Exploring Protest-Related Social Network Dynamics: Combining the Power of Big-Data with Agent-Based Simulation

1st Laurens H.F. Müter*
Netherlands Police
Utrecht University, Nationaal Politielab AI
Utrecht, the Netherlands
l.h.f.muter@uu.nl
*Corresponding author

2rd Christof van Nimwegen
Utrecht University, Nationaal Politielab AI
Utrecht, the Netherlands

3rd Remco C. Veltkamp
Utrecht University, Nationaal Politielab AI
Utrecht, the Netherlands

***Abstract*—While social media usage has taken a prominent role in large social movements, societal constraints are catching up with privacy and transparency in preparation and coordination. To cope with this problem, we propose a method combining agent-based simulation and big data analysis to gain insight into upcoming protests while ensuring a well-weighted and clean analysis process. This simulation method aims to strengthen the information position of governmental officials in advance of a large-scale protest to allocate resources better and take suitable measures to ensure public safety while reducing the preliminary privacy invasive interventions. The proposed method is tested on a real-world case study where posts from Black Lives Matter protests are used to simulate social interaction in advance. Results show that behavioural constructs such as the spillover effect can be predicted based on previous data, contributing to gaining information from a wider perspective.**

***Index Terms*—Social Network Dynamics, Agent-Based Simulation, Social Media**

I. Introduction

The use of social media has risen over the last ten years [1]. Where in the early years only a handful of users marked the landscape of social media, nowadays, platforms such as X (former Twitter) serve an important social role, even overshadowing some newspapers [2].

With interest in social media, X can also play an important role in matters related to large-scale social movements and protest coordination [3]. For this reason, X is becoming important for governmental organisations to gain insight into what is happening before, during and after a demonstration [4]. However, reading and analysing messages on a bulk scale pose ethical concerns regarding privacy and transparency and are unsuited as default practices [5], [6].

However, in specific cases, such as engaging in the available information from these networks, it remains crucial to strengthen the information positions of the police when needed. Thus, models are trained for these specific cases on historical data. One example of real-world social network data is provided by [7] as presented in figure 1. In this study, real-world X posts are used to map a social network graph which can be analysed using state-of-the-art methods such as classification, clustering and topic detection. Although the availability of such data is limited, due to data storage restrictions and conflicts between the usage and original purpose of the data [8].

To solve the problems of availability and usefulness of the data, we propose a method where little data real-world data is combined with agent-based simulation techniques. This way, the amount of data is reduced, whereas officials still have enough information to gain insight into what is happening and what the near future might look like. Although insights gained from real-world data and protests are used to construct the simulation, no additional data is required to run different scenarios. This provides the opportunity for local officials to prepare for a large-scale protest. In contrast, it allows researchers to model complex social interactions within a social network.

II. Related work

Multiple social factors should be considered in understanding the social dynamics of a large-scale protest. Two of these effects, spill-over effect and polarisation, can be measured on social media platforms. The next sections describe these in more detail.

A. Spill over effect

Within the first stages of the process that eventually leads to a large protest, X is often used to gain attention towards a specific topic [9]. When influential users and local news agencies pick up the discussed topics, a quick expanding phase can erupt, leading to many engaging in this topic [10]. By definition of social structures within a social media network,

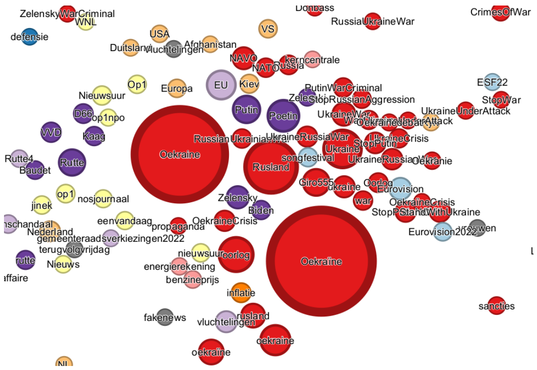

Fig. 1. A social network representation of topic (in the form of hashtags) from real-world protest-related X-posts, using ForceAtlas2 [7].

the spillover effect is often one of the key indicators that a topic will emerge under a large crowd within a short period [11]. The spillover effect happens when a user starts to post about a specific topic, this topic is then picked up by directly connected users, who also start to post messages about the topic, which are picked up by their readers and so on. This effect is visible within a network as an abrupt change in the predominantly discussed subjects.

B. Polarisation

Another effect is shown in the polarization of groups, which is a process of increasing the stance between opposite views on a topic [12]. For example, climate change is denied by some, whereas others are strong supporters. Climate is a topic that influences everyone (since global warming also affects people who do not contribute to its expansion) and can only be tackled if everyone wants to cooperate. This makes it an easy candidate to cause polarisation since both critics and defenders will strengthen their point of view within this discussion, leading to a harder tone, less room for nuance and a strong mentality of "you are either with us or against us" [13]. This results in people with a more nuanced opinion on the matter being driven to either the critics or defenders. Within a social network, a strong polarization is noticeable when topics related to a more nuanced view disappear in favour of topics that more strictly point towards the critics or defenders.

C. Representation on X

Since the spillover effect and polarisation happen within large social networks, these effects can be seen on social media platforms. Moreover, other effect such as how people react in the after-match of large events is also represented on social media [14].

Human interaction happens on a different level in comparison to face-to-face, therefor social media networks are also defined as digital humanities communities [15]. Which specifically refers to digital interaction between actors within the network. Since the data from these networks is available (often for a fee), there are numerous possibilities to analyse the user interactions and to look for patterns beyond the discussed spill-over and polarisation effects.

One of these possibilities is mapping the discussed topics in large topic networks, where factors such as information flow, density, clusters, modularity, and isolates can tell something about which information will reach users and how the sub-networks will interact when for example new topics are introduced [16]. Other methods include the usage of metadata such as followers, likes and retweets to gain information about the degree of information spread posed by specific persons [17]

Finally, a network of users with topics related to X can be represented by a network graph, such as Euler diagrams and treemaps [18]. There are also 3D visualisations available [19] for instance to display the reflection of spatiotemporal mobility on X.

D. Agent Based Modeling

Agent-based modelling (ABM) has proven highly effective for analyzing social media data due to its ability to simulate individual behaviours within complex social systems. Platforms such as X (formerly Twitter) offer a unique, large-scale dataset since messages posted there are publicly available, enabling unprecedented quantitative analysis of social behaviours and dynamics [20].

Originally, ABM techniques were developed for applications such as epidemiology, where they effectively simulate virus spread across a population [21]. A commonly used framework in this context is the SIR model, which captures the influence of direct neighbours in a network on an individual's state [22]. Similarly, social media platforms are frequently studied for phenomena like viral marketing and rumour propagation, given their rapid and wide-reaching dissemination patterns [23]. Rumour-spreading models on social media often employ the Agent-Based Social Simulation (ABSS) approach, which, unlike traditional models, accounts for the behaviour of recovered users who no longer influence their neighbours, reflecting the nature of opinion and behaviour "decay" in networks [24].

By fine-tuning simulation parameters, ABM can approximate real-world information diffusion, linking physical and online social networks. This approach reveals geographic patterns in social media activity, where certain cities may exhibit higher engagement with specific events, reflecting localized social dynamics [25]. Within microblogging platforms, agents' behaviours can be modelled using statistical techniques like Markov Chain Monte Carlo, where transition probabilities are empirically derived from user data. Here, agents can either read or post, with followers receiving messages that may propagate further depending on their sentiment and engagement [26]. To validate such models, researchers commonly employ metrics like Root Mean Square Error (RMSE), with findings indicating that removing either core egocentric users or highly engaged accounts drastically alters message volume, highlighting the influence of key network nodes. Interestingly, the impact of removing the top 100 most engaged users surpasses that of removing the seed node alone.

Repetition is also a key factor in information diffusion on social platforms. Research indicates that sending the same message multiple times can boost message reach, though a threshold exists; repeated messages beyond six iterations risk diminishing returns as users may start to feel annoyed [27]. Retweeting behaviour, explored through the lens of reward mechanisms, shows that users are likelier to share content when they can comment or quote it, reinforcing information dissemination through interactive features [28].

Agent-based models have been used to study polarization, exploring how network topology and contagion probabilities affect opinion dynamics. Using tools like NetLogo and Matlab, studies demonstrate that network connectivity influences the speed and persistence of information spread, potentially leading to either the persistence or resolution of polarizing topics [29].

Several tools facilitate ABM in social contexts, each tool offers unique advantages based on the specific simulation requirements, with some providing ease of use and visualization, while others support high customization or scalability. Tools like HashKat and Mason may offer advantages for complex social network analyses, while NetLogo and Madkit balance accessibility and flexibility.

1) Cormas: [30]: Cormas (Common-Pool Resources and Multi-Agent Systems) is designed for natural and social resource management simulations, emphasizing interactions among agents over resources. It specialises in resource-based modelling, offers visualization tools for interaction with stakeholders and is accessible to domain experts. Conversely, it has limited flexibility for applications outside resource management, making it less useful for large-scale or complex network simulations.

2) Madkit: [31]: A Java-based framework, Madkit (Multi-Agent Development Kit) is highly modular and supports multi-agent systems with extensive libraries for agent communication. It is highly customizable and good for developing complex, distributed agent systems suitable for both research and industrial applications. Using the framework requires a steep learning curve and the Java dependency may limit accessibility for non-technical users. There is also limited built-in support for social network-specific simulations.

3) Mason: [32]: A flexible, general-purpose simulation toolkit in Java, Mason supports complex, large-scale simulations, including spatial and agent-based models. It has a high performance for large-scale simulations, is highly flexible and extendable and provides detailed control over simulation design. Like Madkit, it requires Java programming knowledge and lacks built-in social network analysis tools, which must be implemented separately.

4) NetLogo: [33]: NetLogo is widely used for ABM and provides an accessible interface for creating simulations with its agent-based language. It is user-friendly and suitable for beginners and experts, with strong visualization capabilities and extensive libraries for social and biological simulations. Since the performance shows limitations with large-scale or complex models it is less suited for simulations requiring high computational efficiency.

5) Swarm: [34]: One of the earliest ABM frameworks, Swarm was initially developed to simulate biological and social phenomena and provides a robust foundation for complex modelling. It is well-suited for creating detailed simulations of individual and group behaviours and strong event-based structure. Due to its long existence, there are limited updates and support, it has an outdated interface and performance can be slower with complex or large-scale models.

6) HashKat: [35]: This tool focuses specifically on social network modelling, with a strong emphasis on modelling information diffusion and opinion dynamics. It is Tailored for social network simulations, incorporates statistical measures of information diffusion and is suitable for analyzing network-specific behaviours. Due to this specific usage, it has limited generalizability for non-social network applications and a narrower scope compared to more versatile ABM platforms.

7) Soil: [36]: Soil is an ABM platform aimed at facilitating modular simulation setups with a focus on environmental and social dynamics. It is designed in modules and supports a wide range of environmental and social applications with a relatively easy setup. Less support for high-complexity network modelling, and focus on modularity over performance may limit its scalability for large models.

While various platforms support ABSS, there is a recognized need for domain-specific models. General-purpose ABSS frameworks provide a robust foundation for social simulations, but tailored models can yield more precise insights into specific social phenomena [36].

III. Method

The Social Network Protest Simulation Tool (SNPS) is created as an in-browser simulation tool to simulate the dynamics of social media users in the phases leading up to a protest. To calculate different scenarios, a set of parameters is defined to tune the network behaviour.

A. Design and Development

When developing the SNPS tool, the choice of JavaScript libraries and frameworks plays an important role in the tool's performance, compatibility, and user experience. After consideration of various options, vanilla javascript in combination with the material design was selected to suit the requirements of this project, as being flexible, easy to use and lightweight.

1) Frameworks Considered for Tool Development: In the initial phase, popular JavaScript frameworks such as React [37], Vue [38], and Angular [39] were evaluated. These frameworks are known for their ability to speed up development processes due to their component-based architecture, which facilitates modular and reusable code. Additionally, they are supported by large communities, ensuring that developers have access to extensive resources and updates [40]. However, despite these advantages, the decision was made to develop the simulation tool using vanilla JavaScript. This choice was primarily driven by the need to ensure compatibility with older browsers, which may not fully support the latest JavaScript frameworks. Vanilla JavaScript provides a lightweight and flexible approach, making it easier to maintain broad compatibility without relying on the additional overhead introduced by modern frameworks.

2) Design and Style: For the user interface, Material Design was selected as the design framework. Material Design, developed by Google, offers a clean, intuitive, and familiar user interface that is widely recognized and appreciated by the target user group, including law enforcement professionals [41]. The decision to use Material Design was influenced by its default Google style, which aligns well with the expectations of users and enhances usability through a consistent and modern look and feel.

3) Graph Visualization Libraries: A critical aspect of the simulation tool is the visualization of networks, which requires a robust and flexible graph visualization library. Several libraries were considered:

Sigma.js: Sigma.js is an open-source library licensed under the MIT License. It was initially considered due to its strong community support flexibility. Sigma.js is highly performant, particularly for large-scale graphs, as it utilizes WebGL for rendering, which accelerates performance. It is also stable, making it a viable option for future updates and modifications.

Vis.js: Vis.js has an extensive feature set and is well-documented, making it easy to learn and use. It is used in many projects and still has strong community support, and its performance remains solid for smaller graphs. Additionally, Vis.js is also licensed under MIT, making it a flexible and legally safe choice for development.

D3.js: D3.js is a powerful, commercial-grade tool known for its wide range of features and exceptional documentation. However, it was not selected since it is not inherently designed for in-browser engines, which could complicate integration and affect performance. Although D3.js offers fast rendering and extensive capabilities, this was not needed for the project.

G6: G6 is an MIT-licensed library with many easy-to-use features. While G6 has a reasonable community and offers good performance, the language barrier poses a challenge, since part of its documentation is written in Chinese.

Vis.js: was ultimately chosen for this project since it is lightweight and easy to use, it has a successful track record within many other projects, is well documented and has strong community support.

B. Topic Tuning

During initialisation, each user is provided with a topics array, where each topic is represented by a value between zero and one, as illustrated in equation 1. The values for each topic are generated randomly, but every user's topics list (T_u) sums up to one, as seen in equation 2. This restriction is necessary to observe how the ratio of topics evolves. The number of topics can be adjusted via the graphical user interface by modifying the *Number of Topics* parameter.

Since users interact with each other by reading messages of directly related users, the topics are tuned within every simulation step to reflect the influence users have on each other's topics. An influence factor determines the influence between users, a value between zero and one, which is randomly drawn from a normal distribution as defined in equation 4. The mean (μ) and standard deviation (σ) of this influence distribution can be configured within the graphical user interface using the *User Influence Distribution Mean* and *User Influence Distribution Std* parameter respectively. By default, the mean is set to .2, and the standard deviation to .1 resulting in relatively small influences between users in the short run, whereas effects are visible after 100+ steps under default conditions. Since the influence parameter is restricted to the range of [0, 1], any values exceeding one are capped at 1, and any values below zero are set to 0.

For the calculation of the new values of the topics from a given user, equation 5 is used. Here, the value of a topic t'_u is based on the user's influence factor f_u, the current value of the topic t_u and the value of the same topic from another user (t_o) that is directly connected, see equation 3. The calculation consists of two parts, $(1 - f_u)t_u$ denotes the inverse of the influence factor to keep the original value of the topic, and the average of the t_u and t_o multiplied by the users' influence factor f_u denotes the second factor to determine how the user's topic will change (as shown in equation 5).

This method ensures the sum of topics remains equal to one (as noted in equation 1), while the topics can change in accordance with the user's influence factor.

$$\forall t \in T(0 \leq t < 1) \tag{1}$$

$$\sum_{t \in T_u} (t) = 1 \tag{2}$$

$$t_u \in T_u, t_o \in T_o \tag{3}$$

$$f_u \sim \mathcal{N}(\mu, \sigma^2) \text{ where } 0 \leq f \leq 1 \tag{4}$$

$$t'_u = (1 - f_u)t_u + \frac{f_u(t_u + t_o)}{2} \tag{5}$$

C. Topics Display

Every topic has a distinct colour associated with it, depending on the list index of the topic. Every node within the displayed social network graph is represented by one colour for every user. This colour presents the two topics with the highest relative value for every user.

Within every time step of the simulation, after topics are tuned, colouring is applied to give the user a single colour that represents the two most prominent topics within the users' topics list (T_{u_1} and T_{u_2}). To calculate the blended user topics' colour, the ColorJs library [42] is used. Mixing the two topics with the highest relative value is done using the CIELAB interpolation algorithm [43] as available at the ColorJs library. Only mixing the two most popular topics of a user will ensure that the global representation of topics over the entire user

Index	#HEX	Name
0	#2BC321	green
1	#2174C3	blue
2	#C32121	red
3	#C3C121	yellow
4	#C321B4	pink

TABLE I
THE COLOURS USED FOR COLOURING THE TOPICS; *Index* REPRESENTS THE LIST INDEX OF A TOPIC, *#HEX* DEPICTS THE HEXADECIMAL REPRESENTATION OF THE COLOUR, AND *Name* IS THE COLOUR NAME IN ENGLISH.

group remains clear since colours are easily distinguishable while keeping the possibility to "blend" topics.

D. Real World Data

The data is extracted from X-posts (former tweets) that centre around a Black Lives Matter demonstration in Amsterdam during the COVID-19 pandemic [44]. The posts are manually labelled where aspects such as sentiment, expressions of violent behaviour and related media are included. A dedicated labelling protocol also includes the description of the most important events. After the protest started, there were concerns regarding the measures related to social distancing as the demonstration became overcrowded. This issue gained national attention, eventually leading to societal discontent regarding the handling of the COVID-19 measures and in particular towards the mayor of Amsterdam, Femke Halsema.

The open API of X (former Twitter) was used to collect the posts, from one day before the protest to one week after. Each API call contained the same queries to maintain consistency (terms included "demonstration"). Unique posts are used, thus retweets are excluded. The set included descriptive terms such as hashtags and nouns, but the full texts are not available. To construct a network, the hashed mentions are included as well as the hashed usernames of people that posted the messages, where the hash function on the mentions and usernames is kept consistent to make comparisons possible between the two. Furthermore, the set contained a time stamp to indicate that the messages were posted. This timestamp is used to select an hour-by-hour time series of posts.

The labelling team included six people who had labelled for 18 months to obtain four distinct datasets. The labelling took place under the supervision of representatives of the Netherlands police's open-source intelligence team. A subset of posts was used to determine the inter-annotator agreement which, averaged to 0.70. The initial dataset (including retweets) contained 84,901 tweets from which 6,155 were labelled as described in a dedicated data in brief paper [44].

The Black Lives Matter (BLM) demonstration in Amsterdam is scheduled at 17:00. Around noon, the mayor of Amsterdam and the police unit commander met to determine the demonstration's location, ultimately selecting Dam Square over the larger Museum Square. Following this decision, the demonstration organizers took precautions to encourage social distancing, including marking crosses on the pavement for protesters to stand on, setting up aisles, and posting signs to remind attendees of the coronavirus measures.

By early afternoon, social media activity intensified, as announcements on Facebook helped to attract people who would later join the demonstration. This online mobilization increases Black Lives Matter-related messages across platforms between 13:00 and 18:00. The anticipated 500 attendees quickly increase to over 10,000 by late afternoon.

The key speakers' program starts at 17:00, with prominent figures taking the stage to address the crowd. At 18:15, the mayor of Amsterdam, Femke Halsema, arrived at the demonstration, wearing a Keti Koti button, a symbol of respect for the celebration of the abolition of slavery.

As the evening progressed, public sentiment shifted, increasing discontent surfacing on social media towards the protesters and local authorities. Concerns over coronavirus safety violations became a focal point of criticism, particularly at the demonstration's size and the perceived lack of adherence to social distancing protocols. Public criticism of the mayor and local government continued to build from 19:00 onwards, fueled by frustration over the authorities' handling of the event amidst the ongoing pandemic.

An additional dataset was included to study the robustness of the results. For this purpose, a protest related to "Zwarte Piet" (Black Pete), a tradition in the Netherlands that has sparked controversy, was selected. The KOZP movement (Kick Out Zwarte Piet) argues that the character perpetuates a racist stereotype due to its blackface, while supporters view it as a harmless cultural tradition. Although some incidents were reported, the main topics did not shift during this protest.

E. Evaluating SNPS

To evaluate the social network protest simulator, a real-world data set is used to determine the most prominently discussed topics and how these topics change over time. The experiment consists of multiple steps; first, the data is subdivided into hour-by-hour sets, using the "created_at" field as included in the data. Second, the data is preprocessed by extracting the nouns and hashtags. Third, the users who posted the messages are obtained and the users mentioned in the posts are extracted. fourth, the five most discussed topics (related to five distinct terms) are obtained. Fifth, each of the five topics is counted within the posts of the fifty most active users. Sixth, a simulation round is played based on the previous timestep, using the most prominent topics-counts for each user and a network structure based on the mentions from step 3. Finally, the results of each simulation step are compared with the real network at that time (thus the simulation based on the data from t is compared to the actional data from $t+1$. The steps are explained in more detail.

1) Subdivide Data: The "created_at" field was available within the dataset, which is converted into local date-time representation using the Python Pandas data frame implementation. This resulted in 18 different subsets containing data between 05:00 and 23:00 on the day of the demonstration (other subsets are excluded since they contained too few records to perform the analysis).

2) Data Preprocessing: The hashtags and nouns were already extracted from the full text using REGEX. The nouns are stemmed using the "nl_core_news_sm" pre-trained model of the Python library Spacy. Furthermore, all hashtags and nouns are transformed into lowercase.

3) Obtaining users: The users that posted the messages are provided by the dataset, which also includes the mentions from the posts. The usernames are hashed with the same hashing algorithm to make comparisons possible.

4) Most Discussed Topics: Using the Counter library of Pythen, the most discussed terms are extracted, the top terms also include tokens such as 'aar', 'i', and 'zh' which have no meaning in Dutch, thus these terms were ignored. The resulting topics include 'halsema' (the major of Amsterdam during the protest), 'demonstratie' (demonstration), 'blacklivesmatternl' (Black Lives Matters hashtag), 'anderhalvemeter' (one and a half meter, referring to a CODIV-19 measure), 'corona' (referring to COVID-19).

5) Topic Counts per User: The fifty most connected users are obtained by determining the connectedness of every user. The users with the most connections are selected for further analysis. From the selected users, all nouns and hashtags are obtained from messages that they posted. Only the terms of the previous step are counted, resulting in a tuple of five topics per user. From these counts, the topic lists T_u can be determined by taking the ratio of topic counts for each user.

6) Simulation Rounds: Using the preprocessed data within the constructed network, a simulation round is played on each timestep based on the data from the previous time step, as described by formulas of the topic tuning section.

7) Compairing Results: The most prominent topic for each user from the simulation is compared with the most prominent topics from the real data. Since the most prominent topic can be seen as a category, we used default measures such as F1, accuracy, and precession to evaluate the performances. These measures are obtained using the PyCM python library.

IV. Result

The resulting dashboard is shown in figure 3 with the parameters on the left. When the user presses the "Start" button, the simulation starts, displaying a generated network instantly. The start button is now changed into a "Stop" button which can interrupt the simulation. When the simulation is stopped, a "Reset" button appears to re-instantiate the data and clear the previous network graph.

First, a dry run was played using a randomly generated network. As seen in figure 2 the results represent a spillover effect, visible between timestep 371 and 385, where a red topic is introduced within a predominantly green topic population, which leads to a burst of red topic users. Moreover, in an early step (328), the green crowd did not pick up a purple topic.

After the dry run, a simulation with the BML dataset was conducted, see figure 4 for the resulting networks at 17:00 and 18:00, including a simulation based on the 17:00 data. In this visualisation. the edges to connect the users are included and the nodes are enlarged based on their connectedness (the more connections a user has, the larger its node is displayed). The same seed is used to draw a consistent network layout.

t_{prev} - t_{curr}	**F1 0**	**F1 1**	**Acc 0**	**Acc 1**	**Pre 0**	**Pre 1**
04:00 - 05:00	1.0000	0.0000	1.0000	1.0000	1.0000	0.0000
05:00 - 06:00	0.8000	0.8000	0.8000	0.8000	1.0000	0.6667
06:00 - 07:00	0.8000	0.5000	0.7143	0.7143	1.0000	0.3333
07:00 - 08:00	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
08:00 - 09:00	0.5000	0.7500	0.6667	0.6667	0.5000	0.7500
09:00 - 10:00	1.0000	0.0000	1.0000	1.0000	1.0000	0.0000
10:00 - 11:00	0.8571	0.0000	0.7500	0.7500	1.0000	0.0000
11:00 - 12:00	0.7500	0.0000	0.6000	0.6000	0.7500	0.0000
12:00 - 13:00	0.6154	0.5455	0.5833	0.5833	1.0000	0.3750
13:00 - 14:00	0.3333	0.0000	0.2000	0.2000	0.5000	0.0000
14:00 - 15:00	0.7273	0.0000	0.5714	0.5714	1.0000	0.0000
15:00 - 16:00	0.6000	0.0000	0.4286	0.4286	1.0000	0.0000
16:00 - 17:00	0.8696	0.4000	0.7857	0.7857	1.0000	0.2500
17:00 - 18:00	0.8750	0.6667	0.8182	0.8182	0.8750	0.6667
18:00 - 19:00	0.7692	0.5714	0.7000	0.7000	1.0000	0.4000
19:00 - 20:00	0.8571	0.8571	0.8571	0.8571	1.0000	0.7500
20:00 - 21:00	0.9091	0.6667	0.8571	0.8571	1.0000	0.5000
21:00 - 22:00	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

TABLE II

Measures to compare the topics from the real network at the current hour t_{curr} and the topics from the simulated network based on the previous hour t_{prev}. Each column depicts a measure, **F1** for the F1 score, **Acc** for Accuracy, and **Pre** for Precision. Only the two dominant topics are considered (0 for 'Halsema', and 1 for 'demonstration').

The full results of the measures are included in table II. The time that served as input for the simulation is displayed as t_{prev}, and the current timestamp as t_{curr}. For example, the measures of the first row compare a network of users and topics based on the data available at 04:00 to the actual data available at 05:00. Only two topic categories (0 and 1) are compared since the other topics were not dominant to any user within the time. The topics relate to "Halsema" (the mayor of Amsterdam at the time of the protest) and "demonstratie" (which is the Dutch word for demonstration).

As shown in table II, F1 scores range between 0 (no predictive value) to 1 (perfect predicted), these extreme values occur at the start of the start and end of the demonstration day. When most information is available, the F1 measures values between .6000 and .9091, corresponding with reasonable results. Between 13:00 and 14:00 the measures indicate that the predictive power of the simulated decreases, which corresponds with a fast pace of shifts in different events during the day of the demonstration. Towards the end of the day, the predictive power of the simulation model seems to stabilise, when shifts in topics alternate less quickly.

Other topics from the BLM dataset were selected to evaluate the robustness of the model. When substituting "burgermeester" (Dutch for mayor) for "Halsema", the F1 score improves at the 10:00-11:00 prediction and on all the predictions from 17:00 onwards. Whereas when looking at the alternations between "demonstration" and "Halsema" or "demonstration" and "mayor" the model shows near-perfect F1 scores.

The F1 scores were also near perfect when simulating the data for the KOZP protest, which could be an indication of little alternation within topics.

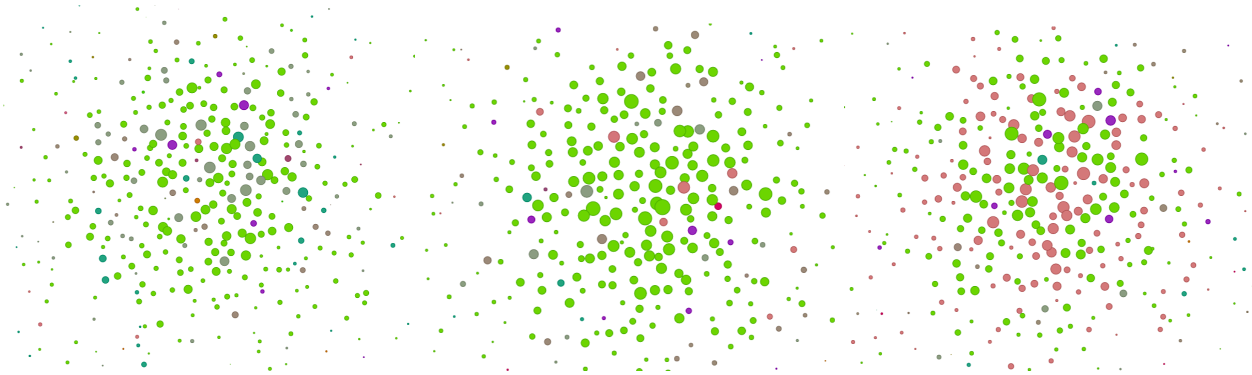

Fig. 2. Three snapshots during simulation, left, during step 328, shows green topics with some purple and dark green topics. Middle, step 371, where most users have green as their main topic. Right, at step 385, the same users have adopted Red as their main topic (which is similar to the real-world spillover effect).

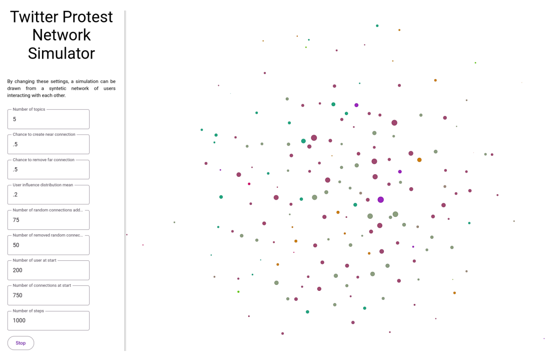

Fig. 3. A screenshot of the Protest Simulator during a run. Parameters of the current run are shown on the left.

V. Concluding remarks

The combination of vanilla JavaScript for broad compatibility, Material Design for a familiar and clean user interface, and Vis.js for its well-documented and feature-rich graph visualization capabilities provided the best balance between performance, usability, and maintainability for the development of the SNPS tool. This selection ensures that the simulation tool is not only functional and efficient but also user-friendly and accessible to a wide range of users.

VI. Discussion

The SNPS tool can help gain insights into protest-related network dynamics and information dynamics before a protest. However, several factors must be considered when interpreting these results, such as the variability in different protest scenarios. Observations carry a specific context, thus, distinct social, political, or cultural contexts could produce varying network behaviours. Therefore, while the tool provides a powerful analysis, its results should be considered in the specific conditions when the simulations were conducted.

A. Parameter Optimization

Additional runs were conducted to examine the tool's performance in various configurations. First, the number of topics per node did not significantly impact the network dynamics. Typically, only the first two topics alternated across the network. Second, the sensitivity factor did play a role in the network dynamics. With high thresholds, the network tended to remain static. Conversely, more network shifts occurred with a lower threshold. Finally, assigning higher thresholds to well-connected nodes resulted in a few users increasing their influence over the network as time progressed.

B. Incorporating Social Norms and Habits

One of the key challenges in simulating social networks is the accurate representation of human behaviour, which is often unpredictable and influenced by many factors, including cultural norms, historical contexts, and individual habits. Although the tool currently allows for the inclusion of certain heuristics to simulate expected behaviours —such as increased activity in response to an overcrowded demonstration— these are inherently simplified representations of reality. For example, the presence of the mayor of Amsterdam wearing a Keti Koti button elicited strong emotional responses such as anger and discontent within specific parts of society, leading to heightened activity in the related network clusters. While these behaviours can be simulated to some extent, the tool cannot fully capture the complexity and nuance of human reactions manual adjustments and the incorporation of domain-specific knowledge are necessary to enhance the accuracy and relevance of the simulations.

C. Potential for Broader Applications

Looking beyond the current focus on protest-related activities, the simulation tool holds potential for a wide range of various platforms and contexts. For instance, it could be adapted to analyze trends on social media platforms like X, to monitor how certain topics from different countries (in this case the death of George Floyd) might contribute to potential social unrest within the Netherlands. By identifying shifts in network dynamics or changes in user behaviour, the tool could provide early warnings of emerging public issues or social unrest. However, these applications would require further development and customization to address different domains' specific characteristics and challenges.

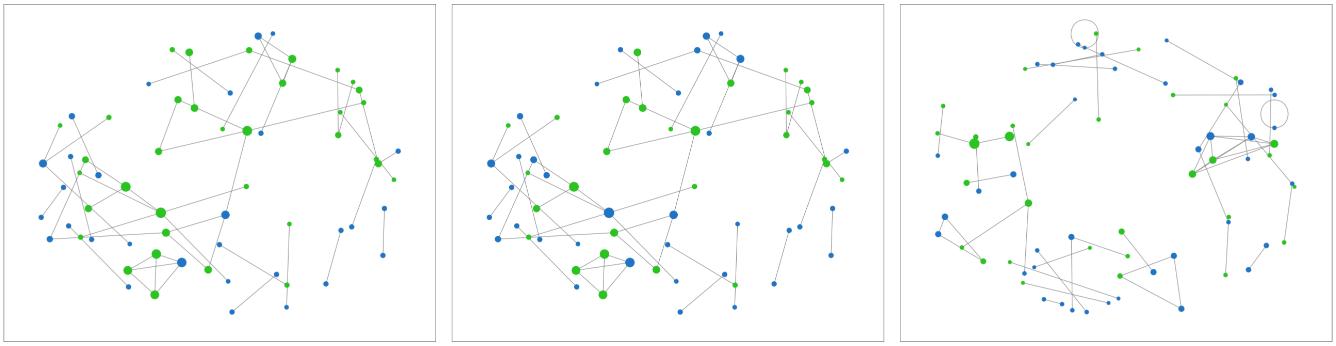

Fig. 4. Three snapshots during simulation, left, the network at time 17:00, showing green and blue topics. Middle, simulation based on the 17:00 data shows some users that have changed their most prominent topic. Right, the network at time 18:00 shows the actual state.

D. Computational Challenges

The tool's effectiveness comes with significant computational demands, due to the iterative processes involved in colour tuning and the addition of indirect nodes. This process becomes increasingly intensive as the size and complexity of the network grow. Similarly, adding indirect nodes—those connected via second-order relationships—requires extensive iterations over all connected nodes to accurately identify and integrate these secondary connections into the network. These computational challenges may limit the tool's scalability and speed, especially when dealing with large-scale networks or real-time data processing. To cope with the computational challenges of finding simple paths of the indirectly connected nodes, at depth first approach should pose a workable alternative for the current implementation. Another improvement could be made on the data representation of the network, for example by using dictionaries to speed up the sorting times. Finally, speed performances might improve by using GPU or process operations in parallel. Future enhancements could focus these optimisations on leveraging more powerful computational resources to mitigate these limitations.

VII. Future Work

Future work can expand in several key areas to enhance understanding of social media network dynamics and their real-world parallels. First, incorporating different types of networks beyond the current scope would provide a richer view of user interactions across platforms, enabling comparative analysis and revealing unique dynamics within each network. Expanding the data sources and examining their interactions would also yield insights into cross-platform behaviour and how information spreads or changes context between networks.

The multi-dimensionality of social media data offers another area for exploration. For example, the connections between individuals and the formation of online groups could deepen understanding of the underlying social structures in online communities. Additionally, examining real-world dynamics during protests and comparing these to online activities could reveal the similarity between digital and physical spaces regarding influence, organization, and response patterns.

Finally, investigating the impact of interventions within the network could be insightful. For instance, studying the effects of someone discussing a sensitive topic or the occurrence of an unexpected event could demonstrate the resilience or vulnerability of social media networks to external stimuli. Understanding these impacts could lead to models that better predict the network's response to real-world events, fostering preparedness and responsiveness in social media management and public communication.

VIII. Ethical Statement

The usage of social media data poses ethical concerns regarding consent. Therefore it is prohibited by the GDPR to use data that can be used to trace individual users or gather personal information [45]. Since the data used in this research is already published, it is anonymised and therefore suitable for the purpose of this study. Furthermore, additional measures are taken to consider storage, reproducibility and data management. For more information regarding data-related subjects, please contact the data protection officer [1].

Another potential issue with this research is the use of the proposed tools to act on extracted behaviour, as concrete actions might be planned based on this method's predictions. To address uncertainty and potential flaws in the predictions, we strongly recommend including a human in the loop when making decisions. Additionally, privacy considerations should be thoroughly incorporated when applying this method. Data should be obtained in compliance with the requirements of the GDPR, and further measures should be taken to ensure user privacy (e.g., anonymising the data).

References

[1] Liu Yi Lin, Jaime E Sidani, Ariel Shensa, Ana Radovic, Elizabeth Miller, Jason B Colditz, Beth L Hoffman, Leila M Giles, and Brian A Primack. Association between social media use and depression among us young adults. *Depression and anxiety*, 33(4):323–331, 2016.

[2] Alice Ju, Sun Ho Jeong, and Hsiang Iris Chyi. Will social media save newspapers? examining the effectiveness of facebook and twitter as news platforms. *Journalism practice*, 8(1):1–17, 2014.

[1] https://www.uu.nl/en/organisation/practical-matters/privacy/data-protection-officer

[3] Yannis Theocharis. The wealth of (occupation) networks? communication patterns and information distribution in a twitter protest network. *Journal of Information Technology & Politics*, 10(1):35–56, 2013.

[4] Laurens HF Müter and Remco C Veltkamp. Analysing protest-related tweets: An evaluation of techniques by the open source intelligence team. In *Future of Information and Communication Conference*, pages 41–60. Springer, 2024.

[5] Lee Humphreys, Phillipa Gill, and Balachander Krishnamurthy. How much is too much? privacy issues on twitter. In *Conference of international communication association, singapore*, 2010.

[6] David Ediger, Karl Jiang, Jason Riedy, David A Bader, Courtney Corley, Rob Farber, and William N Reynolds. Massive social network analysis: Mining twitter for social good. In *2010 39th international conference on parallel processing*, pages 583–593. IEEE, 2010.

[7] Jesse Borghardt. Unveiling protest-related social networks on twitter: A comprehensive analysis of topic differences and temporal patterns within protest-related social networks in the netherlands. Master's thesis, Utrecht University, 2023.

[8] Julia M Puaschunder. Big data ethics. *Puaschunder, JM (2019). Journal of Applied Research in the Digital Economy*, 1:55–75, 2019.

[9] Mark Tremayne. Anatomy of protest in the digital era: A network analysis of twitter and occupy wall street. In *Social Networks and Social Movements*, pages 110–126. Routledge, 2016.

[10] Wei Xie, Feida Zhu, Jing Jiang, Ee-Peng Lim, and Ke Wang. Topicsketch: Real-time bursty topic detection from twitter. *IEEE Transactions on Knowledge and Data Engineering*, 28(8):2216–2229, 2016.

[11] Sarah A Alkhodair, Steven HH Ding, Benjamin CM Fung, and Junqiang Liu. Detecting breaking news rumors of emerging topics in social media. *Information Processing & Management*, 57(2):102018, 2020.

[12] David Gunnarsson Lorentzen. Polarisation in political twitter conversations. *Aslib Journal of Information Management*, 66(3):329–341, 2014.

[13] Max Falkenberg, Alessandro Galeazzi, Maddalena Torricelli, Niccolò Di Marco, Francesca Larosa, Madalina Sas, Amin Mekacher, Warren Pearce, Fabiana Zollo, Walter Quattrociocchi, et al. Growing polarization around climate change on social media. *Nature Climate Change*, 12(12):1114–1121, 2022.

[14] Akemi Chatfield and Uuf Brajawidagda. Twitter tsunami early warning network: a social network analysis of twitter information flows. 2012.

[15] Martin Grandjean. A social network analysis of twitter: Mapping the digital humanities community. *Cogent arts & humanities*, 3(1):1171458, 2016.

[16] Itai Himelboim, Marc A Smith, Lee Rainie, Ben Shneiderman, and Camila Espina. Classifying twitter topic-networks using social network analysis. *Social media+ society*, 3(1):2056305117691545, 2017.

[17] José Antonio Iglesias, Aaron Garcia-Cuerva, Agapito Ledezma, and Araceli Sanchis. Social network analysis: Evolving twitter mining. In *2016 IEEE International conference on systems, man, and cybernetics (SMC)*, pages 001809–001814. IEEE, 2016.

[18] Mithileysh Sathiyanarayanan and Nikolay Burlutskiy. Design and evaluation of euler diagram and treemap for social network visualisation. In *2015 7th international conference on communication systems and networks (COMSNETS)*, pages 1–6. IEEE, 2015.

[19] Joaquín Osorio Arjona and Juan Carlos García Palomares. Spatio-temporal mobility and twitter: 3d visualisation of mobility flows. *Journal of Maps*, 16(1):153–160, 2020.

[20] Manlio De Domenico, Antonio Lima, Paul Mougel, and Mirco Musolesi. The anatomy of a scientific rumor. *Scientific reports*, 3(1):2980, 2013.

[21] Enrique Frias-Martinez, Graham Williamson, and Vanessa Frias-Martinez. An agent-based model of epidemic spread using human mobility and social network information. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 57–64. IEEE, 2011.

[22] Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000.

[23] Emilio Serrano, Carlos Ángel Iglesias, and Mercedes Garijo. A novel agent-based rumor spreading model in twitter. In *Proceedings of the 24th International Conference on World Wide Web*, pages 811–814, 2015.

[24] Chaitanya Kaligotla, Enver Yücesan, and Stephen E Chick. An agent based model of spread of competing rumors through online interactions on social media. In *2015 winter simulation conference (WSC)*, pages 3985–3996. IEEE, 2015.

[25] Z Chen. An agent-based model for information diffusion over online social networks. kent state university, 2016.

[26] Maira A de C Gatti, Ana Paula Appel, Cicero Nogueira Dos Santos, Claudio Santos Pinhanez, Paulo Rodrigo Cavalin, and Samuel Barbosa Neto. A simulation-based approach to analyze the information diffusion in microblogging online social network. In *2013 Winter Simulations Conference (WSC)*, pages 1685–1696. IEEE, 2013.

[27] Manuela López, Carmen Hidalgo-Alcázar, and Paul Leger. The effect of message repetition on information diffusion on twitter: An agent-based approach. *IEEE Transactions on Professional Communication*, 66(2):150–169, 2023.

[28] Yizhou Yan, Fujio Toriumi, and Toshiharu Sugawara. Understanding how retweets influence the behaviors of social networking service users via agent-based simulation. *Computational Social Networks*, 8:1–21, 2021.

[29] Carlos Rodríguez Lucatero. Modelling and simulating a opinion diffusion on twitter using a multi-agent simulation of virus propagation. In *Advances in Computational Intelligence: 16th Mexican International Conference on Artificial Intelligence, MICAI 2017, Enseneda, Mexico, October 23-28, 2017, Proceedings, Part II 16*, pages 102–119. Springer, 2018.

[30] Pierre Bommel, Nicolas Becu, Christophe Le Page, and François Bousquet. Cormas: an agent-based simulation platform for coupling human decisions with computerized dynamics. In *Simulation and gaming in the network society*, pages 387–410. Springer, 2016.

[31] Olivier Gutknecht and Jacques Ferber. The madkit agent platform architecture. In *Workshop on infrastructure for scalable multi-agent systems at the international conference on autonomous agents*, pages 48–55. Springer, 2000.

[32] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, Keith Sullivan, and Gabriel Balan. Mason: A multiagent simulation environment. *Simulation*, 81(7):517–527, 2005.

[33] U Wilensky. *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with Netlogo*. The MIT Press, 2015.

[34] Gregory Madey, Vincent Freeh, and Renee Tynan. Agent-based modeling of open source using swarm. 2002.

[35] Kevin Ryczko, Adam Domurad, Nicholas Buhagiar, and Isaac Tamblyn. Hashkat: large-scale simulations of online social networks. *Social Network Analysis and Mining*, 7:1–13, 2017.

[36] Jesús M Sánchez, Carlos A Iglesias, and J Fernando Sánchez-Rada. Soil: An agent-based social simulator in python for modelling and simulation of social networks. In *Advances in Practical Applications of Cyber-Physical Multi-Agent Systems: The PAAMS Collection: 15th International Conference, PAAMS 2017, Porto, Portugal, June 21-23, 2017, Proceedings 15*, pages 234–245. Springer, 2017.

[37] Cory Gackenheimer. *Introduction to React*. Apress, 2015.

[38] Erik Hanchett and Ben Listwon. *Vue. js in Action*. Simon and Schuster, 2018.

[39] Shyam Seshadri. *Angular: Up and running: Learning angular, step by step*. " O'Reilly Media, Inc.", 2018.

[40] Elar Saks. Javascript frameworks: Angular vs react vs vue. 2019.

[41] Aryo Pinandito, Hanifah Muslimah Az-zahra, Lutfi Fanani, and Anggi Valeria Putri. Analysis of web content delivery effectiveness and efficiency in responsive web design using material design guidelines and user centered design. In *2017 International Conference on Sustainable Information Engineering and Technology (SIET)*, pages 435–441. IEEE, 2017.

[42] L Verou and C Lilley. Supported Color Spaces • Color.js — colorjs.io. https://colorjs.io/docs/spaces#lab, 2020. [Accessed 16-08-2024].

[43] K McLaren. Xiii—the development of the cie 1976 (l* a* b*) uniform colour space and colour-difference formula. *Journal of the Society of Dyers and Colourists*, 92(9):338–341, 1976.

[44] Laurens HF Müter, Christof van Nimwegen, and Remco C Veltkamp. Dutch tweets before, during, and after a black lives matter demonstration in amsterdam: Expert annotated data, protocol, and labelling tool. *Data in Brief*, 55:110527, 2024.

[45] Bart W. Schermer, Dominique Hagenauw, and Nathalie Falot. Handleiding algemene verordening gegevensbescherming en uitvoeringswet algemene verordening gegevensbescherming, January 22, 2018. URL: https://www.rijksoverheid.nl/onderwerpen/privacy-en-persoonsgegevens/documenten/rapporten/2018/01/22/handleiding-algemene-verordening-gegevensbescherming (accessed: 16 February, 2023).