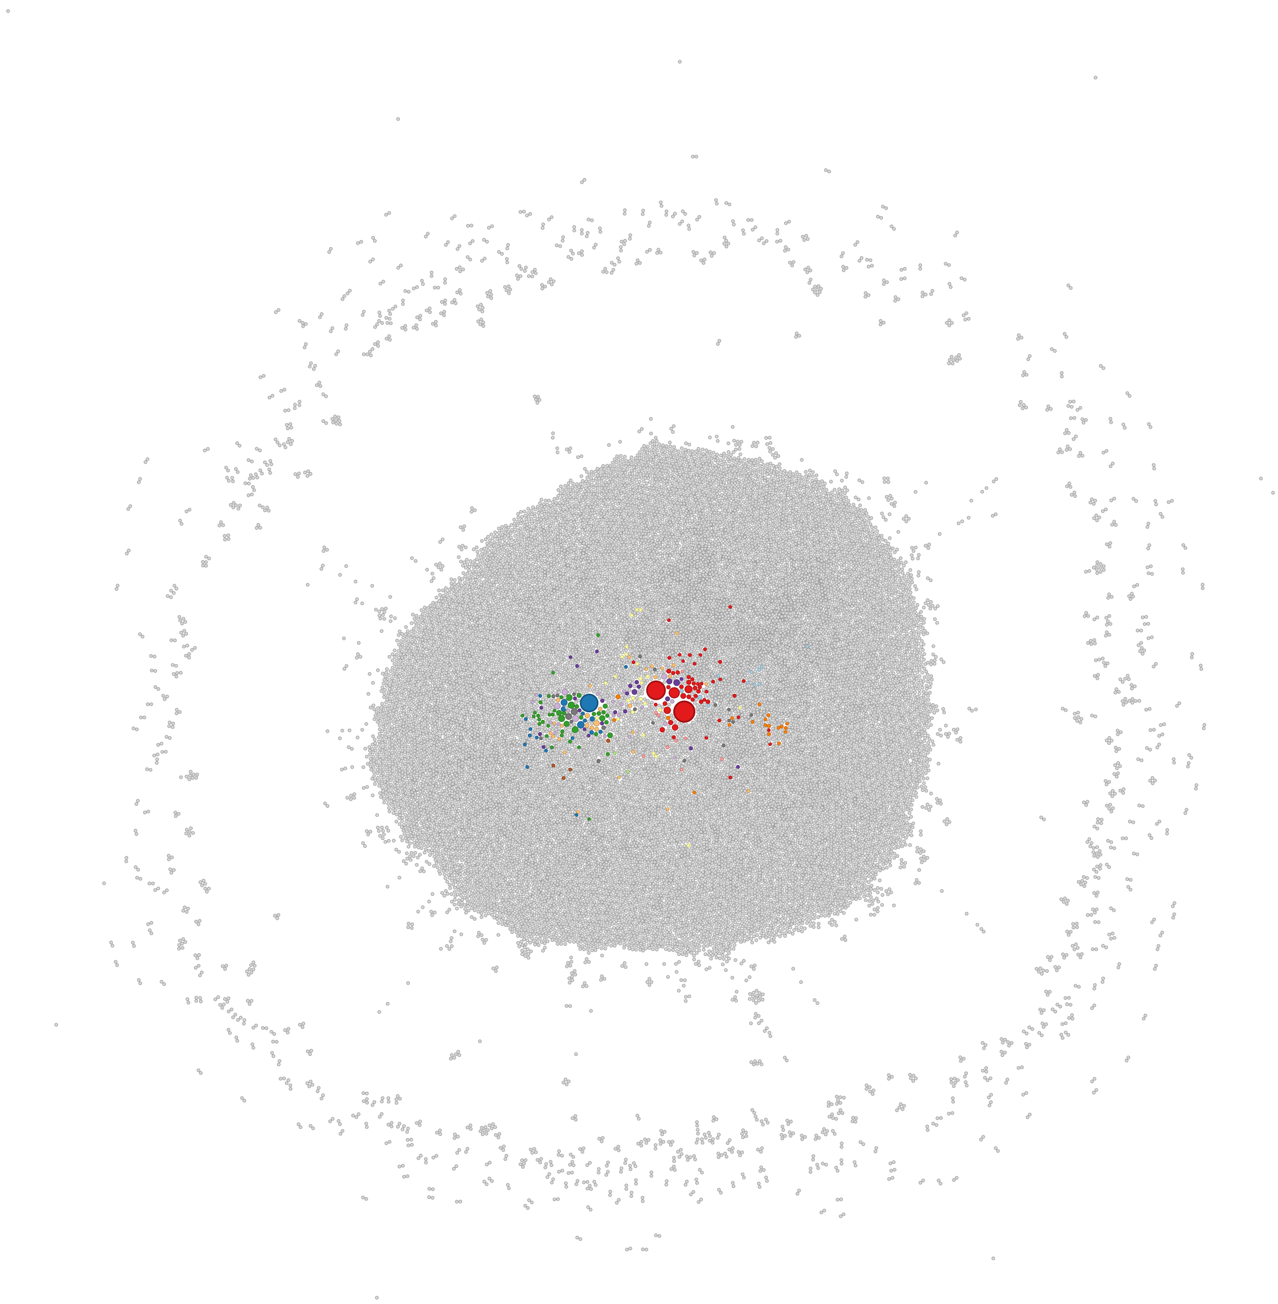

FIGURE A.1: Categorized social network of hashtags. Filter(s): -. Remaining nodes/edges: 24,319N (100%) Coloring used as defined in Figure 4.5.

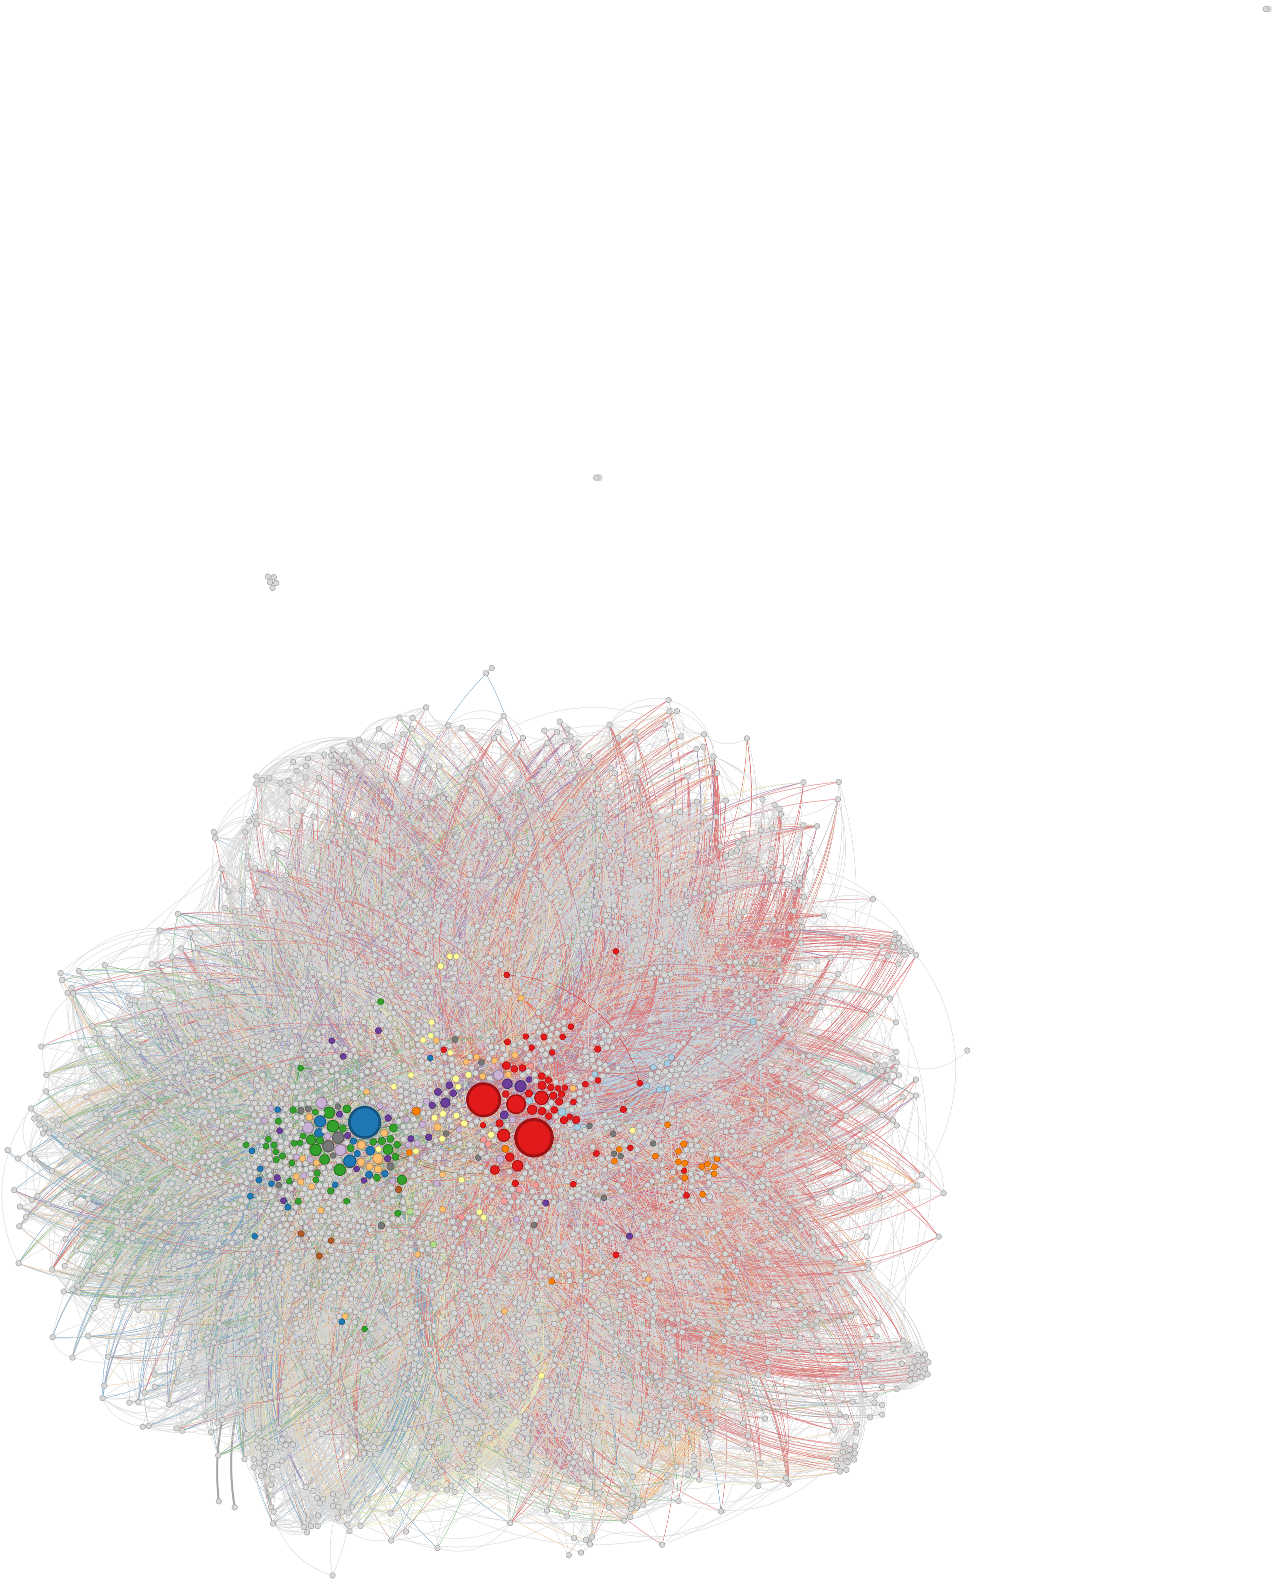

FIGURE A.2: Categorized social network of hashtags. Filter(s): $WD \geq 20$. Remaining nodes/edges: 3,312N (13.6%). Coloring used as defined in Figure 4.5.

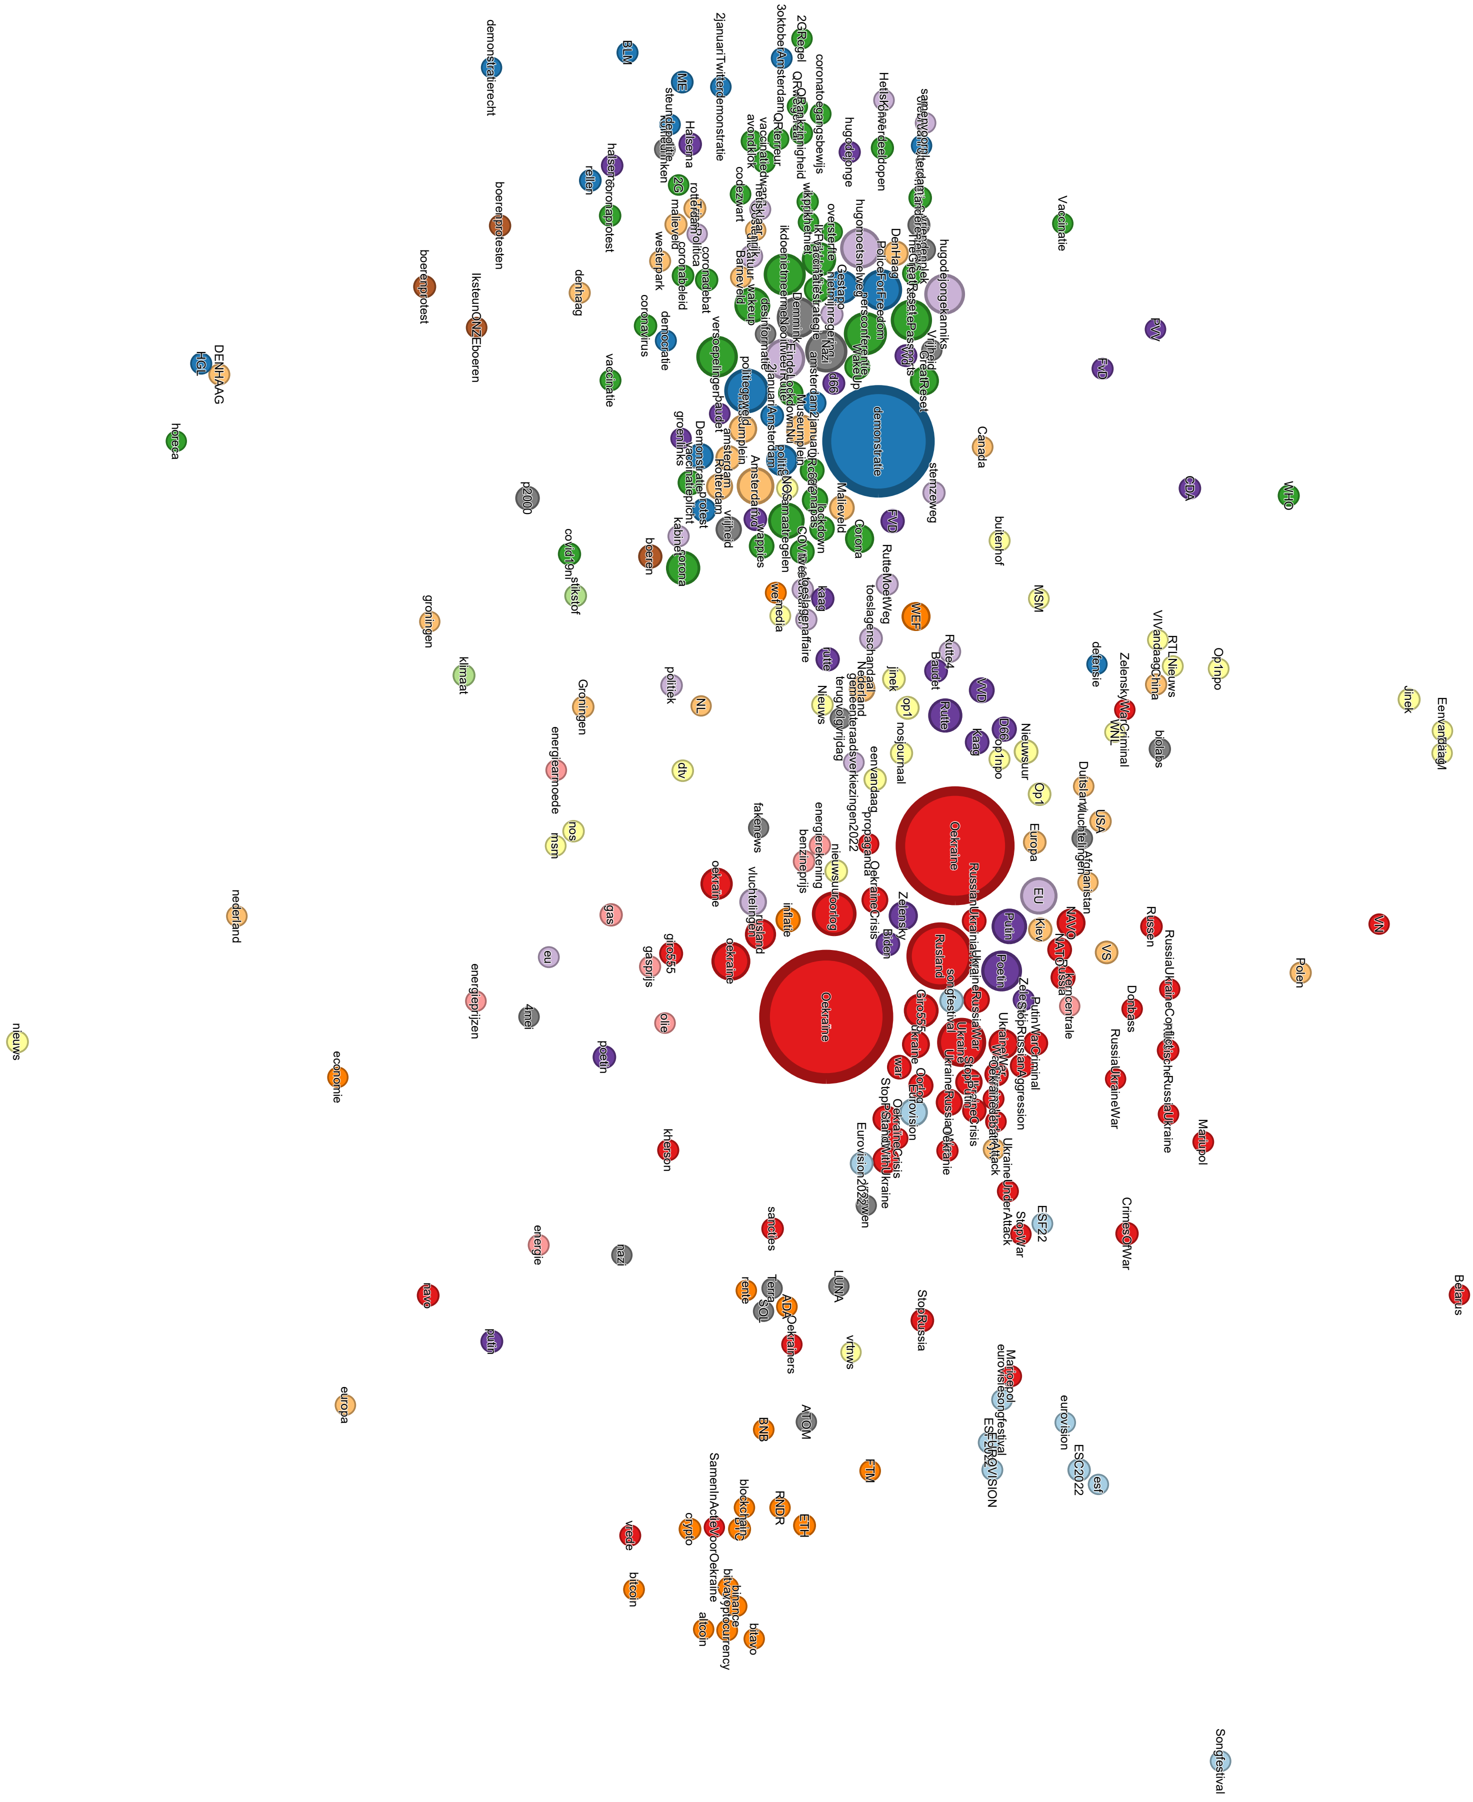

FIGURE A.3: Categorized social network of hashtags. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N (1.2%). Coloring used as defined in Figure 4.5.

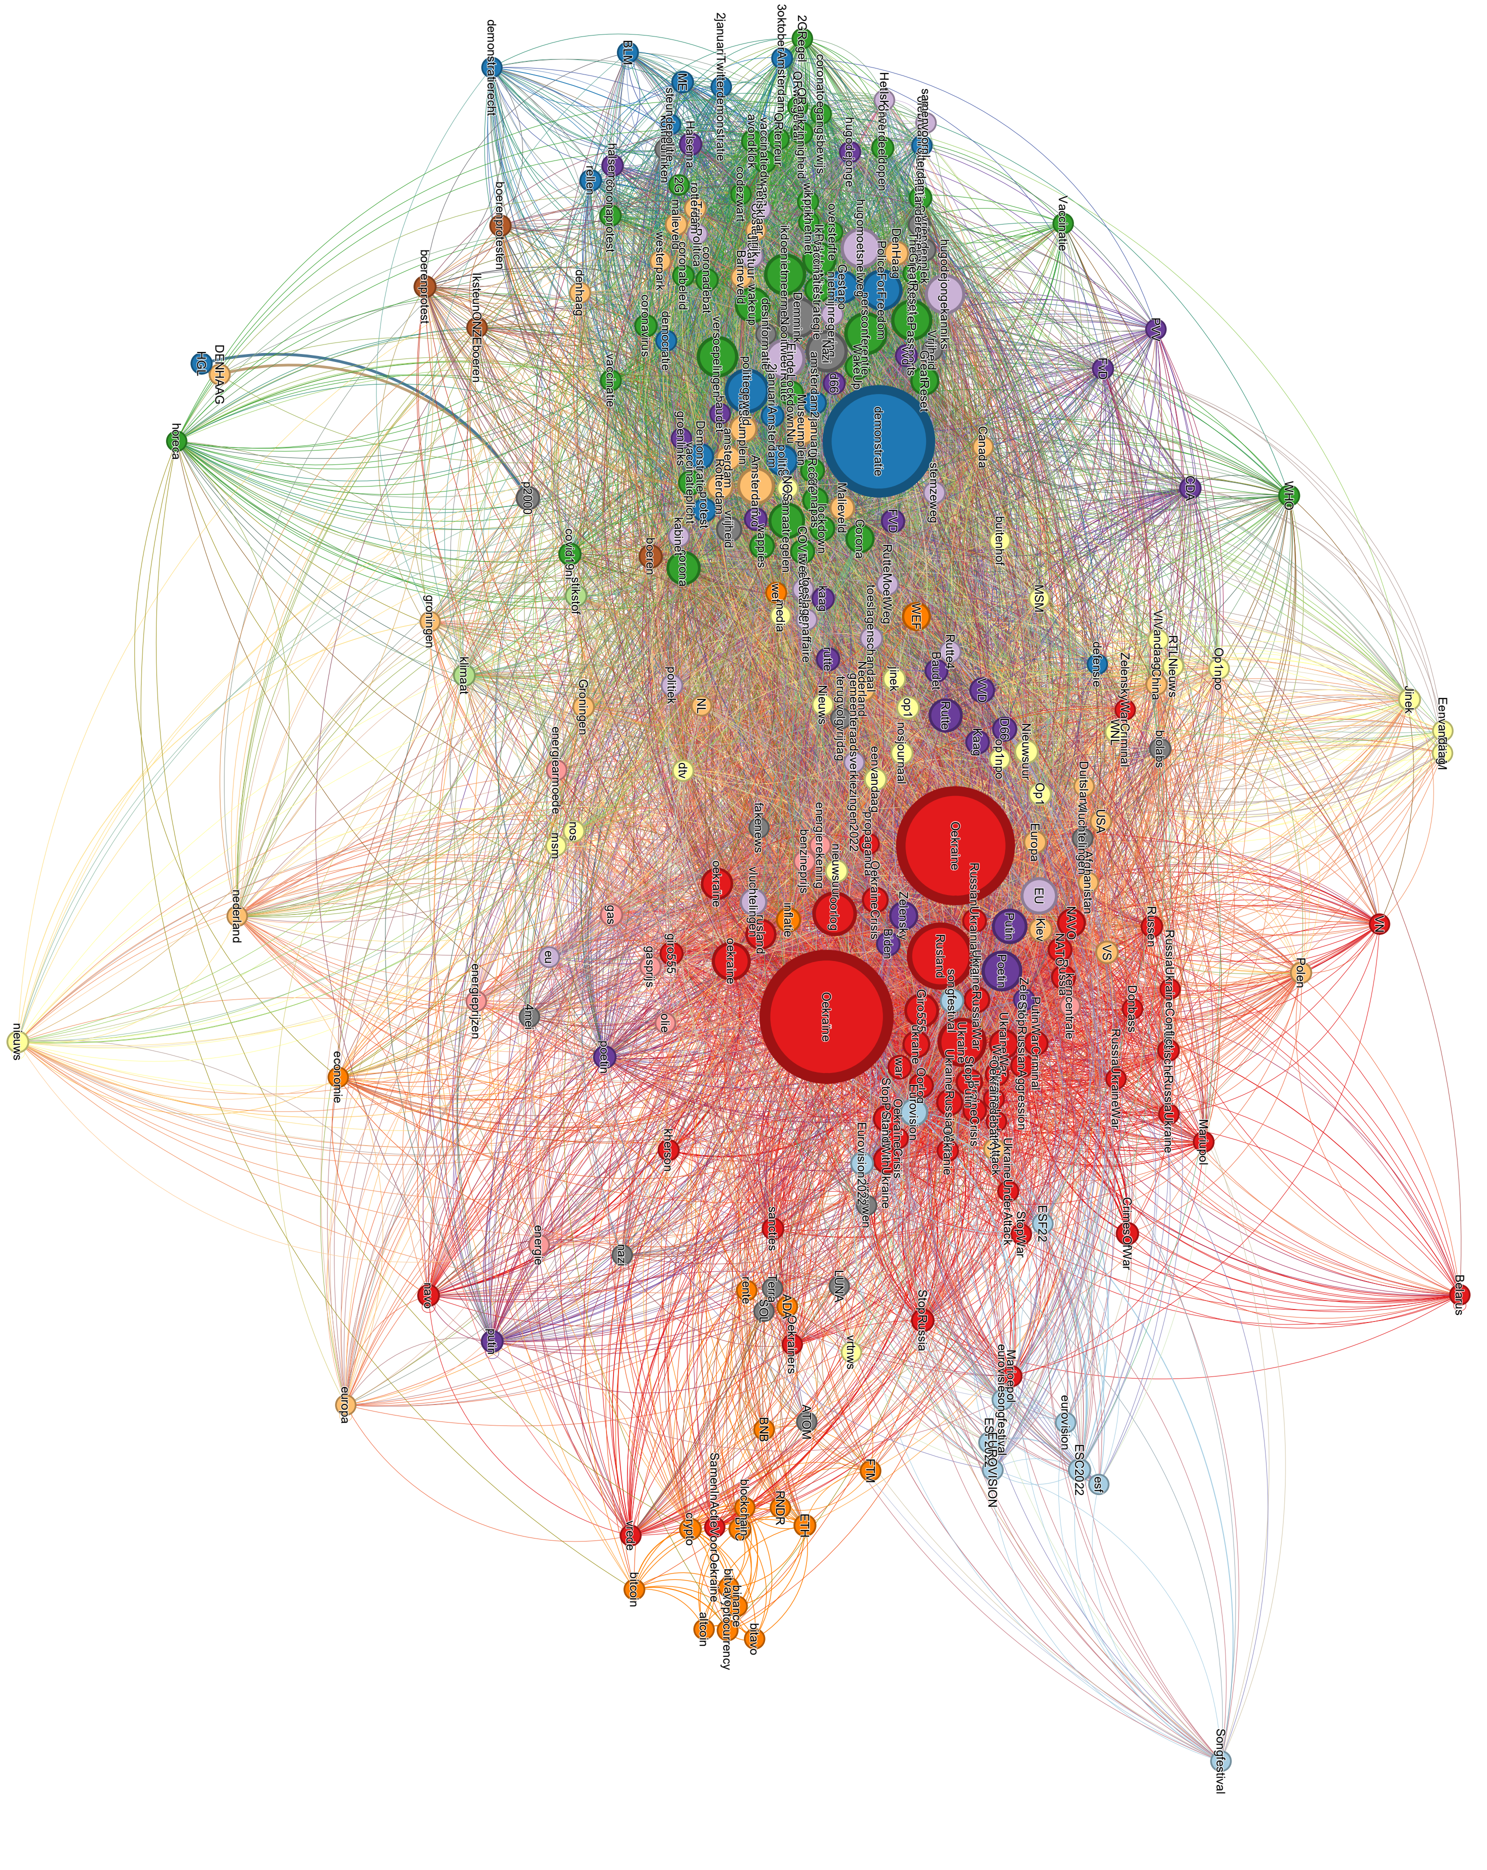

FIGURE A.4: Categorized social network of hashtags. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N(1.2%) - 10,465E (6.6%)). Coloring used as defined in Figure 4.5.

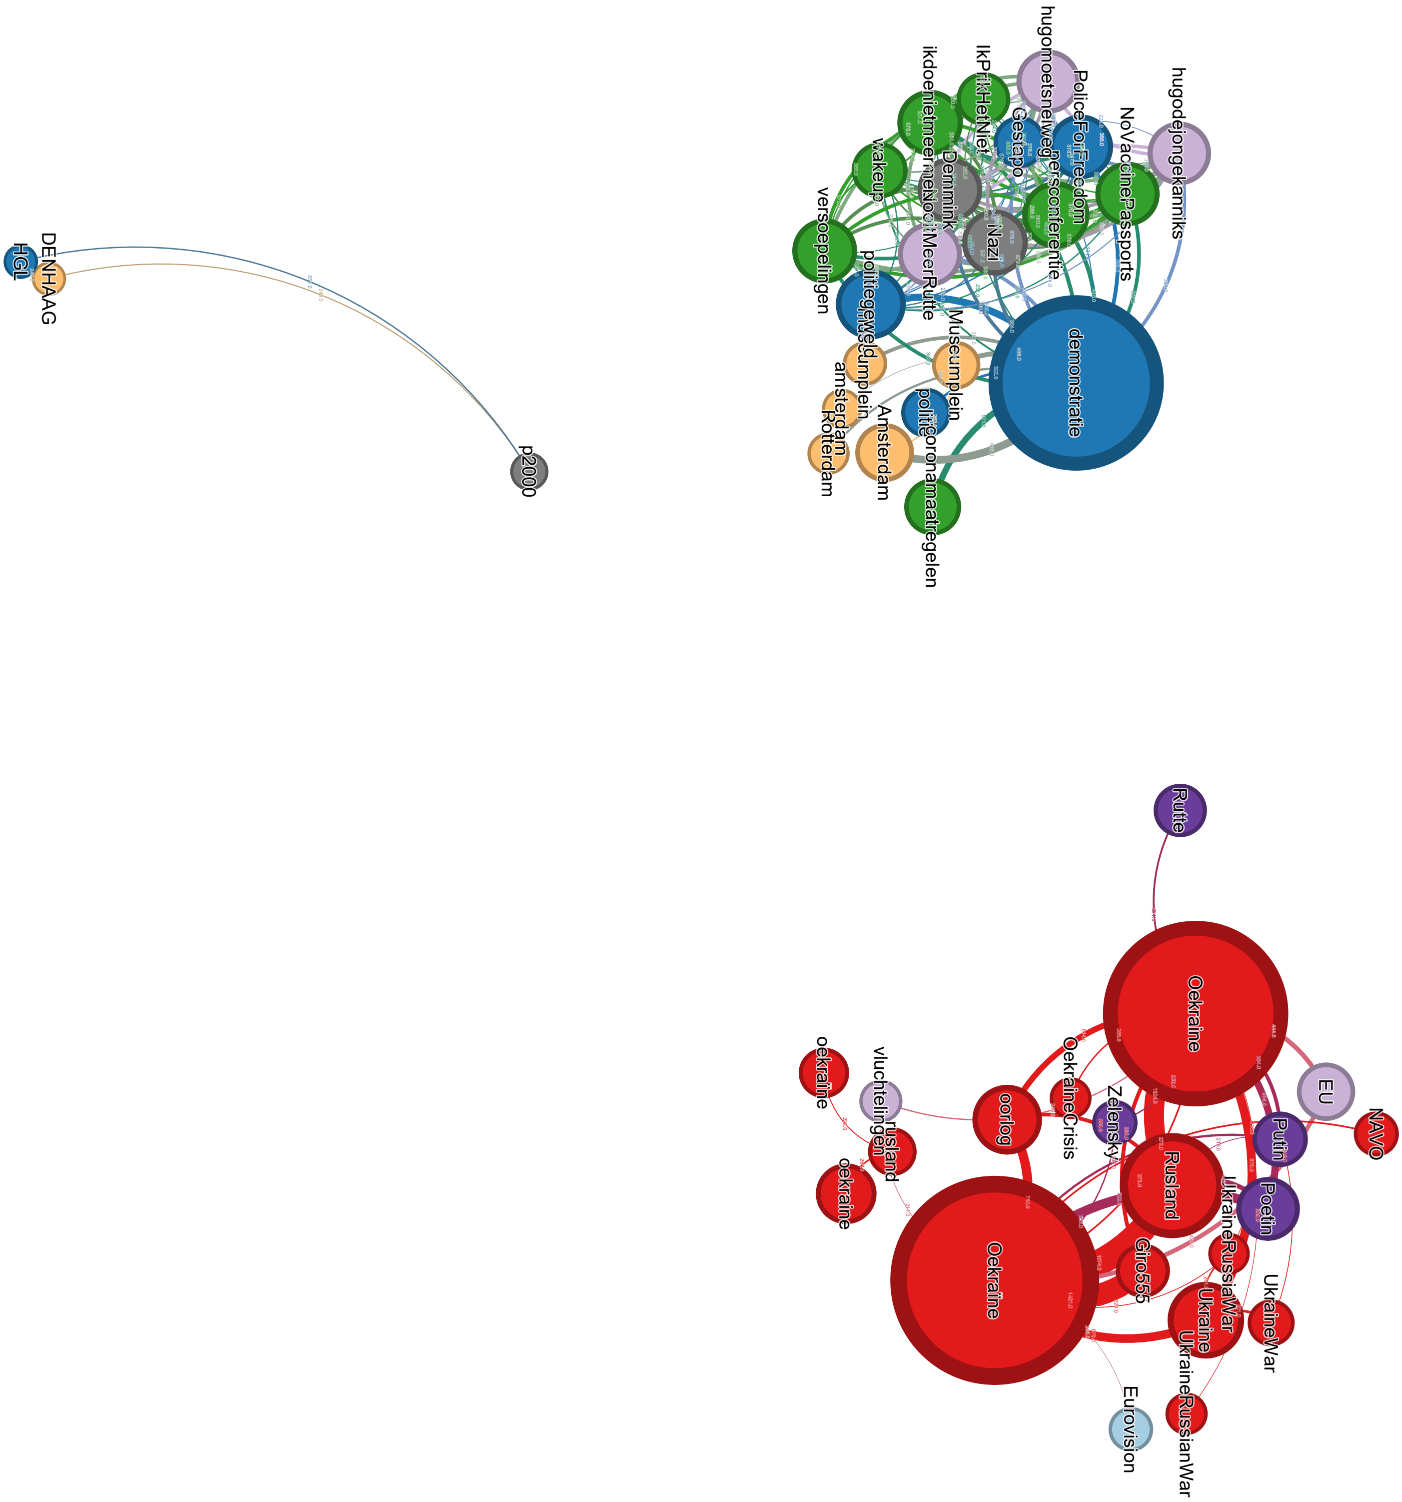

FIGURE A.5: Categorized social network of strongest relationships between hashtags. Filter(s): $EW \geq 200$. Remaining nodes/edges: 46N(0.2%) - 150E(0.09%). Coloring used as defined in Figure 4.5

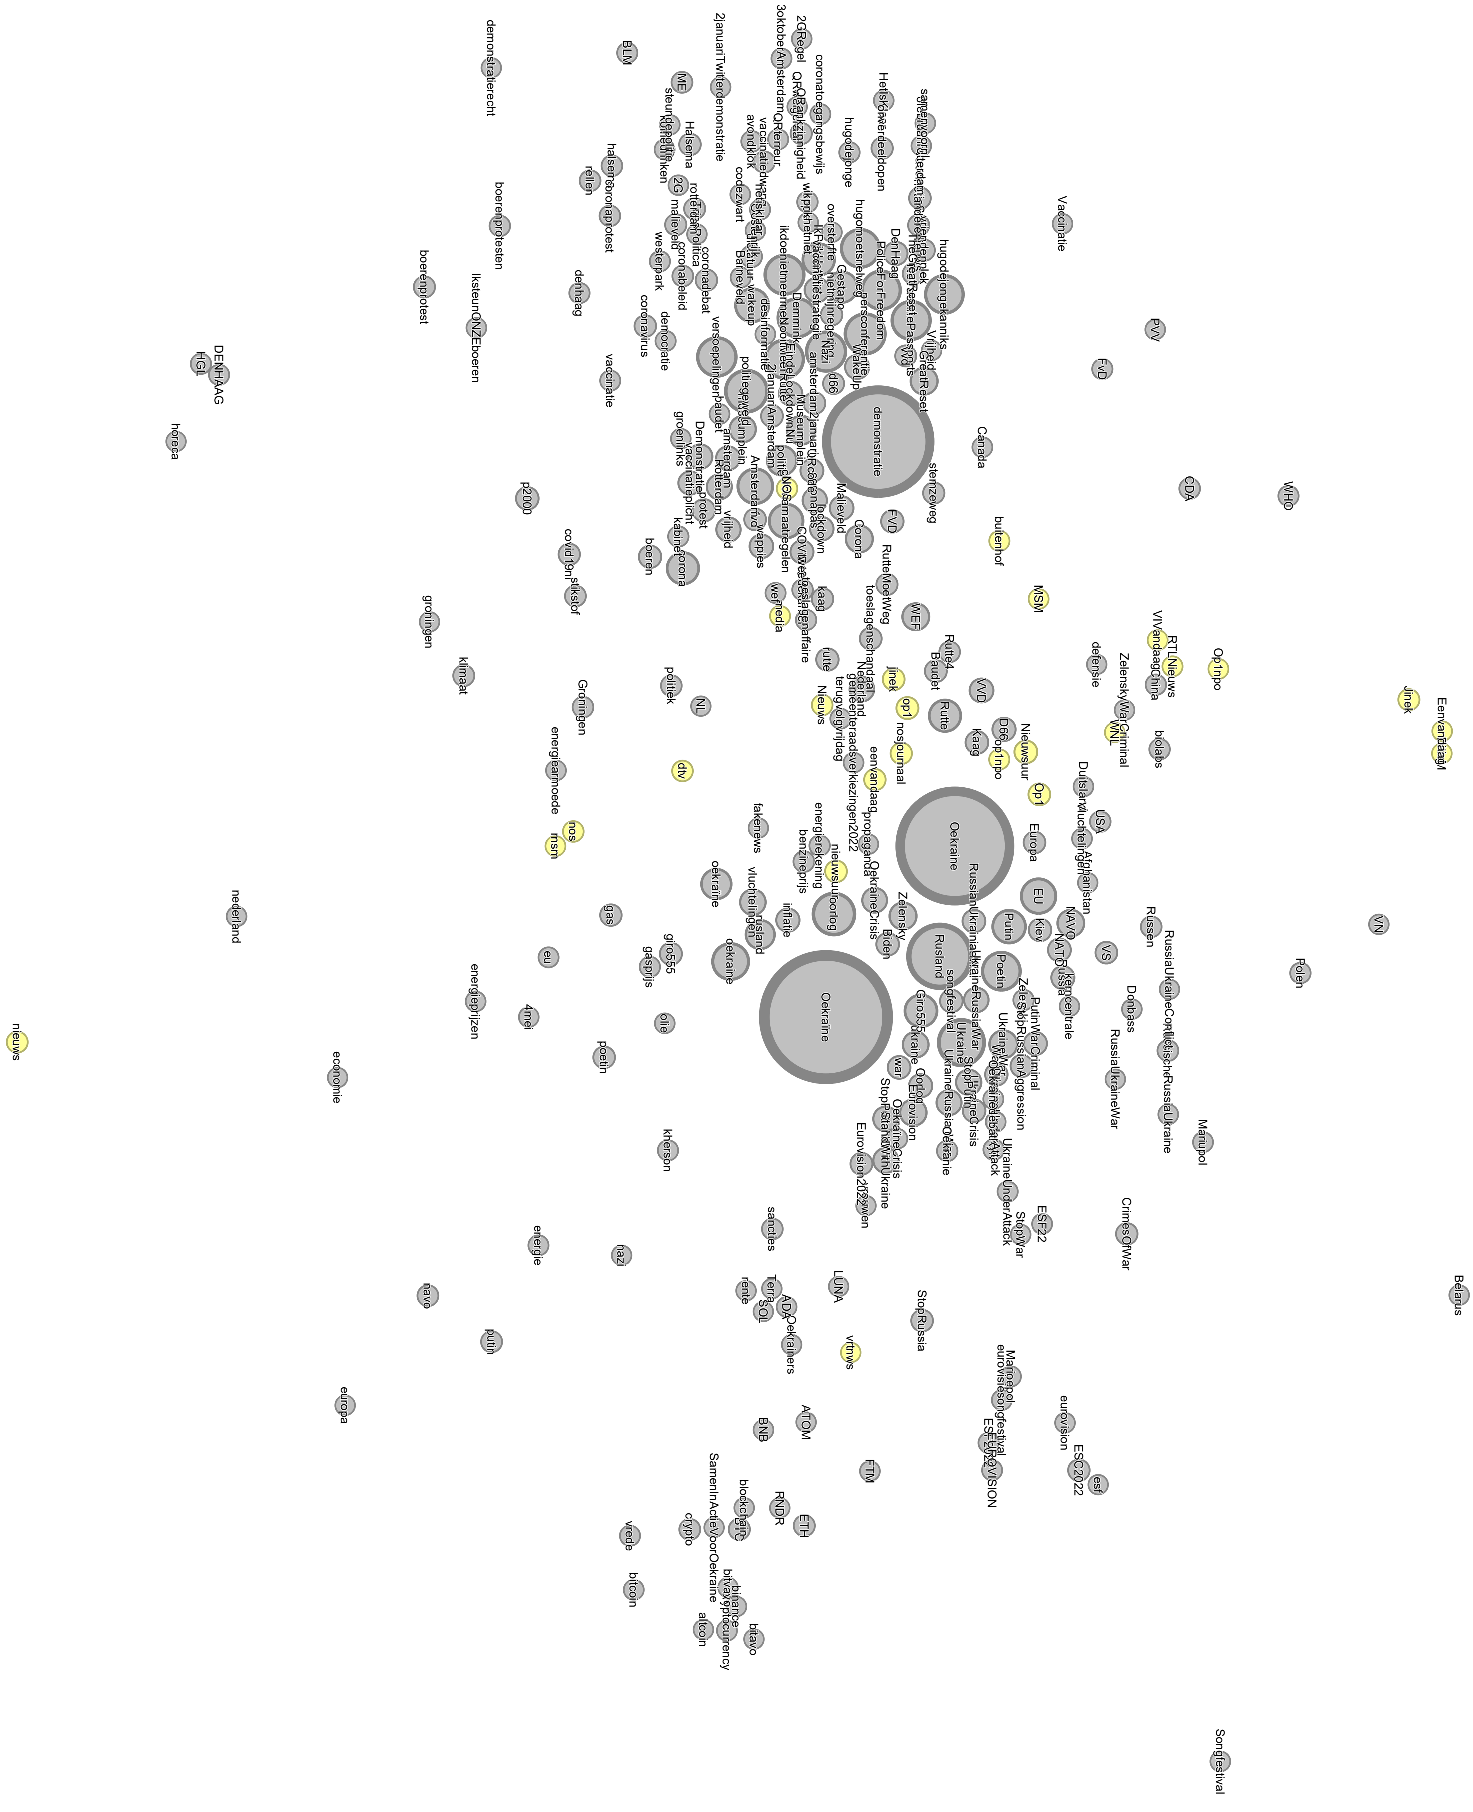

FIGURE A.6: Position of hashtags categorized as media. Filter(s): $WD \geq 300$. Remaining nodes/edges: 284N(1.2%)

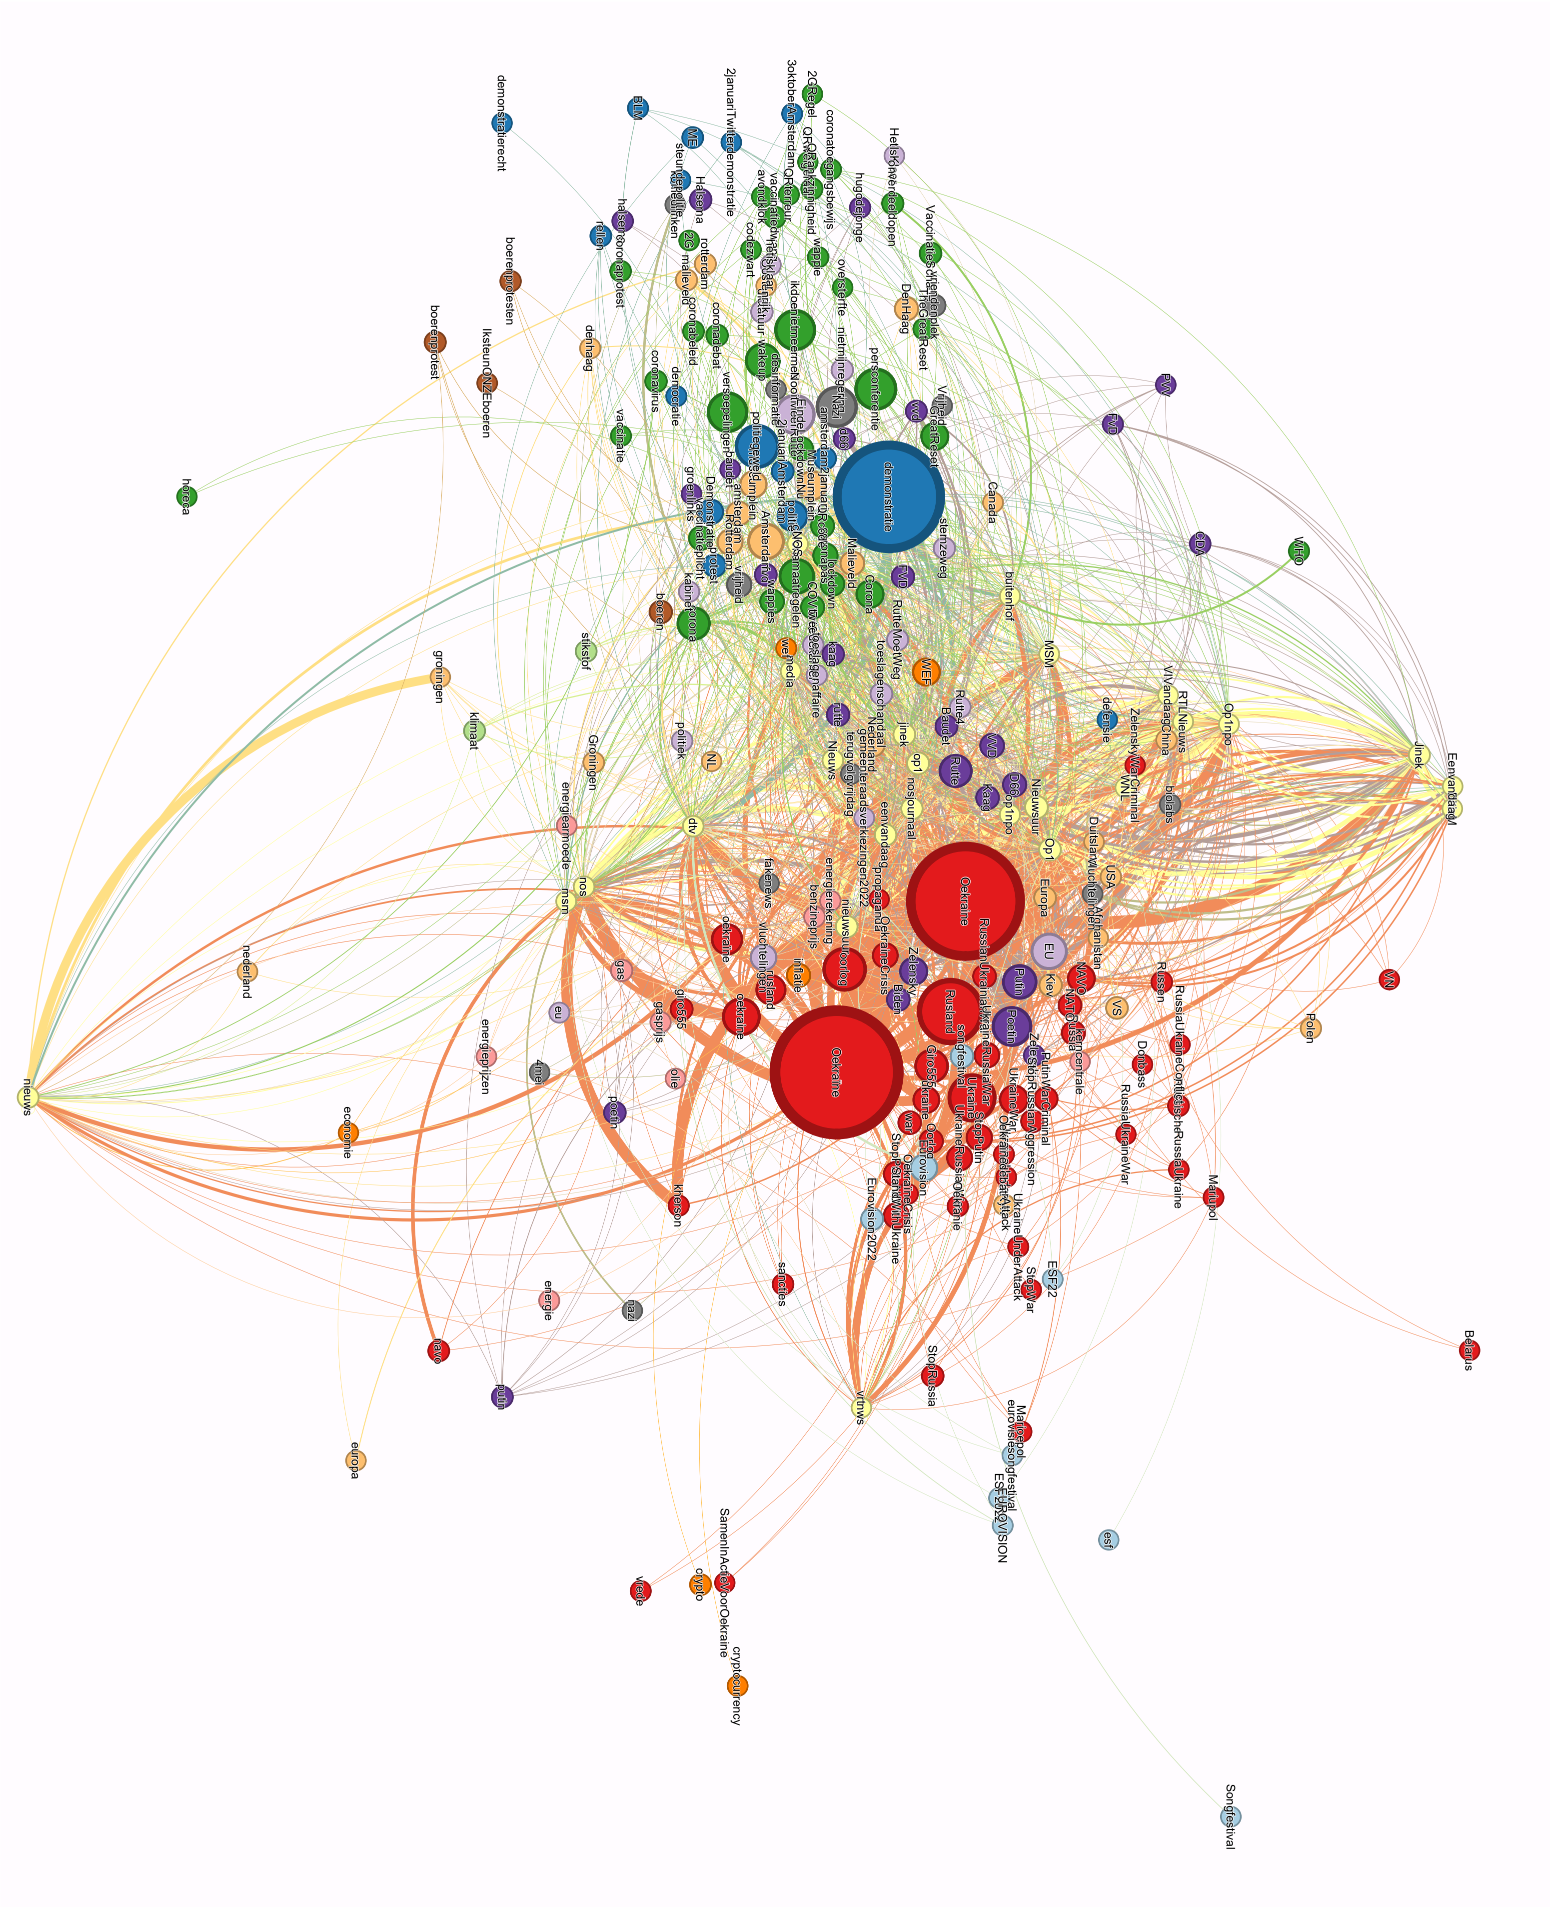

FIGURE A.7: Relationships of hashtags categorized as media. Filter(s): $WD \geq 300$ AND & ES OR ET = Media AND $D \geq 1$. Remaining nodes/edges: 239N(1.0%) - 1,745E(1.1%). Coloring used as

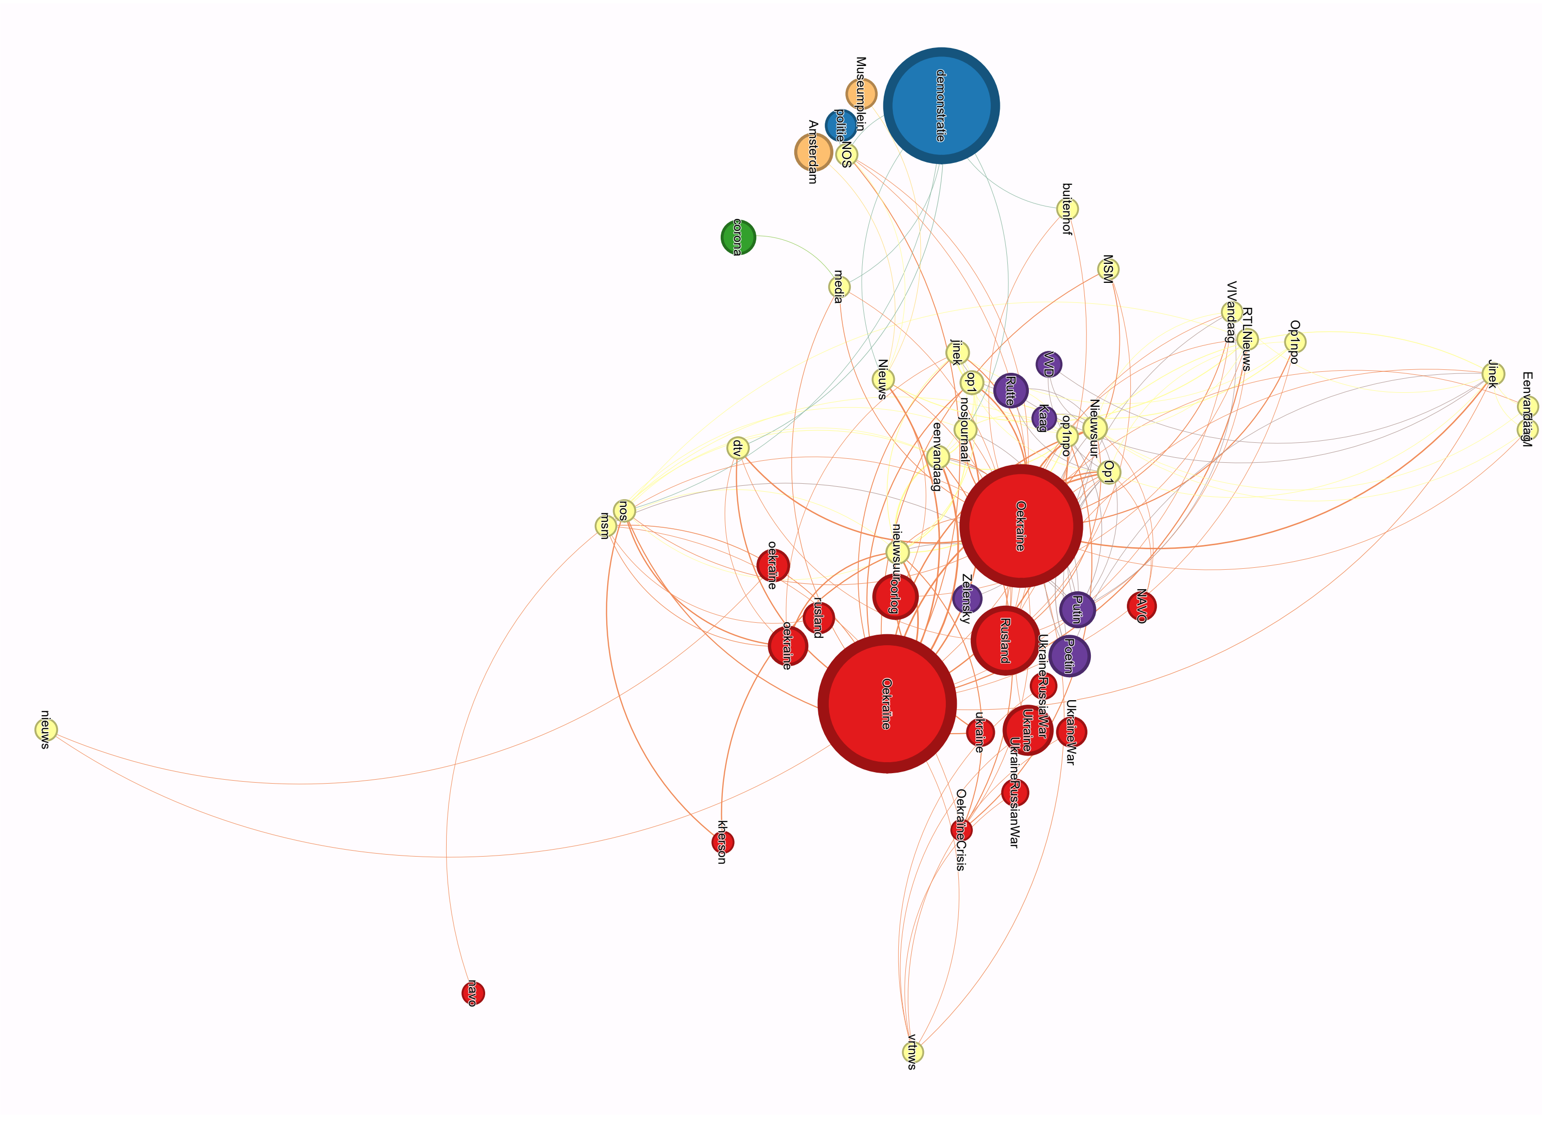

FIGURE A.8: Tights relationships of hashtags categorized as media. Filter(s): $WD \geq 300$ $EW \geq 11$. Remaining nodes/edges: 51N(0.2%) - 150E(0.09%). Coloring used as defined in Figure 4.5

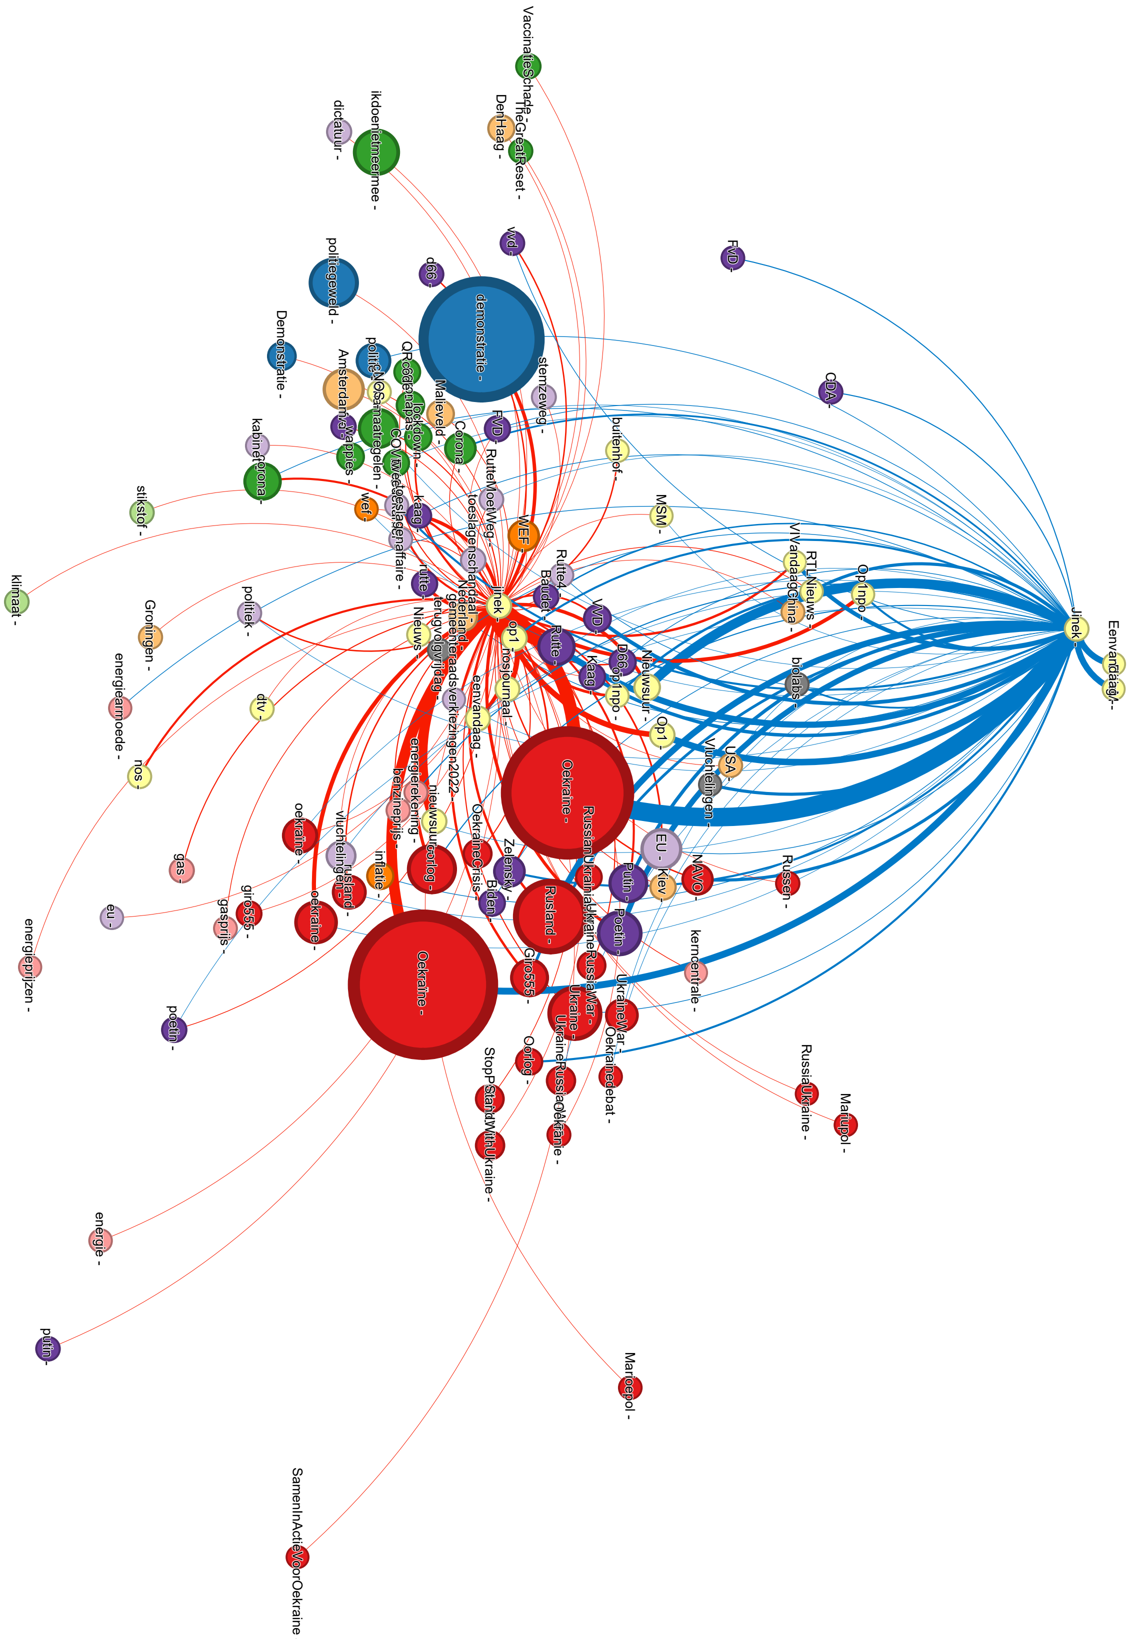

FIGURE A.9: Relationship differences between the hashtags Jinek and jinek. Edges connected to Jinek visualized in blue. Edges connected to jinek visualized in red. Thickness of edges indicating the edge weight (rescaled to a minimum of 1 and a maximum of 50). Coloring used as defined in Figure 4.5

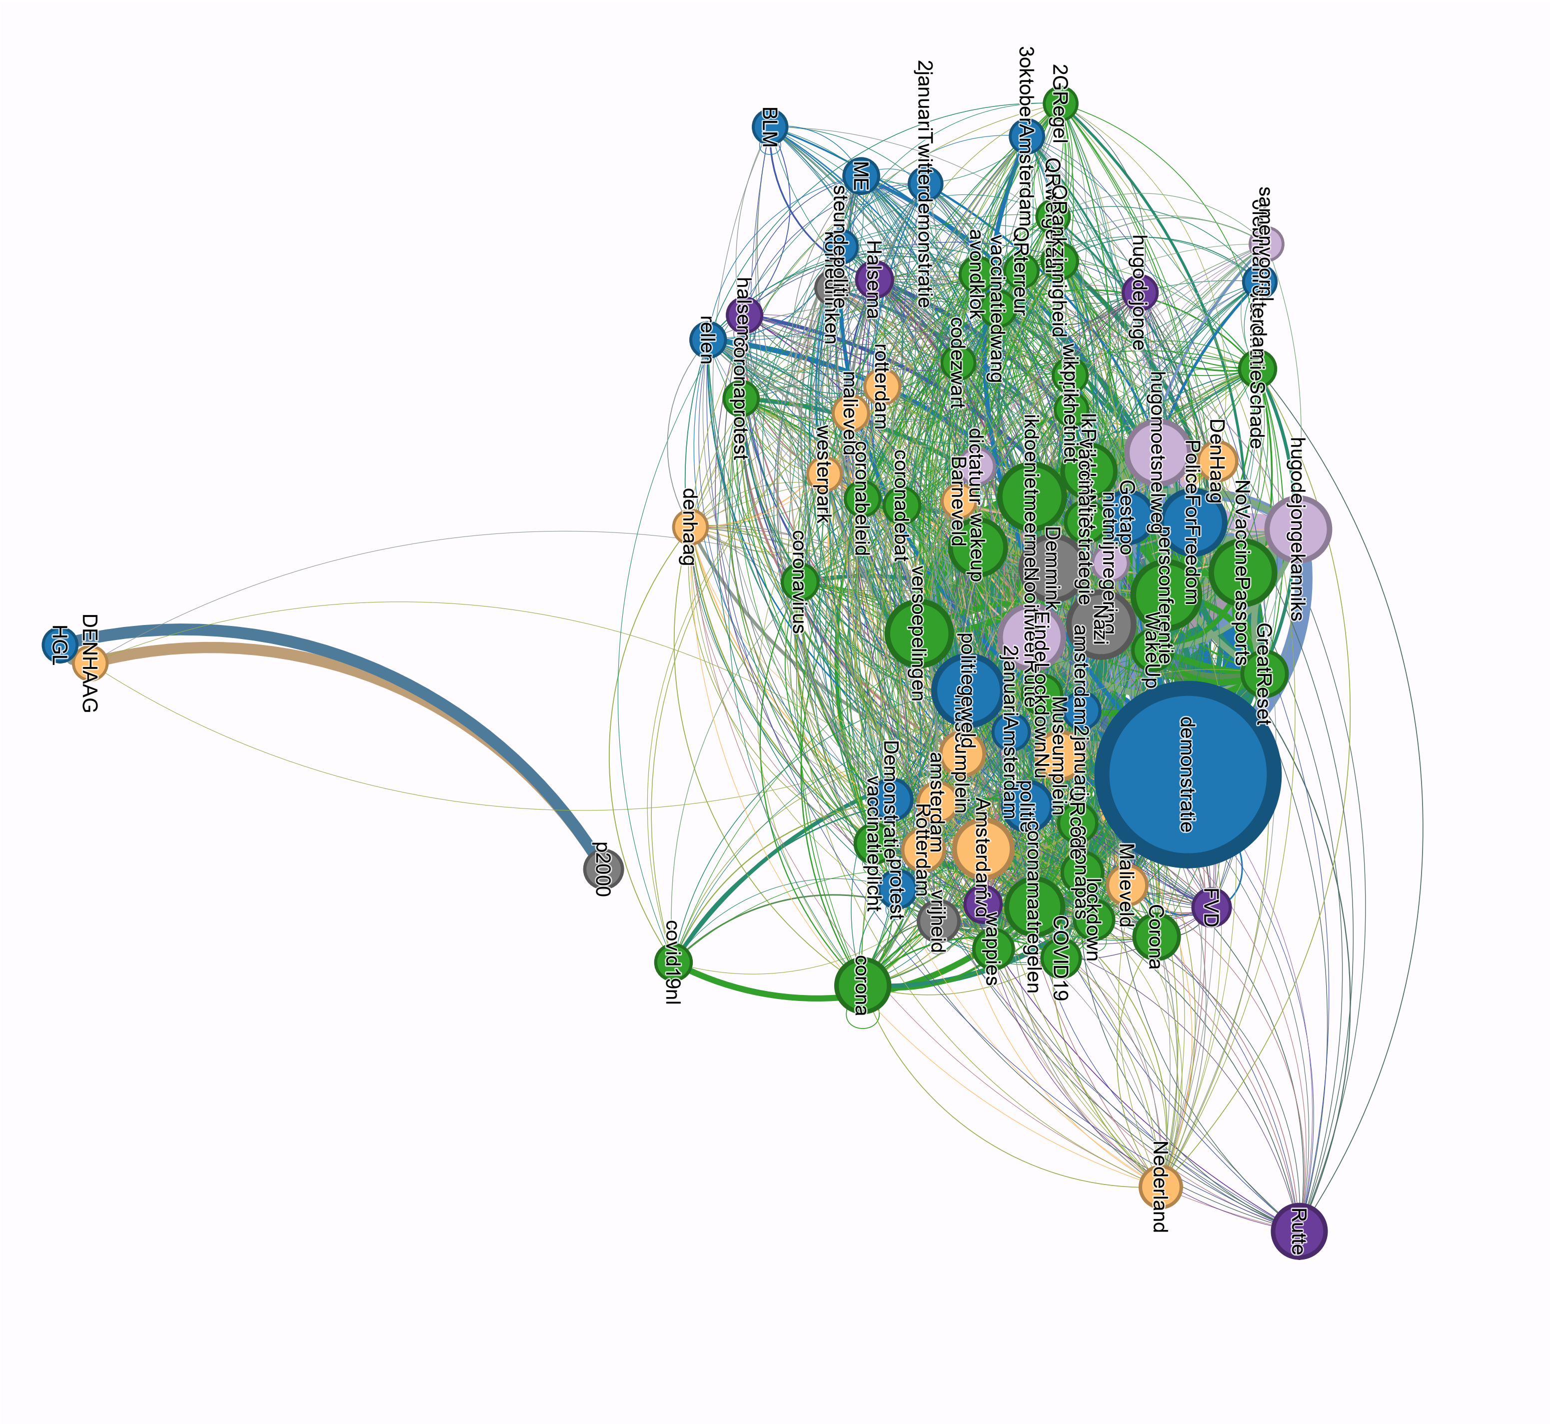

FIGURE A.10: Active social network of hashtags between 1 feb 2021 and 31 jan 2022. Weighted degree >300. 82 nodes (0.34%) (Scaled edges visible, scaled X2.0, prevent overlap)

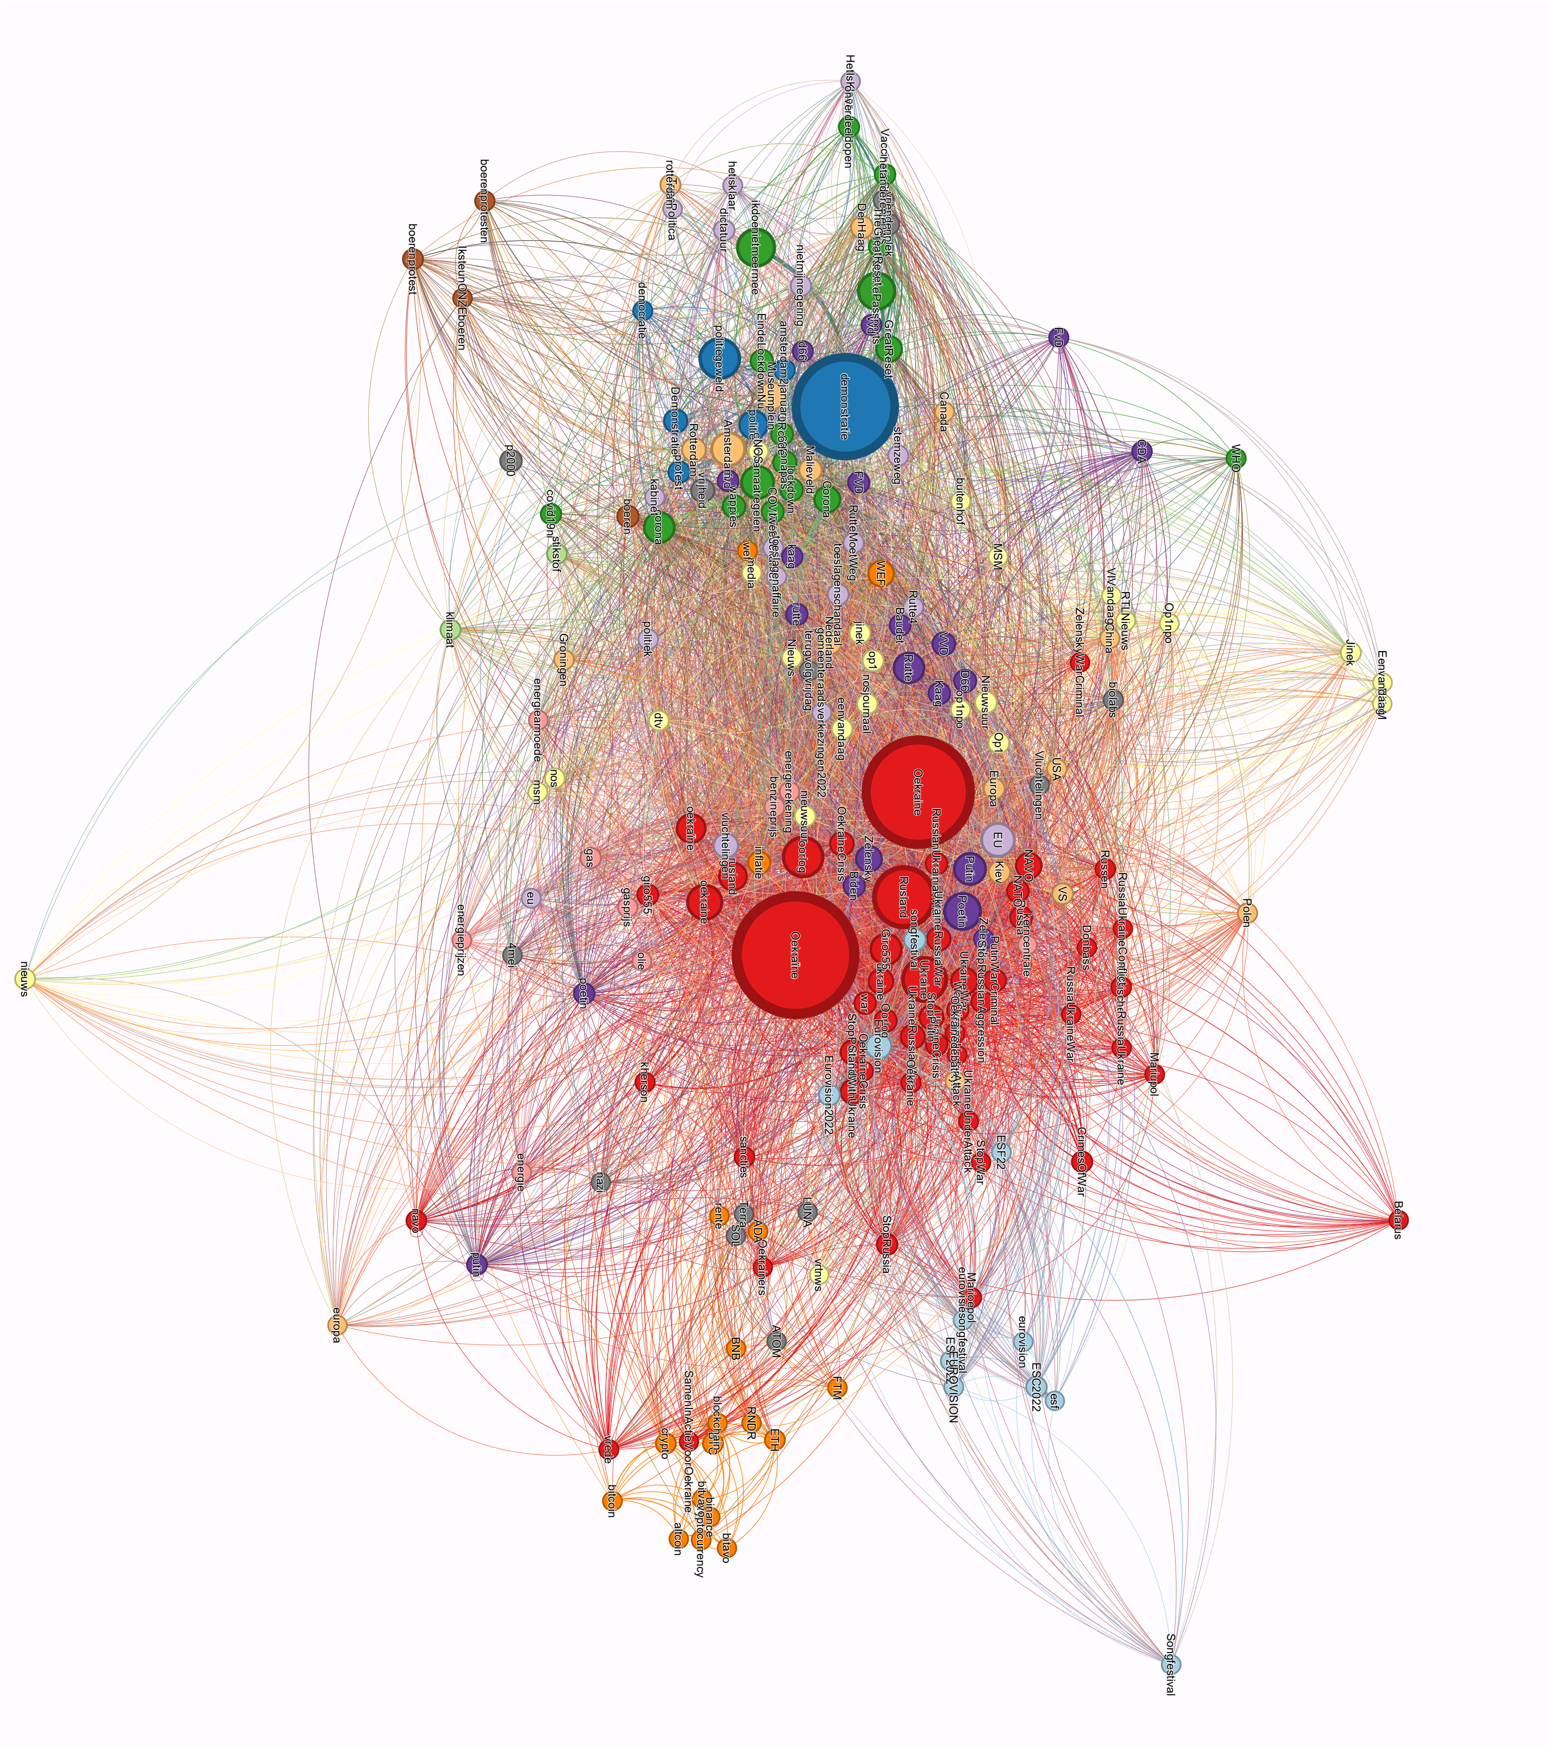

FIGURE A.11: Active social network of hashtags between 1 feb 2022 and 31 jan 2023. Weighted degree >300. 284 nodes (1,2%) (Scaled edges visible, scaled X2.0, prevent overlap)

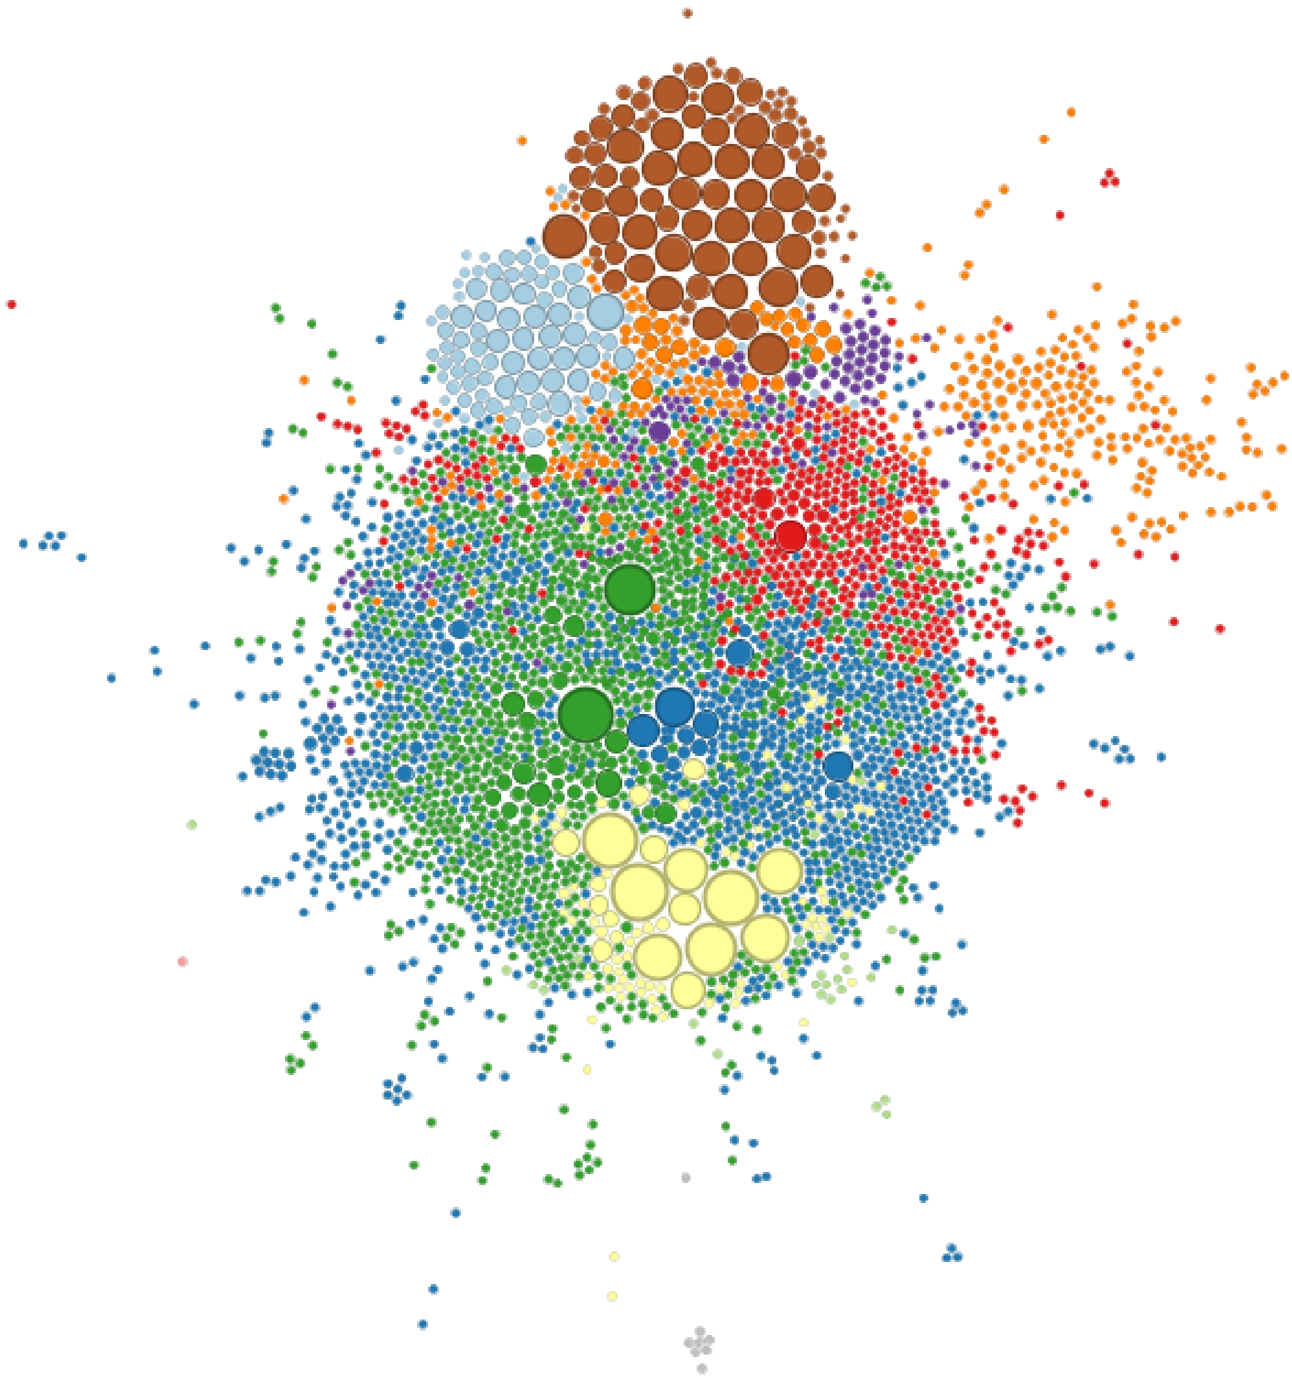

FIGURE A.12: Social network of users classified in modularity classes based on mentions on Twitter. Filter(s): $W \geq 2$.

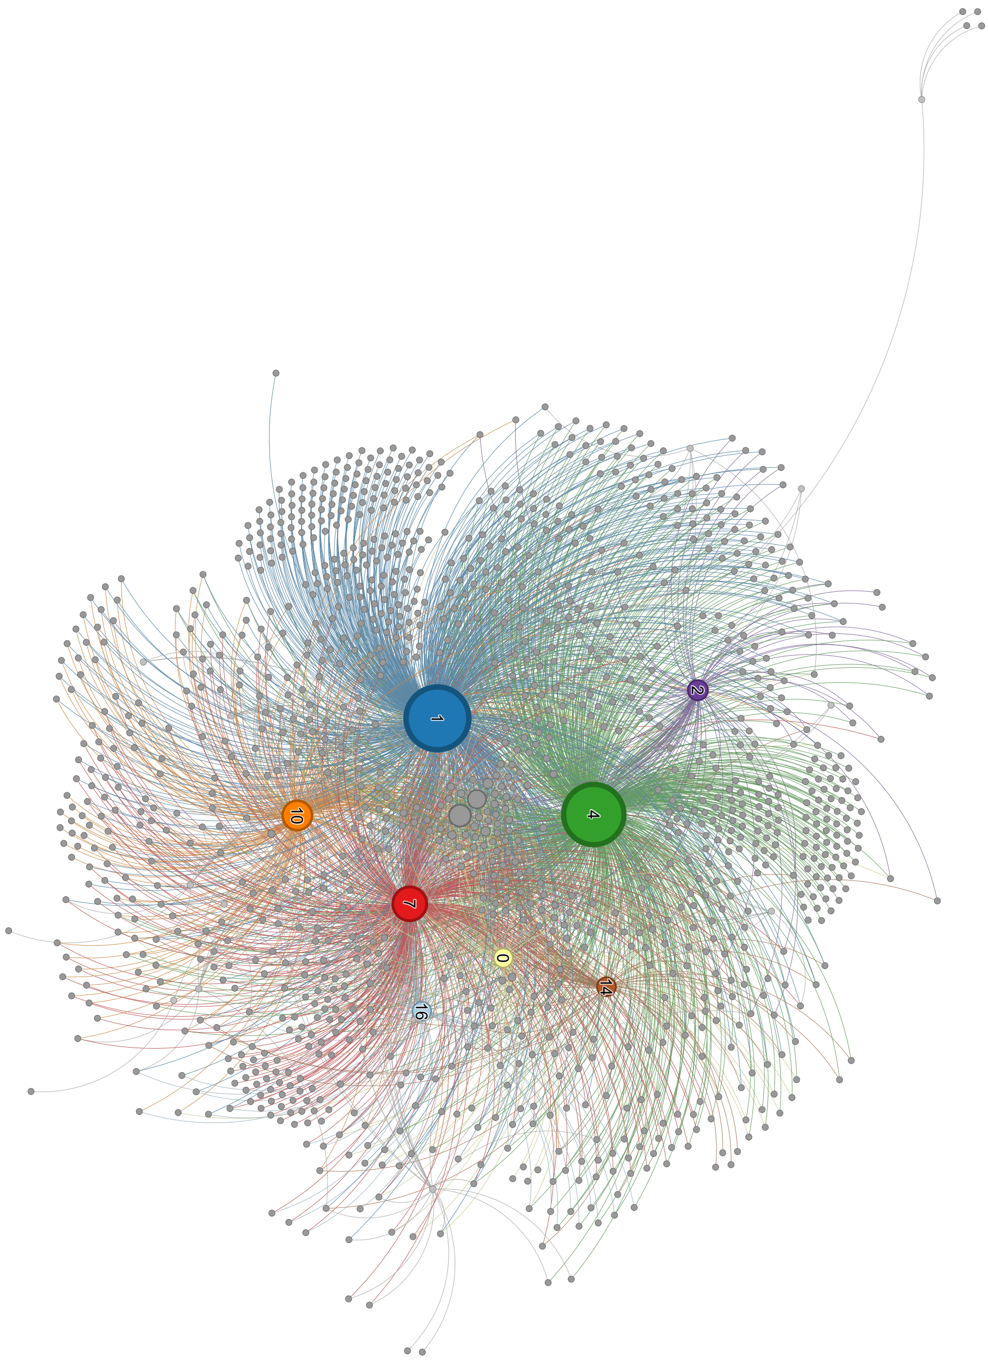

FIGURE A.13: Social network of hashtags used by the defined modality classes with filter EW >2 . 1309 nodes and 2968 edges. All edges rescaled to 1 for visibility purposes. Network scaled X100.

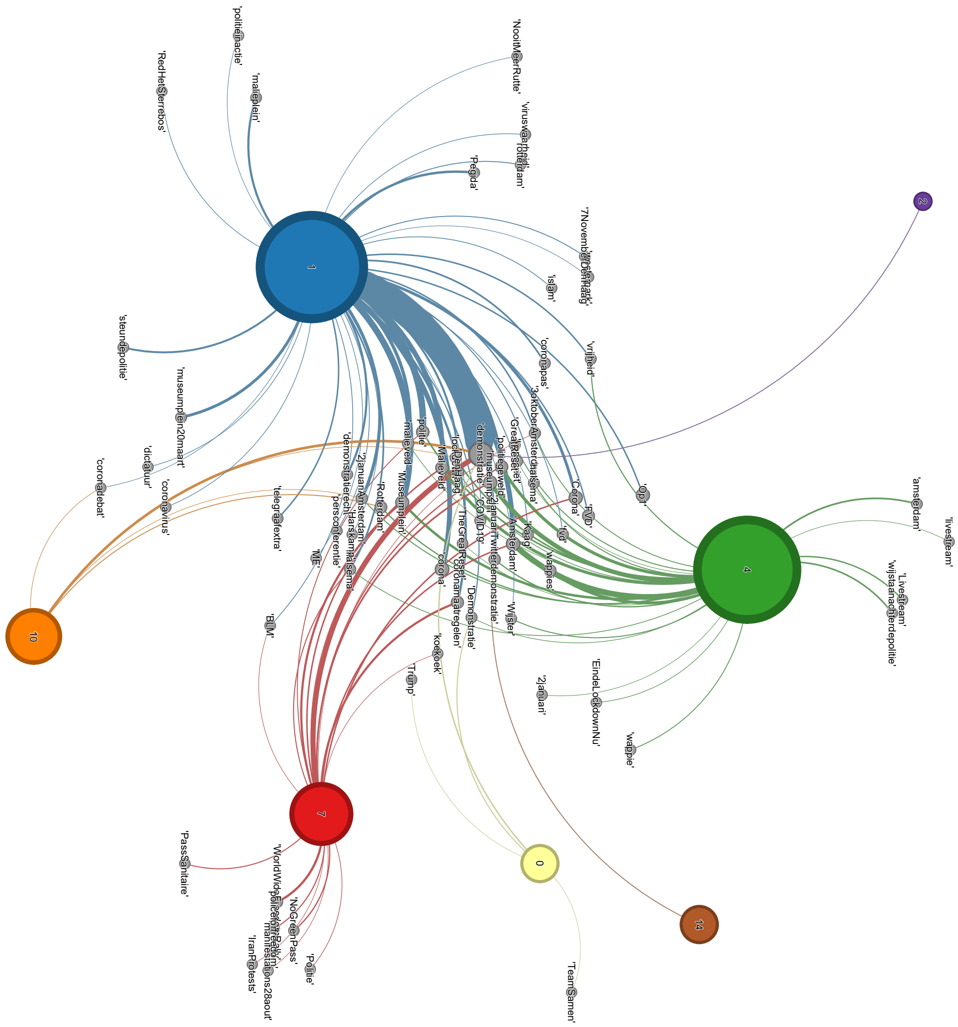

FIGURE A.14: Social network of hashtags used by the defined modality classes between 1 feb 2021 and 31 jan 2022. Edge weight filter of >5. 77 nodes (1,84%), 109 edges (1,86), (Scaled edges visible, scaled X50, prevent overlap)

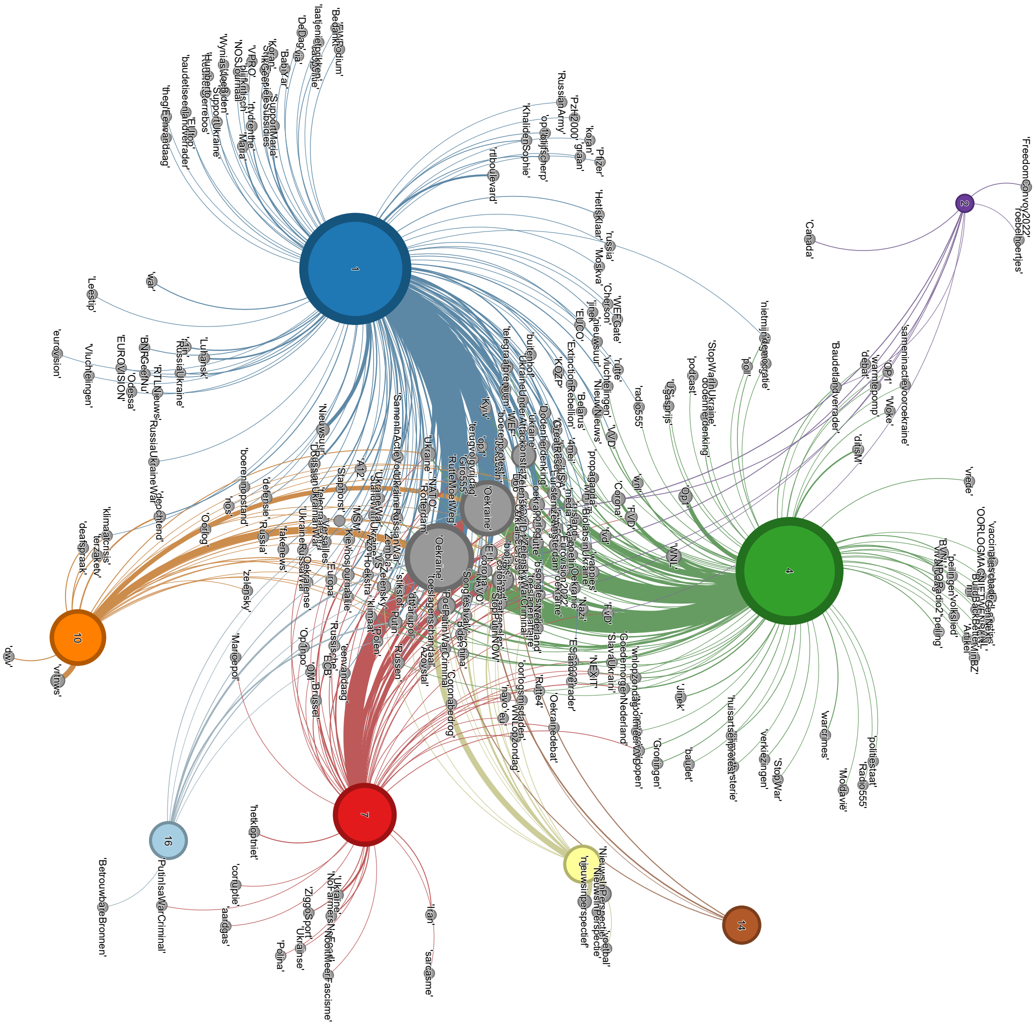

FIGURE A.15: Social network of hashtags used by the defined modality classes between 1 Feb 2022 and 31 Jan 2023. Edge weight filter of >5. 277 nodes (6.63%), 433 edges (7.4%), scaled edges visible, scaled X50

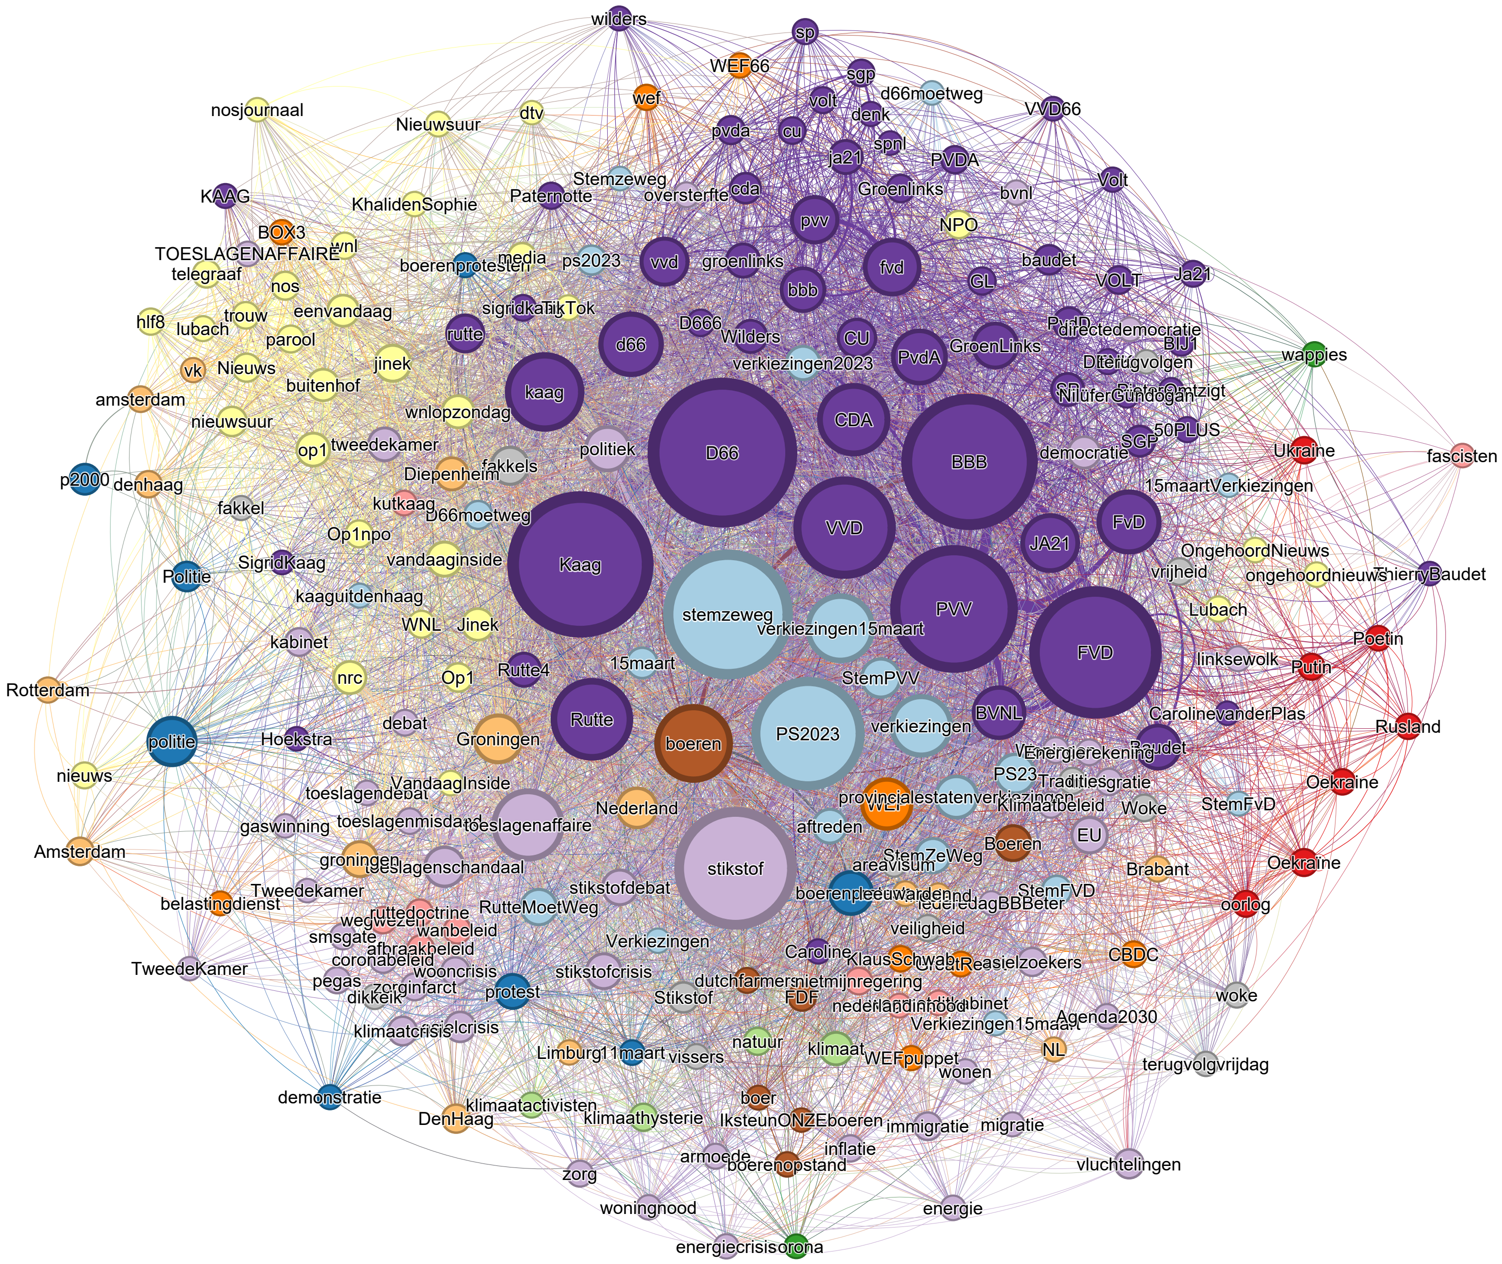

FIGURE A.16: Social network of hashtags used between February 12 and March 10 2023. Position based on only the edges that are visible. Filter(s): $WD \geq 300$. Remaining nodes/edges: 233N(1.1%). Network scaled X300. Coloring used as defined in 4.27

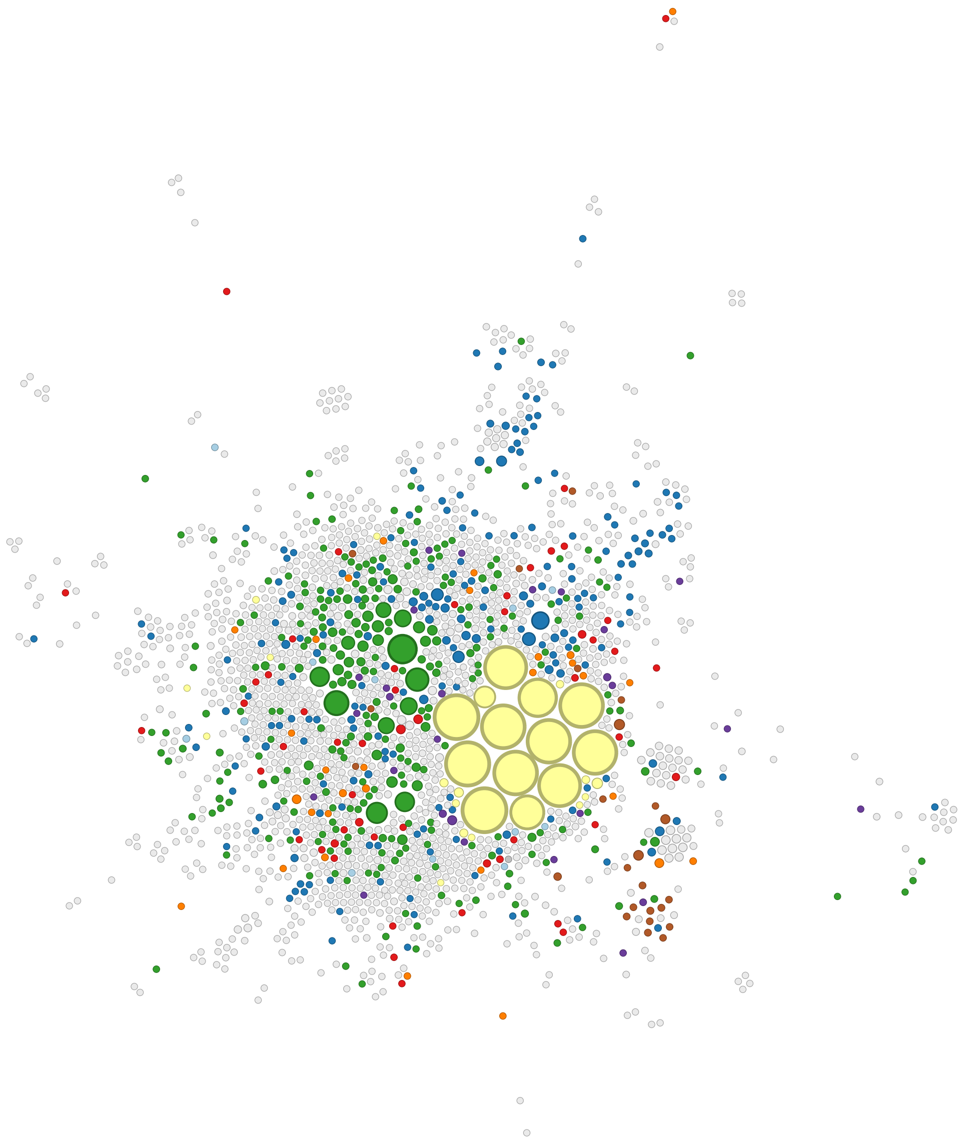

FIGURE A.17: Social network with the users of the earlier identified classes in which the position is based on the mentions in the evaluation data set. Filter(s): $EW \geq 2$. Legend visible on the right.

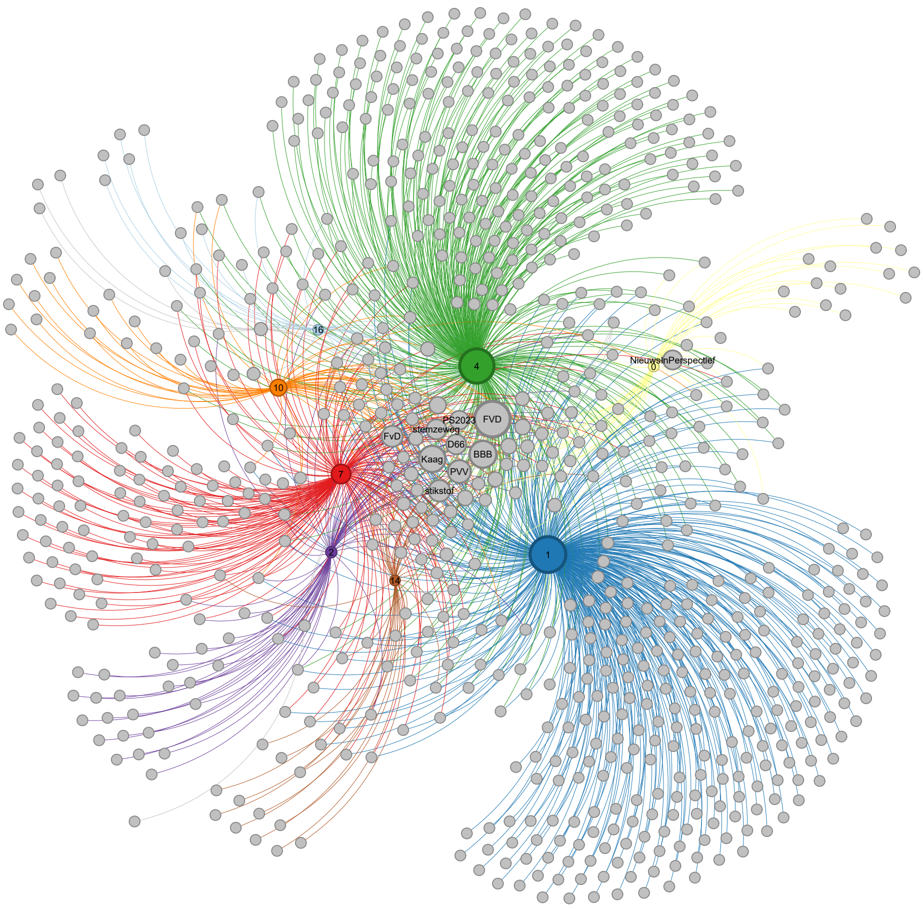

FIGURE A.18: Social network of hashtags used by the defined modularity classes between 12 Feb 2023 and 10 Mar 2023. Filter(s): $EW \geq 2$. Remaining nodes/edges: 814N(25.44%) - 1,249E(28.76%). Edges rescaled to min 0.1 and max 1 for visibility purposes. Network scaled X1000.

Appendix B

Appendix - Table representing capital letter hashtag usage differences

Hashtag	EW (Jinek)	EW (jinek)	Capitalization similarity
Corona	5	2	true
corona	2	6	true
D66	2	4	
d66		4	true
Demonstratie		1	
demonstratie		10	true
Eenvandaag	18		true
eenvandaag	1	11	true
FVD	2	1	true
fvd		2	true
Giro555	8	7	true
giro555		1	true
Kaag	10	1	true
kaag	2	10	true
Nieuwsuur	28	10	true
nieuwsuur		27	true
NOS		2	
nos		6	true
Oekraine	54	38	true
oekraine	1	1	
Oorlog	6		true
oorlog	3	9	true
Op1	19	16	true
op1	6	55	true
Op1npo	6	11	
op1npo	7	8	true
Poetin	9	16	
poetin	1	2	true
Putin	10	6	true
putin		1	true
Rusland	18	7	true
rusland		1	true
Rutte	18	12	true
rutte		9	true
Vluchtelingen	9		true
vluchtelingen	1	1	
VVD	11	5	true
vvd	2	4	true
WEF		1	
wef	1	1	
EU	2	5	
eu		1	true
			true percentage = 32/42 (76%)

TABLE B.1: Edge weights displaying relationship strengths between hashtags that were present with both a capital and a non capital letter. Last column displaying if (non)capitalized hashtag was connected more often to (non)capitalized Jinek.

Appendix C

Appendix - Table top 10 hashtags per modularity class in 2021 and 2022.

MC.	Hashtag 1				Hashtag 2				Hashtag 3			
	2021		2022		2021		2022		2021		2022	
	#	A	#	A	#	A	#	A	#	A	#	A
1	demonstratie	246	Oekraine	738	politie	57	Oekraine	612	Museumplein	55	Rusland	125
4	Amsterdam	50	Oekraïne	536	museumplein	38	Oekraine	438	politiegeweld	27	VWNL	151
7	demonstratie	45	Oekraïne	362	WordWideFreedomRally	20	Oekraine	308	coronamaatregelen	19	Poetin	112
10	demonstratie	23	Oekraïne	179	Rotterdam	8	vrtnws	134	2januar iAmsterdam	5	Oekraine	111
2	demonstratie	8	Oekraine	15	3oktoberAmsterdam	4	Oekraïne	13	politiegeweld	5	FreedomConvoy2022	11
0	demonstratie	12	NieuwsInPerspectief	222	Demonstratie	9	Oekraïne	144	TeamSamen-	5	nieuwsinperspectief	117

MC	Hashtag 4				Hashtag 5				Hashtag 6			
	2021		2022		2021		2022		2021		2022	
	#	A	#	A	#	A	#	A	#	A	#	A
1	museumplein	52	Ukraine	101	corona	43	Poetin	90	Amsterdam	43	Giro555	85
4	Malieveld	26	demonstratie	143	wappies	26	Rusland	116	Museumplein	22	Op1	100
7	Museumplein	19	Rusland	98	museumplein	15	EU	57	Malieveld	13	Ukraine	44
10	coronadebat	5	deochtend	28	lockdown	5	EU	18	KoffieDrinken	4	terzaketv	17
2	politiegeweld	2	Canada	11	2januariTwitterdemonstratie	2	Corona	11			Rusland	8
0	Trump	5	Oekraine	44	museumplein	4	NieuwsInPersp...	31	Wijster	3	Rusland	26

MC	Hashtag 7				Hashtag 8				Hashtag 9			
	2021		2022		2021		2022		2021		2022	
	#	A	#	A	#	A	#	A	#	A	#	A
1	Malieveld	42	Nieuwsuur	84	politiegeweld	39	EU	71	malieveld	29	oekraine	70
4	amsterdam	15	EU	71	lockdown	14	oekraine	66	politie	14	Corona	58
7	Amsterdam	13	NAVO	35	Corona	13	Zelensky	27	NoGreenPass	11	nosjournaal	27
10	Museumplein	3	deafspraak	16	3oktoberAmsterdam	3	dwv	16	coronapas	3	Zelensky	15
2			Baudetlandverrader	5			FvD	5			Poetin	5
0	Miranda	3	Poetin	12	vluchtelingen	3	Zelenskyy	10	Kerstmarkt	2	Amsterdam	9

MC	Hashtag 10											
	2021		2022		2021		2022		2021		2022	
	#	A	#	A	#	A	#	A	#	A	#	A
1	2januariAmsterdam	28	WEF	60								
4	wijstaanachterdepol..	14	Poetin	46								
7	politiegeweld	11	Rutte	25								
10	coronamaatregelen	3	klimaatcrisis	14								
2			roebelhoertjes	5								
0	MuseumPlein	2	voetbal	9								

TABLE C.1: Top 10 Hashtags per Modularity Class (MC) divided over 2021 and 2022. Values that are left out have not been used more than two times.

Appendix D

Appendix - Queries used for tweet retrieval.

D.1 Exploring data set

terms = [’demonstratie’]

D.2 Evaluation data set

terms = [’demonstratie’, ’protest’, ’rellen’, ’betogers’, ’demonstranten’, ’onrust’, ’sociale onrust’, ’gevaar’, ’bezetting’, ’op straat’, ’onvrede’, ’oproer’, ’politie’, ’opstoot’, ’hooligans’, ’relschoppers’, ’actievoerders’, ’actievoerders opgepakt’, ’Extinction Rebellion’, ’klimaatactivisten’, ’kwetsbare natuurgebieden’, ’Hannah Prins’, ’fossiele subsidies’, ’FDF’, ’Farmers Defence Force’, ’boerenstandpunten’, ’boerenactiegroepen’, ’Mark van den Oever’, ’toeslagenaffaire’, ’gascrisis’, ’verzakte huizen’, ’waterschade’, ’11 maart’, ’de grootste demo aller tijden’, ’Uitkopen boeren’, ’stikstofbemiddelaar’, ’Johan Remkes’, ’gedwongen uitkoop’, ’stikstof’, ’PVV’, ’Geert Wilders’, ’Groep Van Haga’, ’Wybren van Haga’, ’demonstreren’, ’actievoeren’, ’rellen’, ’11 maart demo’, ’Milieu activisten’, ’Boeren activisten’, ’Politieke partijen’, ’Stikstof’, ’Toeslagenaffaire’, ’Groningengas’, ’BBB’, ’BoerBurgerBeweging’, ’plattelandbewoners’, ’Caroline van der Plas’, ’FvD’, ’Thierry Baudet’, ’Gideon van Meijeren’, ’Pepijn van Houwelingen’, ’Tjeerd de Groot’, ’Jan Paternotte’, ’Kaag’, ’Wierd Duk’, ’PAS-melders’, ’big agro’, ’Immigratie’, ’vluchtelingen’, ’Fakkel’, ’tractor’, ’trekker’, ’vlag’,]