Thesis Proposal

Contextualized Conversational Network Dynamics on Social Media

Thomas Magelinski

April 28, 2022

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Kathleen M. Carley, Chair, Carnegie Mellon University
Renaud Lambiotte, University of Oxford
Patrick Park, Carnegie Mellon University
Osman Yağan, Carnegie Mellon University

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Copyright © 2022 Thomas Magelinski

Abstract

Network Science provides a framework to understand the large-scale discussions that happen on social media and their impact on society. However, a standard network model of a conversational network destroys the context that users are interacting within. First, the interactional context is destroyed. The interactional component of context includes the content of the conversation in which the users are interacting. When interactional context is not accounted for, separate discussions are combined into one big network, artificially inflating the number of nodes and edges in the network. This leads to inaccurate information about conversation structure and important actors. Next, the personal context is destroyed. The personal component of context includes the attributes of the users involved, as observed through their self-descriptions. Long-standing social theory of offline social communities such as self-categorization place great importance on personal context. Thus, this context needs to be accounted for to test these theories in the social media setting.

This thesis provides the theory and methodologies needed to account for both interactional and personal contexts which were previously lost in network analysis of social media conversations. Specifically, I study the importance of these contexts as they relate to community dynamics. I find that network structure is indeed dependent on interactional context, indicating that existing non-contextualized analyses could be improved. When investigating personal context, I find that the long-standing theory of self-categorization can be extended from offline social communities to massive online communities, with some important limitations. Taken together, the dynamic contextualized analysis outlined in this thesis furthers our understanding of attribute salience in online interactions. Each of these analyses is performed on multiple case studies, providing both validation and a set of examples used to detail a list of best practices for contextualized network analysis.

Introduction

Overarching Thesis Goal

Many important problems rely on a strong understanding of online communities, especially in the area of social cybersecurity. However, current methods give an obscured view of online communities because they do not account for different contexts of interactions occurring in the data. This phenomenon is illustrated in Figure 1, using the Reopen America discussion as an example. The "mixed" network in Figure 1, corresponds to the type of network that current methodologies use to understand social communities. Nodes represent users, and they are connected when they interact directly with each other, e.g., when one user replies to another user's Tweet. The network is "mixed" in that it combines all of the data into a single network. However, the data is actually comprised of three separate discussions, one about protests, one about strategy, and one about BLM. Upon separating out the networks by the context in which connections are made, we see that each conversational context has a different network structure and different user communities.

Current methods only have access to the mixed-context view of interactions where none of the contextualized communities can be observed. Thus, non-contextualized analysis leads to an obscured view of communities. Further, any detected communities are more difficult to analyze, since the full context of their interactions is not understood. Thus, the first part of this thesis is concerned with developing methods for performing this network separation based on context. From there, methods are developed to better understand the resulting inter-related contextualized networks.

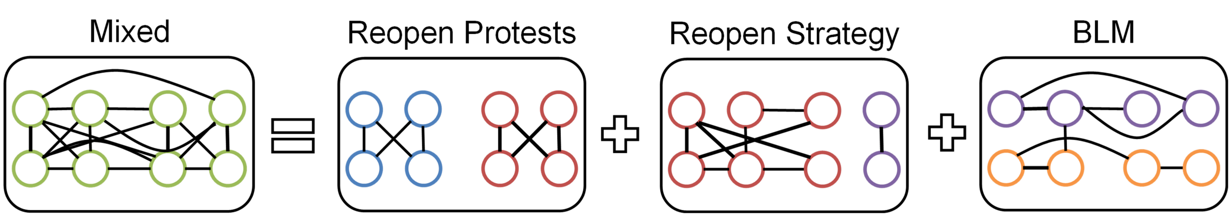

Figure 1: Cartoon illustration of how contextual mixing can hide community structure. Networks represent users connected through conversation. Nodes are colored by their network community. The "mixed" network includes links from all three contexts. This lack of contextualization hides the community structure seen in the contextualized networks. After contextualization, we understand that Reopen Protests and Reopen Strategy have similar structure, while the BLM conversation is very different.

While Figure 1 illustrates “context” in terms of the information that is being exchanged between pairs of users, there is another important aspect of context that is also often overlooked: personal attributes. Personal attributes play a key role in sociological theories of offline interaction, and there has been limited work in testing these theories within social media. The same edge between two people may take on different meanings if the two users describe themselves to be very similar or very different on key issues. Zooming out, communities can undergo change without changing their individual edges when the members collectively change their self-image. For example, a group of political actors may keep their alliances the same while collectively moving in ideology.

In this thesis, I develop a series of tools and analyses to contextualize the interactions that we observe in large social media datasets, and thereby better understand online community dynamics. This begins with methods for extracting the edge-based contexts in order to perform the network separation observed in Figure 1. From there, methods for analyzing the social dynamics within and between these contexts are developed. Following this, a method for studying the relationship between self-descriptions and online social communities is given, shedding light on the applicability of offline sociological theories to the online domain. Lastly, a pipeline combining both edge-based and description-based contextualization is presented.

Literature Review

Social Cybersecurity

Perhaps the most pressing area of research relying on a solid understanding of online communities is the area of social cybersecurity, which is defined in [19] as follows:

> Social Cyber-security is an emerging scientific area focused on the science to characterize, understand, and forecast **cyber-mediated** changes in human behavior, social, cultural and political outcomes, and to build the cyber-infrastructure needed for society to persist in its essential character in a **cyber-mediated** information environment under changing conditions, actual or imminent social cyber-threats.

Thus, this area of work encompasses a number of important challenges in the information environment including the spread of disinformation and the measurement of polarization [18]. Early work on the science of “fake news,” for example, calls for further work to understand its spread and how it is received [33, 51]. At the individual level, network centrality measures are often used to determine important actors in a conversation. However, these centrality measures have been found to be sensitive to the quality of the observed network compared to the underlying network; thus, a de-contextualized network adds “noise” edges to the point which centrality analysis may be unreliable [14]. These problems are typically studied at the community-level, and as such, community-detection is often including in information operation analysis pipelines [86].

Analysis of polarization, too, often relies on analysis of interaction networks and could thereby benefit from a contextualized approach. Specifically, distinct communities within retweet networks are often used as evidence of polarization [22, 30, 81]. However, these analyses do not *contextualize* the observed retweets. Without this contextualization, it is difficult to distinguish if communities are polarized because they are supporting opposing ideas, or they are simply involved in different discussions.

A final example is the problem of stance-detection, where social media data is used to label users' position on a topic, e.g., pro or anti-gun control. Note that stance detection is closely related to the study of polarized communities, however it is methodologically distinct in that it uses content-based approaches to label users before performing network analysis [50]. Stance detection methods such as that in [50], assume that all observed data is on-topic enough to leverage, however early results of the contextualization process developed in this thesis suggests that this is not the case. The contextual mixing that occurs in social media datasets could be harming results of these analyses, and thereby could be improved by this work.

Community Detection and Clustering

Community detection is the problem of dividing a network into sub-networks, or "communities" where nodes are more closely related to other nodes within the community than they are to nodes in other communities [93]. Community detection is a core problem in the study of network analysis, and as such many methods have been developed in the space with the dominant approaching being modularity maximization [10, 64, 83].

Two sub-areas of community detection are of particular relevance to this thesis. First, is the application of community detection to spatial networks, or those where nodes are fixed at a location in space [27, 66]. Spatial networks are relevant but distinct from contextualized networks. While spatial networks have nodes embedded in space, contextualized networks can be modeled to have edges embedded in space, where their spatial position indicates the context occurring in that interaction. Thus, the methods from spatial networks will not be directly applicable, but may be worth considering in the development of new contextualized methods. When it comes to edges embedded in a vector space, classic clustering techniques such as DBSCAN and its variants are applicable [26, 29, 58, 59, 77]

More directly related is the area of multi-view, multi-layer, or multi-slice networks, which expands upon traditional networks with the addition of distinct edge types [1, 61]. A contextualized conversation could be modeled as a multi-view network where users are nodes and edge types represent interactions within different contexts. For example, two users might have one edge indicating their conversation about sports, and another edge indicating their conversation about politics. A series of specialized techniques for clustering multi-view networks have been developed that give a single definition of communities that combines information from all views [23, 47, 62]. This is a useful approach to incorporating contextual information to improve the quality of detected communities. This is particularly important due to work that indicates that multi-layer cluster structure drives the diffusion of information over a multi-layer network [95, 97]. In our case, this could mean that decomposing social media conversations into a multi-layer network could uncover diffusion patterns that were obscured through contextual mixing. However, multi-view clustering's output of a single definition of communities will not allow for the comparison of contextualized communities or the analysis of communities shifts between contexts, as is the focus of this thesis.

Dynamics of Network Communities

Understanding the dynamics of network communities is another core area of work in Network Science, however the prevailing models are difficult to apply to social media data. A popular approach to modeling network dynamics is to use network snapshots, which model the dynamic network as a sequence of static networks, usually constructed from the edges occurring within fixed time windows [70, 71, 82]. Snapshot-based approaches have also been developed on temporal networks [43, 44, 57]. These approaches then compare the snapshots, either at the network level or based on community structure. Such comparison is not possible in social media datasets, where there may be little overlap in the users present in different snapshots, and where adjacent snapshots may have networks that differ in size by orders of magnitude. While snapshot-based approaches have some, but not enough, tolerance for the transience of nodes seen on social media, statistical methods are even more restrictive. Many statistical approaches, such as the stochastic actor-oriented model assume near perfect knowledge of node connections, measured at regular intervals, which is far from the data seen on social media [65, 78, 79].

Trails are another approach to understanding network community dynamics that is relevant for this thesis [7, 16]. Trails can be used to model nodes' transitions between semantic states while accounting for the time between these states. For example, in [16], trails modeled how terrorist organizations transitioned between different types of attacks. In this thesis, trails will be used to understand how users transition between contexts.

Story and Topic Detection

For the problem of content-based contextualization, the areas of topic detection and story detection are very relevant as they both make use of social media text to better understand the context of a post. These methods are slightly different than the notion of context that I will use in this thesis, as will be explained. Also, there is little work that goes beyond the detection and analysis of a topic or a story to understand how they relate to community dynamics.

Topic detection seeks to uncover patterns, or "topics" in a collection of text documents [9]. These topics are typically characterized by their most prominent and frequent words. There are many topic detection models that have been developed, including a number of methods that have been designed specifically for social media by leveraging the brief nature of social media posts and the presence of hashtags [4, 21, 28, 45, 54, 90, 92, 98].

While topics can be distilled to a series of words stories often have a notion of a topic tied to a specific event. Story detection methods build on topic-based approaches to find temporally prominent topics corresponding to events that occur during data collection [2, 3, 25, 67, 80]. Again, these models predominantly leverage social media text, but also use temporal patterns.

For this thesis, a method of accounting for context will be developed similar to the techniques used in topic and story detection. However, the method will expand on the usage of text by including both hashtags and URLs, as well as the conversational structure directly.

Machine Learning on Networks

Recent developments in machine learning will enable the content-based contextualization method proposed in this thesis. The machine learning community has seen increasing interest in deep learning methods applicable to graphs. These approaches work by converting network-based data into vector-based data. Some approaches use random walks to generate sequences of nodes which can then be fed to a skip-gram architecture to embed nodes in a vector space based on network structure alone [34, 72] More commonly, node attributes (in vector form) are required. Nodes can then aggregate information from their local neighborhoods to obtain their vector embedding [17, 20, 63, 91]. This has led to a large area of research into what type of aggregation scheme nodes should adopt in different scenarios [15, 36, 49, 87].

A graph-based framework can be used to model Twitter posts, as each post may be connected to other posts (replies or quotes), hashtags, and URLs. Thus, a twitter dataset can be seen as a heterogeneous network connecting tweets, hashtags, and URLs. Representing this network in a vector space enables contextualization of interactions observed in Tweets. The majority of information in a tweet is encoded in the tweet's text, which could be represented by a vectorized node-feature using a variety of different natural language techniques [11, 24, 48, 60].

Until recently, feature-based methods required supervision or some labeled training data to work with. However, deep graph infomax has been developed as a framework for learning feature-based representations of graphs in an unsupervised manner using mutual information [88]. Further, this has been expanded to heterogeneous networks [75, 88]. These methods will be at the core of the contextualization model proposed in Chapter 2.

Networks and Identity

For the other type of context considered in this thesis, personal descriptions, there is a wide area of prior work. Sociology has long been concerned with how internal processes play out at the community level. The specific theories most relevant to the connection between individual attributes and community dynamics are social identity theory and self-categorization theory [8, 37, 39, 40, 41, 42, 46, 69, 84, 85]. These theories posit that the concept of self is defined in terms of attributes and these attributes are selected with respect to the community that an individual is or wants to be a member of. Self-categorization theory outlines the idea of a "community prototype" or a collection of attributes that would belong to a prototypical member of that community.

Social theory states that members are aware of these prototypes and are aware of how their attributes compare to it. The theory posits that these relationships are key factors in tie formation and group dynamics. Specifically, people with prototypical attributes have higher potential for leadership roles. Conversely, community members who are poorly aligned with the group prototype will seek to conform to the group to improve their status. The theory of prototype adoption is quite similar to models of correlated information spread, where abstract bits of information are spreading along a network, but the adoption of these bits of information can be correlated [96]. This is similar to prototype adoption in that a number of attributes are potentially being adopted across a network, pairs of attributes within a prototype are positively correlated, while those between prototypes are negatively correlated.

It has been found that Twitter users do signal their social identity in their biography [69]. Further, there is evidence that user self-description alignment is associated with content propagation on Tumblr [94]. These studies provide evidence that community prototypes may exist on large social media platforms like Twitter, but they offer no method for directly testing this hypothesis, as I outline in Chapter 4. Beyond testing the presence of community prototypes, further tenets of the social theory can then be tested and connections between personal attributes and contextual dynamics can be explored.

Community-Aware Centrality

An emerging area of research of relevance to the study of networks and identity is that of community-aware centrality [55, 73]. Traditional centrality measures, such as Pagerank, are concerned with quantifying the importance of nodes in a network [12, 68, 93]. These measures are a function of network structure only. However, it is well understood that community structure is an important feature of real-world networks. Thus, community-aware centrality quantifies each node's importance with respect to the given definition of the network's community structure [31, 32, 35, 73].

This field within Network Science is relevant to the thesis as it can allow for the measurement of how important attributes are relative to communities of users. Existing community-aware centrality measures, however, do not allow for the measurement of contribution (a signed quantity), and do not allow for the measurement of importance with respect to a specific community, instead giving a single score for the full network. Thus, Chapter 4 in part develops modularity vitality to solve these issues, building on the concepts of network vitalities and the key-player problem [12, 13]. Modularity vitality has since been published and has been verified as an important quantity by outside researchers [55, 74].

Data

This thesis makes use of 5 core datasets throughout its chapters. These each offer unique features meant to best test the methods being developed. Further, the use of multiple datasets results in multiple case studies of the community dynamics under investigation, providing a more robust understanding of the phenomena being examined. Each dataset and its purpose are now explained.

Reopen America

The "Reopen America" Twitter dataset was collected from April 1 to June 22 in 2020 to understand the discussion of the reopen America protests [6]. The dataset was collected using a keyword search using terms such as "reopen" and "openup," including each US state's abbreviation appended to the terms, e.g., "reopenNY." One year after collection, the reply trees were crawled to get a better view of the full conversation. The resulting dataset has 10 million unique tweets across 3.3 million users. At the time of the collection, the Black Lives Matter movement became a major point of discussion and resulted in significant context mixing. The context mixing occurring in this dataset makes is a prime candidate for analysis, both to test the developed method's ability to distinguish context and to demonstrate its importance. Thus, this dataset is used in all analysis chapters of the thesis, Chapters 2-5.

2020 US Elections

The "2020 US Election" Twitter dataset captures online discussion of the most contentious elections in recent US history. False claims of voter fraud and a stolen election were rife on Twitter and are present in this dataset. These claims have since been named "The Big Lie" and have had a lasting impact on American politics [1]. The dataset was captured using a keyword-based stream of Twitter's API from November 2 2020 to November 8 2020. This allowed for the capture of data one day before election night, which was November 3 2020, and one day after major news outlets declared Joe Biden the winner on November 7 2020. The keywords were selected in order to maximize conversation around the election. This includes general hashtags, campaign hashtags, and mentions of prominent figures in the election such as Trump, Pence, Biden, and Harris. It also includes hashtags relating to anticipated election-related issues, such as the Black Lives Matter movement, US Sanctions on Iran, issues with voting-by-mail, and claims of voter fraud. The collection resulted in 4.5M tweets. Unlike the reopen dataset, there is no competing discussion present. Thus, the election dataset presents the "normal" scenario where keywords search alone provides moderately successful contextualization. Also, this dataset spans a much shorter time period than the Reopen dataset, which can give examples of dynamics occurring on different time scales. Because this dataset offers contrast to the Reopen America dataset, it is also used in all analysis chapters of the thesis, Chapters 2-5.

Ukraine Legislature

The Ukrainian Legislature dataset is the record of all Ukrainian legislative votes cast in the 22-month span of the 7th convocation of the Rada. The Ukrainian revolution of 2014 occurs midway through the convocation, drastically changing the political allegiances observed. While this network is extremely different than the data seen on social media, it provides an example of a large, ground-truth change in communities to detect. As such, it is used to validate the community dynamics method developed in Chapter 3.

Coronavirus

The Coronavirus Twitter dataset is the largest dataset examined in this Thesis. It was collected using a keyword-based stream of the Coronavirus discussion resulting in 77 million tweets. The Twitter API does not allow for retroactive collection of a user's profile information. Instead, a user's profile information can only be obtained by direct query or by observation when a user's tweet enters a collection. This makes tracking the evolution of a collection of a group of users' attributes over time difficult. It is also effectively impossible to track the evolution of attributes over an unexpected event.

[1] https://www.npr.org/2022/01/05/1070362852/trump-big-lie-election-jan-6-families

The long-standing collection of the Coronavirus discussion, however, resulted in a longitudinal picture of users' profiles who were active in the discussion of the virus. The coronavirus discussion was general enough to include users across many different interests. Further, the dataset spans the murder of George Floyd and the subsequent rise of the Black Lives Matter Movement. Thus, this dataset is uniquely positioned to study the dynamics of user attributes at the community level, and to specifically study the adoption or lack of adoption of attributes in support of Black Lives Matter, a highly polarizing issue. This will be studied in Chapter 4.

Specialized News Discussion

As part of the analysis pipeline detailed in Chapter 5, a series of best practices in data collection will be provided. The first goal is to provide a set of procedures that researchers can follow to yield the best results from the tools outlined in the previous chapters. The second goal of this dataset is to directly demonstrate the robustness of the methods as well. The collection will maximize the conversational connections between Twitter users through the new Conversation Collection feature of Twitter's V2 API [2]. This feature enables the collection of full reply trees, which were previously unobtainable. With more of the conversational structure available, it is expected that this specialized dataset will be contextualized in a clearer way than the other datasets, providing a useful case study to demonstrate the full contextualized analysis pipeline.

Initially, all tweets containing news links from 6 news agencies will be collected within a short time period, likely one day. The six agencies have been chosen to represent different types of popular news agencies, which may drive different types of conversations, thereby acting as a robustness test for the analyses. The news agencies are as follows: two direct reporting agencies (Reuters and Associated Press), one American left-leaning (CNN), one American right-leaning (Fox), and two state-sponsored agencies (CGTV and RT). The conversation collection will then be used to obtain the full threads surrounding news-related posts on that day. This will result in large number of conversations talking about different topics from different points of view. This dataset will be collected and used in Chapter 5 which will cover application of all the tools developed in preceding chapters.

[2] https://blog.twitter.com/developer/en_us/topics/tools/2020/introducing_new_twitter_api

Research Plan

Chapter 2: Contextual Dynamics of Social Media Discussions

Guiding Questions

A central premise of this thesis is that online interactions occur within an observable context. All of the information surrounding an interaction serve to provide this context. For example, an interaction between two Twitter users can occur when one user mentions another in a tweet. This tweet may contain text, hashtags, and URLs which provide information about the context that user interaction is occurring within. Further, the tweet may be responding (replying or quoting) to another tweet which in turn has even more contextual information. Network analysis can tell us information about a conversation's structure and key participants, however understanding of the content and context of these conversational interactions is critical. This leads us to the first guiding question of this chapter:

- **RQ2.1: How can conversational context be operationalized?**

I assume that answers to this question will enable the separation of tweets into different groups for analysis.

As previously discussed, the conversational network obtained using an entire dataset is actually a summation of many contextualized networks. This summation is likely to be altering the results of conversational network analysis. The results from RQ2.1 enable these differences to be studied:

- **RQ2.2: How do contextualized network characteristics differ from the larger conversational network?**

Here, "network characteristics" include nodeset overlap, central members, and community structure. Major differences in any of these areas will be a significant finding highlighting the need for contextualized analysis.

Once contextualization is possible and its effects are understood, I turn to analyzing the dynamics of the conversational contexts themselves:

- **RQ2.3: How do users transition between conversational contexts and what does that tell us about the conversation as a whole?**

The answers to these questions will prove the importance of contextualized network analysis both for improving the existing workflows and for answering new questions, such as how the conversations in a dataset develop over time.

Proposed Approach

To operationalize conversational context, two approaches are proposed a simple label propagation approach and a more sophisticated deep learning approach. The first approach yields interpretable contexts which serves to validate the deep learning approach. From there, network analysis including community detection and discovery of central members (through measures such as PageRank and degree), will be performed both on contextualized networks and on the full conversational networks. The comparison of results will give the answer to research question RQ2.2. Finally, the flow of users between conversational contexts will be modeled through a Markov transition matrix. This matrix gives the probability that a user, who is active in one conversation, will transition to another. Properties of this matrix will be leveraged to infer the relationships between contexts. For example, the presence of a “sink” in the transition matrix indicates an important conversational context which draws in users from many other contexts.

The Semi-Automated Approach: Label Propagation

This simple method makes the assumption that conversational contexts will be well-captured through popular URLs. That is, posts which provide commentary on a URL are considered to be part of the same conversation as the URL itself. It further assumes that tweets replying to a conversation are also a part of that same conversation. Conversational drift poses a threat to this assumption; two users who repeatedly reply to one another may slowly get off the topic at hand and drift into an unrelated conversation. In the Reopen America dataset, the average component in the tweet-tweet network was 11.6 tweets, with 90% of tweets within 2-tweets of the initial tweet. This gives very little room for conversational drift to occur in the data.

This simple model takes hand-annotated URLs as input. These URLs are annotated to correspond to the same conversation. For example, a URL pointing to the New York Time’s COVID tracker is part of the same conversation as CNN’s COVID tracker, so both URLs will receive the same label, perhaps “COVID updates.” Top tweets (those with the most favorites and retweets), can also be useful sources of conversational contexts as they may include first-hand information (text, images, video, etc.). As such, these can also be included in the seed labeling stage. From there, label propagation is performed. In the first round, only tweets which directly use a URL are labeled. For example, all tweet containing CNN’s covid tracker will be labeled as a part of “COVID updates.” In the following rounds, the replies and quotes will be iteratively labeled. Tweets which are connected to tweets of two different conversations will not be labeled.

The Markov transition matrix for the conversational contexts discovered in the Reopen America dataset using this method is shown in Figure 2. These transitions can be used to understand the overall flow of the conversation. The presence of three sinks: “Trump’s Job Reopening,” “Reopen Strategy,” and “Black Lives Matter,” indicate status differences between conversational contexts which will be further analyzed in this chapter.

The Fully Automated Approach: Deep Tweet Infomax

This simple model is fast and leads to human-understandable contexts, but is incapable of reaching all tweets in a dataset, does not leverage much of the information available (text, hashtags, etc.), and requires human-annotation of data which is very time consuming. Alternatively, I

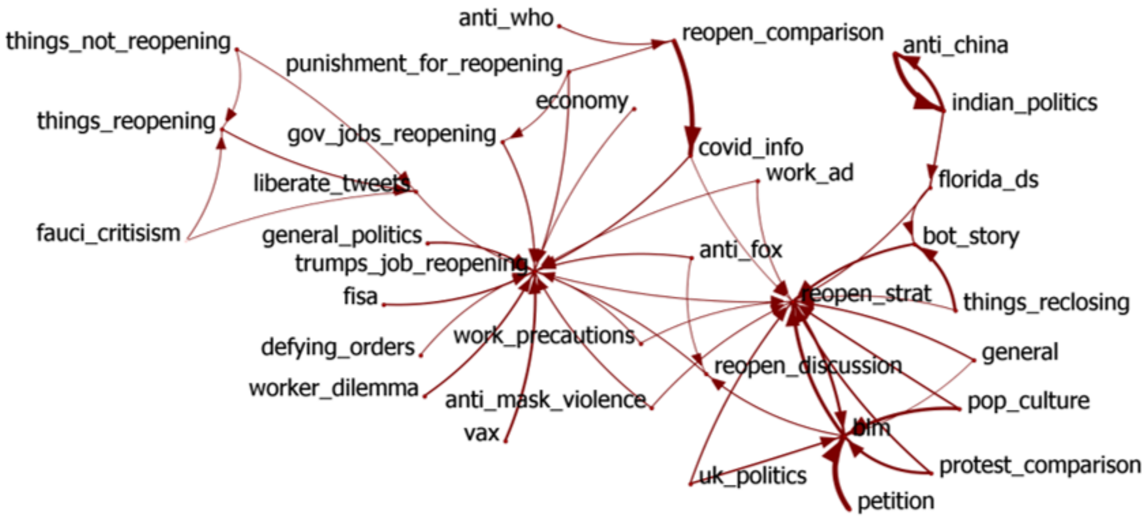

Figure 2: Markov transition matrix between conversational contexts in the Reopen America dataset. Transition probabilities below 0.2 have been trimmed for readability.

propose a second approach which leverages more information and does not require any hand-labeling. The downside of the completely automated approach is that it is up to the analyst to make sense of the contexts after they are discovered.

This second model is a deep graph learning model, which takes into account tweet text, hashtags, URLs, and conversational structure. The flow of information in 1 step of the graph neural network architecture can be seen in Figure 3. A tweet aggregates information from tweets that it is connected to (replies, or quotes in either direction), hashtags, and URLS. Hashtags and URLs, obtain information by aggregating from the tweets that they are used in. This approach allows for all tweets using the same hashtags and URLs to pass information to one another in a memory-efficient manner, while obtaining hashtag and URL representations simultaneously. This model is trained using Deep Graph Infomax, leading to the informal name of the approach of Deep Tweet Infomax (DTI). The architecture will now be described in detail.

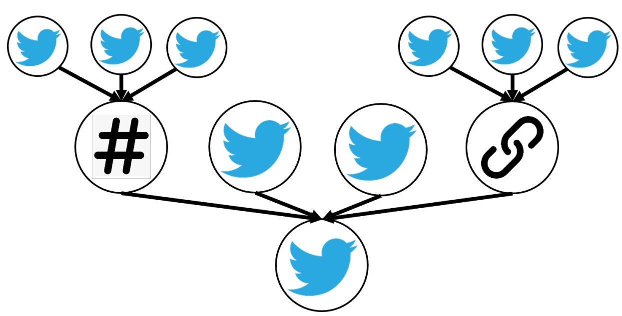

Figure 3: Visualization of how tweets gather information in the proposed model. Tweets gather information from other tweets they reference (replies, or quotes), hashtags they use, and URLs they use. Hashtags and URLS gain their representation from the tweets that use them.

Let t, u, and h represent nodes of the type tweet, URL, and hashtag, respectively. They will be indexed using subscripts, e.g., t_i corresponds to the i^{th} tweet. Feature vectors are represented with the letter x, using subscripts to indicate the corresponding node and superscripts to indicate the layer. For example, $x^0_{t_i}$, represents the 0^{th}-layer vector (otherwise known as the feature vector) for the i^{th} tweet. I will make use of a neighborhood function $\mathcal{N}$, which takes in a node and returns the set of its neighbors. Subscripts of the neighborhood function allow for the return of only a specific type of neighbor. For example, $\mathcal{N}_u(t_i)$ returns all of the URLs connected to the i^{th} tweet.

First, tweet-features are derived from the tweet's text. To enable multi-language analysis, the pre-trained[3] and language-aligned vectors trained using fastText on the Wikipedia corpus were used [11, 48]. The use of language-aligned vectors allows us to place similar tweets in the same discussion, even if they are tweeting in different languages. For each tweet, a normalized tf-idf weighting of the fastText word vectors was used to obtain a 300-dimensional tweet-text embedding. This procedure to embed tweets in Arabic, English, French, German, Hebrew, Italian, Portuguese, Romanian, Russian, Spanish, and Turkish, covering over 95% of the reachable tweets. Finally, feature propagation is deployed to obtain a feature vector for the remaining tweets [76].

Hashtags and URLs themselves do not have an initial feature representation. Instead, they aggregate the feature representation from all of the tweets that they are used in. As tweet-representation improves, so too does hashtag and URL representation. Thus, the first step of the representation process is for hashtags and URLs to aggregate information from the tweets that they appear in, as seen in Equations 1 and 2, where AGG a learnable aggregation function, and σ is an activation function.

$$x^0_{h_i} = \sigma(\text{AGG}(\{x^0_{t_i}, \forall t_i \in \mathcal{N}_t(h_i)\})) \tag{1}$$
$$x^0_{u_i} = \sigma(\text{AGG}(\{x^0_{t_i}, \forall t_i \in \mathcal{N}_t(u_i)\})) \tag{2}$$

Now that all nodes have feature vectors, tweets aggregate from their heterogeneous neighborhoods. Separate aggregation functions are learned for the tweets, hashtags, and URLs that a tweet is connected to, which are then averaged, and an activation function is applied, as seen in Equation 3.

$$\begin{aligned} x^1_{t_i} = \sigma(\frac{1}{3}(&\text{AGG}(\{x^0_{h_i}, \forall h_i \in \mathcal{N}_h(t_i)\}) \\ &+ \text{AGG}(\{x^0_{u_i}, \forall u_i \in \mathcal{N}_u(t_i)\}) \\ &+ \text{AGG}(\{x^0_{t_i}, \forall t_i \in \mathcal{N}_t(t_i) \cup \{t_i\}\}))) \end{aligned} \tag{3}$$

The process thus far defines the network over which features are passed, and the order in which to pass them. The selection of the aggregation function, AGG, is the main topic of debate within graph neural network research. In future work, AGG, can be expected to be substituted for the new state-of-the-art aggregation schemes. For now, I employ the GraphSAGE aggregation, which is the initial aggregation scheme applied in the Deep Graph Infomax work [36]. This

[3] https://fasttext.cc/docs/en/aligned-vectors.html

aggregation scheme is detailed for the tweet-to-tweet relationship in Equation 4, where $\mathbf{W}$ are trainable weight matrices, and $\mathbf{b}$ is a trainable bias vector. Note that the first term is only present when a feature vector from the previous layer is available. So for the initial aggregation steps of hashtags and URLs, this term is not present. Finally, a nonlinear activation function must be selected. Again following the original Deep Graph Infomax work, the PReLU activation function is selected [38].

$$x'_{t_i} = \mathbf{W}_1 x^0_{t_i} + \frac{1}{|\mathcal{N}(t_i)|} \sum_{t_j \in \mathcal{N}(t_i)} \mathbf{W}_2 x^0_{t_j} + \mathbf{b} \tag{4}$$

The process up to here details a single-layer of the architecture. Tweets will only obtain information from 1-hop away, and hashtags and URLs will only receive information from the initial feature vectors. Stacking these layers enables further information spread, and thus better representations of tweets, URLs, and hashtags. Here, three layers are stacked. However, in our case, tweet networks themselves are shallow. The vast majority of Twitter replies are replies to a base-tweet, rather than replies to replies. The choice of encoding tweet relationships with undirected edges also informs this depth. Incorporating tweets within 2-hops in either direction, is quite a lot of context to consider.

The architecture is trained using Deep Graph Infomax (DGI), an approach for learning unsupervised node representations by maximizing mutual information between patch representations and corresponding high-level summaries of graphs [88]. The DGI training process involves four steps. First, a normal forward pass on the data is performed, giving tweet representations, $\mathbf{x_t}$. Next, a readout function is applied to give a graph-level summary vector, s. Here, the Set2Vec operation was used: $\mathrm{s} = \sigma(\text{Set2Vec}(\{x_{t_i} \forall t_i\}))$, where σ is the sigmoid function [89]. Third, a forward pass is performed on corrupted data, giving corrupted tweet representations, $\tilde{\mathbf{x}}_t$. The same corruption function as the original work is used where tweet features are shuffled while edges are kept intact. Finally, a scoring function is applied to classify the tweets as corrupted or not. Equation 5 details the scoring function, where $\mathbf{W}$ is a trainable scoring matrix and σ is the sigmoid function, providing a score between 0 and 1. Binary cross entropy loss was used on the score, d, and the label (corrupted or not) to train the model.

$$d_{t_i} = \sigma(x^T_{t_i} \mathbf{W} s) \tag{5}$$

The application of this model leads to a vectorized representation of tweets, hashtags, and URLs. From there, contexts can automatically be extracted through clustering techniques such as DBSCAN and its variants [29, 59].

Challenges

First, a number of variations on the proposed DTI architecture are being studied. Without direct testing, it is impossible to know whether certain parts of the architecture will lead to more interpretable results than others. For example, whether or not the network should be directed, how deep the model should be, and the optimal embedding size all still need to be determined through testing.

The interpretability gap between the automatically extracted contexts of DTI and the hand-labeled contexts of label propagation poses the most important challenge to overcome in this chapter. While analysis of top tweets, hashtags, and URLs will characterize the context, it is unclear whether or not these will give insight to the *distinguishing* factors. If not, specialized methods will need to be developed. One potential method is to analyze the tweets, hashtags, and URLs that are the most "central" to the content within the vector space.

Next, the implications of a transition between contexts is not well understood. There may be an event which drives users to shift from one context to another. Alternatively, conversations may come to a natural end, and users look to find something else to discuss. Through the case-studies performed in this chapter, these different types of transitions can be better understood.

Chapter 3: Structural Changes in Contextualized Communication Networks

Guiding Questions

While the previous chapter uncovered and analyzed contextualized communication networks, this chapter seeks to understand their implications for user communication communities. There are two guiding research questions regarding these dynamics:

- **RQ3.1 What are the characteristics of communication dynamics within a single context?**
- **RQ3.2 Do communities of users transition between conversational contexts at similar times?**

These questions will be studied separately in the two parts of this chapter. Taken together, the results will lead to a better understanding of the dynamics of user communities within and between conversational contexts.

Proposed Approach

Node persistence, or lack thereof, is a major roadblock for applying dynamic network approaches to social media data. Many dynamic network approaches require all nodes to be present at each time step [65, 70, 71, 78, 79, 82]. On social media, however, there are often large swings in the number of users active within a social network. Before a major event can lead to tens of thousands if not hundreds of thousands of users entering or leaving a discussion. I propose to tackle this issue in two parts. First, only dynamics *within* contextualized conversations are considered, which should have less of a node persistence issue. In the second part, a network trail approach is taken to understand the dynamics of users between contexts, as trail-based methods do not have the same node persistence requirements [7, 16].

Intra-Context Dynamics: Network Snapshots

Within contexts, a snapshot-based approach is taken. Snapshot-based approaches model a dynamic network a series of static networks, generated from all the edges occurring within fixed time windows. When analyzing community dynamics, snapshot-based methods first compare the community structure of different snapshots, and then merge time periods which are similar. Existing measures tend to compare the networks directly by computing distance measures on the adjacency matrix [57]. Such approaches do not account for nuances in community structure. For example, if it is known that a dataset is best clustered using a specific community detection algorithm, there is no way to integrate that knowledge in these approaches. Further, multi-view or multi-modal data is not easily integrated.

Instead, I measure similarity on the co-group matrix. This matrix has a value of 1 for all pairs of nodes that are in the same group, and a value of 0 for all pairs of nodes in different groups. Now, snapshot similarity corresponds directly to community similarity, rather than other details of the network's topology. After pairwise similarity is computed on the co-group snapshots, it is partitioned using a 2-step optimization procedure. This procedure first maximizes the internal similarity of a block (that is, a time period should have as little community change as possible), and then minimizes the external similarity between blocks (that is, a new time period should begin when there is the most change in community structure. This procedure has been fully outlined and applied to the Ukrainian Legislature Network in a publication in Applied Network Science [52]. Figure 4 shows the co-group similarity matrix and the derived partition for the Ukrainian Legislature Network. The detected break-point occurs during the Ukrainian revolution, giving validity to the approach. Prior to the revolution, communities were seen to be stable and polarized, while afterwards they are only moderately stable and dispersed.

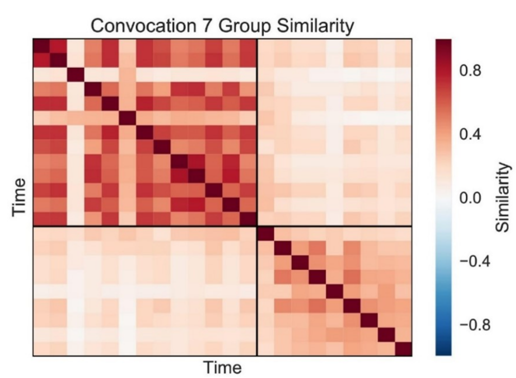

Figure 4: Visualization of the temporal group similarity matrix can be partitioned to find periods of community stability.

This approach will be used to study the internal dynamics of contextualized networks. This will result in a number of case studies, detailing how communities evolve within a conversational context.

Inter-Context Dynamics: Context Trails

For the analysis of inter-context dynamics, I propose to use network trails [7]. Network trails model transitions between networked states. Expanding on the Markov analysis proposed in Chapter 2, trail analysis includes details about *when* transitions occur. This addition enables trail clustering, which clusters users based on the similarity of their trails. Here, similarity includes both the similarity of the states and the similarity of when they transitioned between these states. I propose to employ these methods on contextualized networks to discover dynamic topic groups. This analysis will give clusters of users who are engaged in similar conversations at similar times. Depending on the strictness of time-similarity, this could be a useful form of detection of coordinated actors [56]. A method will be developed for returning the "average" trail of a cluster member, leading to a quick interpretation of a cluster. By applying this technique to a number of datasets, the behavior of dynamic topic groups will be better understood.

Challenges

There are three main challenges to be overcome in this chapter. First, the snapshot-based approach relies on node persistence, and it is unknown how persistent nodes are in contextualized networks. If users are found to be insufficiently persistent even in the contextualized setting, this will itself be an important contribution. In this case, structure of the networks will be studied not at the community level but instead at the metric level by investigating node and network level properties over time.

Second, it is unclear how to best model contextualized trails. Discretizing contexts through clustering enables direct application of trail techniques: each context can be considered a "state". However, a shortcoming of standard trail analysis is its inability to account for similarity between discrete states. For example, "classroom 1," "classroom 2," and "airplane" are all equally distinct under a standard trail model, despite the two classrooms being much more similar than the airplane. The vectorized contexts obtained in Chapter 2 pose an opportunity to rectify this shortcoming, but also pose a challenge. The vectorized representation can quantify the fact that the classrooms are very similar to each other, however specialized techniques must be developed to actually incorporate this knowledge into the trail paradigm.

Lastly, there currently is no method of extracting an "average" trail from a trail cluster. For this to be created, a definition of "average" must be developed which spans both the states visited and the patterns of transitions between them.

Chapter 4: Dynamics of Online Community Prototypes

Guiding Questions

This chapter turns to a different type of context can be leveraged to better understand online communities: individual attributes. Twitter offers a number of ways for users to signal their personality, with the most obvious being the "bio" field, a free-text field for users to describe themselves in. But there are other available fields including an area to post a URL and a location. Theories of social identity posit that individual attributes are closely related to group dynamics.

Specifically, self-categorization theory posits the presence of a "community prototype" which is a collection of attributes that are likely to be signaled from a prototypical member. The theory states that individuals will conform to the community's prototype to gain status, and those not conforming will thereby have lower status in the group. This theory has been proposed and tested for small, offline groups, but has yet to be tested on large online communities. This leads to the first major questions guiding the chapter:

- **RQ4.1: Do large online communities exhibit prototypes?**
- **RQ4.2: If so, do offline findings from self-categorization and social identity theory apply to large online communities? Specifically, is prototype adoption associated with increased status?**

If the answers to these questions are yes, an important bridge will have been formed between offline and online theories of groups. If the answers are no, this opens an opportunity to study why the offline theories do not hold, and to consider what should replace them.

Proposed Approach

A community prototype can be defined more formally as the collection of attributes which contribute maximally to a user community's structure. I propose to use a multi-view projected modularity vitality to quantify this procedure. First, I developed modularity vitality to measure node contribution to community structure in the unimodal setting [55]. Modularity allows us to quantify the quality of group structure in a network, and network vitalities allow us to measure an individual node's contribution towards a network-level quantity. Taken together, modularity vitality quantifies each node's contribution to community structure. Those with negative contributions are community bridges, while those with positive contributions are community hubs. In our work and in external work, modularity vitality has been found to identify important nodes in a number of networks [73].

The formula for the modularity vitality of a node i within a network G and partition C is given in Equation 6, where $G \setminus \{i\}$ indicates the network G after the deletion of node i and Q is the modularity function. In the modularity equation (Equation 7), A represents the adjacency matrix, k_i represents the degree of node i and c_i represents the community label of node i.

$$V_Q(G, C, i) = Q(G, C) - Q(G \setminus \{i\}, C \setminus \{i\}) \tag{6}$$

$$Q(G, C) = \frac{1}{2M} \sum_{i,j} \left(A_{i,j} - \frac{k_i k_j}{2M} \right) \delta(c_i, c_j) \tag{7}$$

This relates to attributes if a user-to-attribute bipartite network is considered. While the modularity equation used in 7 holds for unipartite networks, we can substitute Arthur's modularity for bipartite projections, i.e. the user-to-user shared attribute network [5]. Now, the contribution of an *attribute* to a community of *users* can be quantified. Note that user communities can be defined in a method separate from the bipartite attribute network.

I propose to first find communities in the user communication network, and then study the contribution of their attributes to those communities. The bipartite projection modularity can be used to measure the association of attributes with user interactions. Thus, we can directly answer research question R4.1 by observing the modularity values, with high values indicating evidence for prototype's existence in online communities.

I propose to extend this framework beyond a single attribute type by taking a multi-view network approach. In this network, each view corresponds to a different attribute type. Attribute types that will be examined include: hashtags in a user's biography, mentions in a user's biography, unigrams in a user's location, and hashtags that a user tweets. Within each view, the analysis will be performed as previously described. Expanding to a multi-view modularity vitality, the overall presence of prototypes can be quantified, as well as the presence of prototypical behavior within each view [23].

Beyond quantifying the presence of prototypes, the prototypes themselves will be constructed by ordering the attributes which maximally contribute to a community's structure. The multi-view framework allows for attributes to be compared across attribute type. The result is an ordered collection of attributes that a prototypical member of a community might display. In Figure 5, a prototype of a user community in the Reopen America dataset is displayed, attribute type is recorded with a prefix and is displayed with a unique color. The resulting prototype is easily interpretable; a typical member of this community will have things like "#theresistance", "#VoteBlue", "joebiden" in their biograph, list their location somewhere in the USA, and will retweet people like Joshua Potash and Joe Biden. Based on these facts we can understand this community to be mainstream Democrats.

Figure 5: Visualization of a community prototype in the Reopen America dataset.

For research question R4.2, I propose to use the steps previously outlined to quantify a user's prototypicality within their community. Based on their prototypicality, we can test findings from self-categorization theory. First, the correlation between prototypicality and status can be compared. I propose to quantify status through network centrality measures such as degree and

Pagerank. A high correlation between these prototypicality and network centrality would verify findings from smaller offline studies.

The second phenomena that I propose to examine is that of prototype adaptation. Social theory states that members will seek to increase their prototypicality in hopes of increasing their status within the community. This can be studied by breaking a dataset into two time periods. In the first, communities and prototypes are uncovered, and users with low prototypicality are identified. The hypothesis is that low-prototypicality users will adopt prototypical attributes by editing their biographies. I propose to compare the overall shift in prototypicality of these users and compare results using overall adoption rates.

Challenges

The first challenge to overcome in this chapter is how to best quantify prototypicality when it comes to the prototypicality vs. centrality study. While each attribute itself already has a score, a simple sum would reward users who have many attributes in their bio. Typically, it is outside of the norm to, for example, put 10 hashtags in your bio. So this would lead to a measurement error. On the other hand, and average score is better, but could result in a biography with only a single being scored as very prototypical, when in fact there isn't much information to make that distinction from. A comparison of a few different measurements will be performed to overcome this.

The second challenge to overcome in this chapter is within the study of prototype adoption. I have proposed studying the dynamics in two steps, one to create the prototypes and the second to see if they are adopted. The problem with this approach is that prototypes themselves are changing in time, so the success of this method will be dependent on the stability of the observed prototypes. If, for example, a community is rapidly changing the hashtags that they put in their biography, there will only be a short time window where that specific prototype is adopted. This effect may be limited by using a large initial time period and a small second time period, but the success of this approach remains to be seen.

Chapter 5: Pipeline for Contextualized Conversation Dynamic Analysis

Guiding Questions

The previous chapters will have furthered our understanding of online community dynamics through the analysis of the context of online interactions and self-descriptions of the users involved. However, these analyses have been performed separately. The main focus of this chapter is to combine the approaches previously developed, leading to the main guiding research question of the chapter:

- **RQ5.1 How can the tools from contextualized network analysis and online community prototypes be integrated into a dynamic community analysis workflow?**

Contextualized network analysis will benefit from attribute analysis, giving a deeper understanding of the derived communities. At the same time, the limits and applicability of prototype analysis will be better understood when considering the trail clusters of Chapter 4.

After the pipeline is constructed, the data requirements will become clear, leading to a secondary research question:

- **RQ5.2 What are the best data collection practices when applying the developed pipeline?**

Proposed Approach

Trail clusters naturally integrate with community prototype analysis because prototype construction is agnostic to the method of obtaining user clusters. That is, the contextual trail clustering approach can be directly substituted in for the previously used community detection approach on the all-communication network. Because contextualized communities are more precise and contain different membership than the standard approach, I expect prototypes to be different than those initially uncovered.

This effect will be examined through a case study, where the entire pipeline is run on a specialized dataset. This dataset will be constructed according to the requirements of the full pipeline. Based on the requirement of user attributes, and conversational connections, this dataset collection procedure will likely collect all available node attributes, and full conversation trees on Twitter. This will result in far more data than a typical keyword collection, so will require a smaller scope of investigation either through topic, time, or both.

Challenges

There are two main challenges in this chapter. First, dataset collection requirements limit the length of time that data can be collected over. Depending on the results from Chapter 4, it may be that prototype adoption occurs on a time-scale too large to be well-integrated into an ideal contextualized network analysis. To combat this, a static prototype could be constructed instead. This analysis is still useful, as it provides details to the internals of dynamic conversational communities.

The second challenge is how to understand contextualized prototypes if trail clustering is not used. Each contextualized conversation can easily be broken down into communities, and prototypes of these communities can be calculated. But how can we compare and contrast community prototypes from different contexts? and is this a useful thing to do? These questions are difficult to answer without actually applying the pipeline to a real dataset.

Contributions and Limitations

Theoretical Contributions

The proposed work makes a number of contributions to the theory of online social community dynamics. First and foremost, the importance of contextualization and its effects on network communities will be demonstrated. While a number of methods, such as topic detection and multi-view clustering imply that multiple conversational contexts are present in social media datasets, the effect that this has on our understanding of social communities has been previously unknown.

Naturally, the dynamics of online communities between conversations has also been previously unknown. Through the application of Markov transitions and trail based analysis, a number of case studies will give insights into how social media users move between points of discussion. Further, trail clustering will uncover the presence of "dynamic topic groups," which, if present, imply that some communities of users move *collectively* between points of discussion.

Another major theoretical contribution is the connection between longstanding offline theories of social identity and more recent work into self-presentation online. Through the work outlined in this proposal, whether or not communities exhibit community prototypes will be understood. Early results indicate that communities do in fact exhibit this behavior. Following this, the limits of the applicability of these theories to online behavior will be better understood as the hypothesis that prototypical members are poised to be community leaders will be tested.

On the Network Science side of theory, modularity vitality has been developed which more strongly connects the two core research areas of centrality and community structure [53]. Modularity Vitality has been externally validated as a useful measure of node importance using the linear threshold model [74].

Computational Contributions

Open-sourced code for the tools developed in this thesis will be made available for further use and improvement by the research community, including:

- modularity vitality (unipartite, bipartite, projections, multi-view)
- unsupervised tweet representation
- contextualized network dynamics
- online prototype construction

Academic Contributions

This thesis has and will result in a number of academic contributions. The contextualization method and analysis seen in Chapter 2 is under revisions for resubmission to the proceedings of ICWSM 2022. The findings of the basic contextualization method have been presented orally at the CMU IDeaS Conference 2021.

The core method of dynamic community detection has been published in Applied Network Science [52]. A follow-on study has validated the method further, which was published in the proceedings of SBP-BRiMS 2019 [53].

The development, proofs of scalability, and validation of Modularity Vitality was published in IEEE Transactions of Network Science and Engineering [55]. A follow-on study was performed which examined the modularity vitality-based method of node filtering in application to hashtag-based topic detection, which was published in the proceedings of SBP-BRiMS 2020 [54]. A talk was given at Networks 21 on the development of modularity vitality for bipartite networks and projections.

The remaining work in this thesis will result in more academic publications and presentations at academic conferences. The examination of community prototypes in Chapter 4 will be submitted to Nature Human Behavior after discussion with the editors about their interest in the work. Talks on contextualized networks and on community prototypes have also been accepted at Sunbelt 2022. Finally, the work on contextual trail clustering and the full contextualized pipeline will result in academic papers. For a timeline of the academic contributions, see Figure 7.

Limitations

A natural limitation of this thesis is its scope of applicability on social media. While core concepts and ideas like contextualized networks and contextual trails are applicable to any form of social media, the current tools are only developed for Twitter. This is for two reasons. First, the question of unsupervised social media representation is a domain-specific open research question. It is a difficult enough question that a significant portion of Chapter 2 is developed to method development specifically for Twitter. Further, comparison of specialized deep learning models across platforms are unlikely to lead to meaningful results. The second major reason for this limitation is the availability of data. Some social media platforms, like Facebook, do not allow for the collection of specific interactions between users. Other platforms, such as Reddit, simply do not have the features seen on Twitter, such as the ability to add personal identifiers to one's profile. Twitter provides enough data to fully examine contextualization in ways that other platforms do not.

With that said, each dataset has specific drawbacks, which have been discussed in the data section. These drawbacks will be minimized through the specialized data collection of Chapter 5. With that said, the ability to directly collect on changes in user attributes is limited by Twitter.

Both contextualization methods developed in Chapter 2 require significant qualitative analysis. The label-propagation approach is more burdensome for reviewers, as it requires an extensive amount of labeled data (URLs, and tweets). The unsupervised deep learning model eliminates the need for annotators, however it results in clusters which need to be interpreted. Methods will be developed to limit the difficulty of this task.

As with all time-window based approaches, the dynamic community detection method developed in Chapter 3 requires a selection of time window. This issue can be limited in a number of ways which will be discussed in detail in the chapter. The simplest approach is to perform analysis on multiple window sizes to examine dynamics occurring over different time scales. Other approaches may include using overlapping windows.

Lastly, causal conclusions on the formation of community prototypes cannot be drawn from the proposed analysis. It is possible that a user's Twitter feed is ranked such that interactions with similar-profile users is more likely. This may also result in community prototypes. Without direct access to this algorithm, which is unavailable, the effect of ranking cannot be determined. Still, results about the presence of community prototypes will hold. Further, an observation of higher-than-average adoption for prototypical attributes would bring added confidence to the presence of sociological effects.

Progress Report and Timeline

Significant progress on the proposed work has been made. The proposed timeline can be seen in Figure 6, with an enhanced and annotated timeline given in Figure 6. Within these timelines we can see that the core methods have been developed for all chapters, including unsupervised network contextualization, snapshot-based community dynamics, and bipartite modularity vitality. Their completion minimizes the uncertainty of the timeline for remaining work. The remaining work of developing supplementary methods, such as the trail clustering - based analysis, as well as the case studies on specific data has been staggered throughout the coming months. Chapter 2 is scheduled to be completed in May. Chapter 3, in June, Chapter 4 in July, and Chapter 5 in October. This gives a reasonable buffer for writing, paper submissions or revisions, and presentation creation culminating in a thesis defense in December 2022.

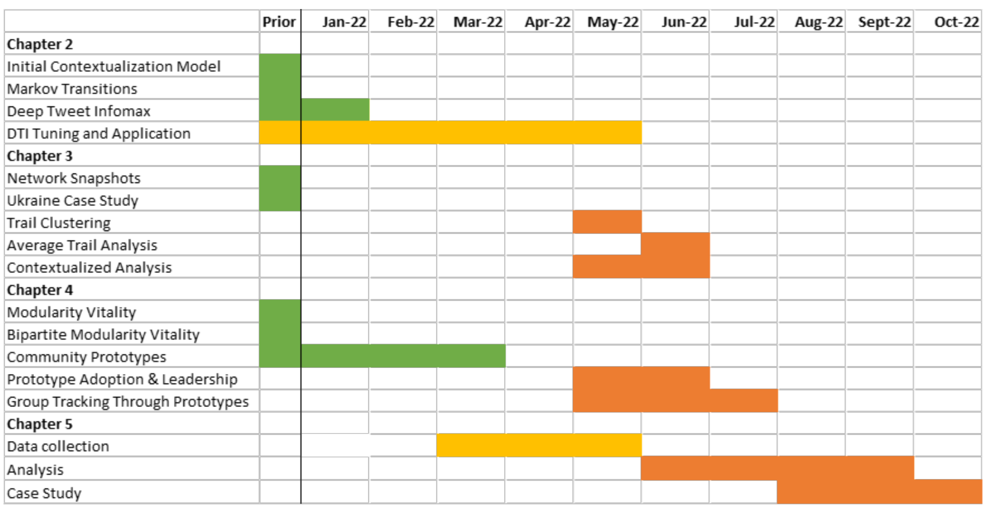

Figure 6: Proposed timeline. Green indicates completed tasks. Yellow indicates tasks that are in-progress. Orange indicates yet to be started.

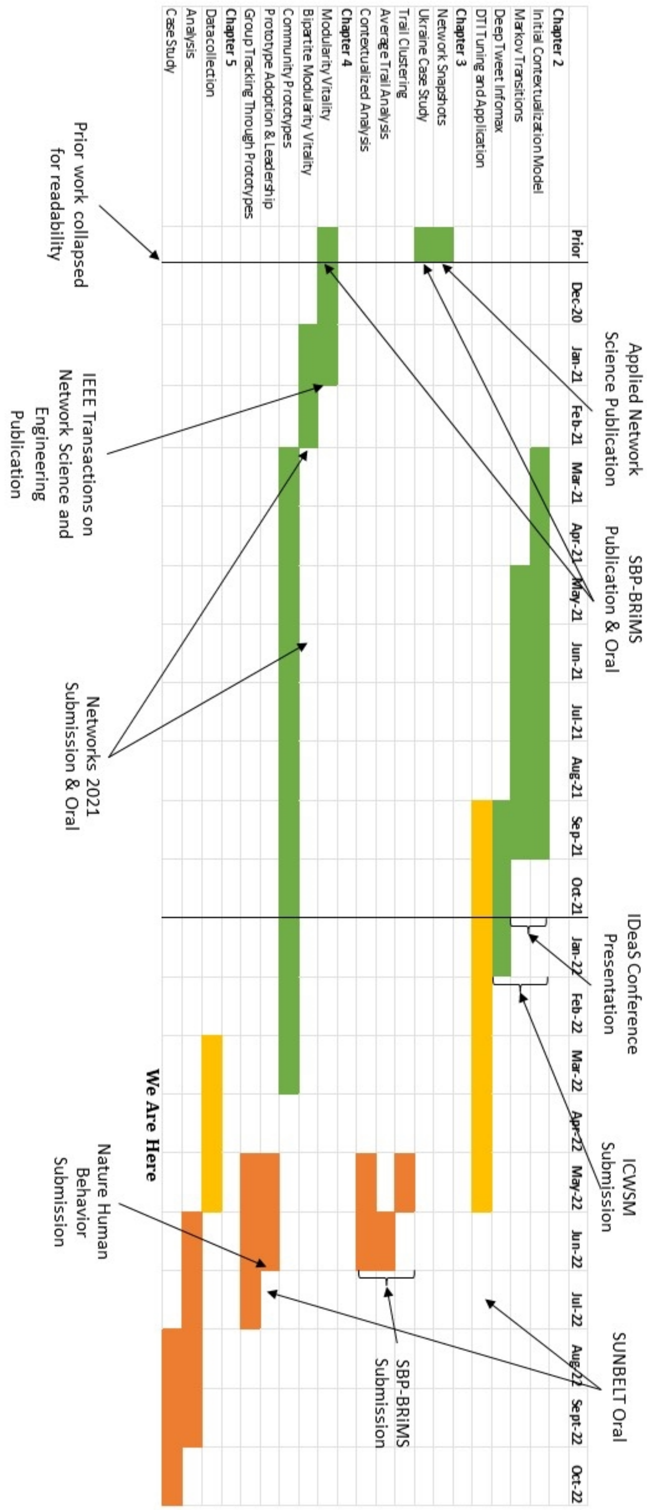

Figure 7: Annotated and extended timeline.

Bibliography

[1] Alberto Aleta and Yamir Moreno. Multilayer networks in a nutshell. *Annual Review of Condensed Matter Physics*, 10:45–62, 2019. [Cited on page 4.]

[2] Thayer Alshaabi, Jane L Adams, Michael V Arnold, Joshua R Minot, David R Dewhurst, Andrew J Reagan, Christopher M Danforth, and Peter Sheridan Dodds. Storywrangler: A massive exploratorium for sociolinguistic, cultural, socioeconomic, and political timelines using twitter. *Science advances*, 7(29):eabe6534, 2021. [Cited on page 5.]

[3] Thayer Alshaabi, Michael V Arnold, Joshua R Minot, Jane Lydia Adams, David Rushing Dewhurst, Andrew J Reagan, Roby Muhamad, Christopher M Danforth, and Peter Sheridan Dodds. How the world's collective attention is being paid to a pandemic: Covid-19 related n-gram time series for 24 languages on twitter. *Plos one*, 16(1):e0244476, 2021. [Cited on page 5.]

[4] David Alvarez-Melis and Martin Saveski. Topic modeling in twitter: Aggregating tweets by conversations. In *Tenth international AAAI conference on web and social media*, 2016. [Cited on page 5.]

[5] Rudy Arthur. Modularity and projection of bipartite networks. *Physica A: Statistical Mechanics and its Applications*, page 124341, 2020. [Cited on page 18.]

[6] Matthew Babcock and Kathleen M Carley. Operation gridlock: opposite sides, opposite strategies. *Journal of Computational Social Science*, pages 1–25, 2021. [Cited on page 7.]

[7] Mihovil Bartulović. *On Trail Comparison, Clustering, and Prediction: Building a Framework for Working with Sequential Network Data*. PhD thesis, Carnegie Mellon University, 2021. [Cited on pages 5, 15, and 17.]

[8] Parantapa Bhattacharya, Muhammad Bilal Zafar, Niloy Ganguly, Saptarshi Ghosh, and Krishna P Gummadi. Inferring user interests in the twitter social network. In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 357–360, 2014. [Cited on page 6.]

[9] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003. [Cited on page 5.]

[10] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008. [Cited on page 4.]

[11] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word

vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017. ISSN 2307-387X. [Cited on pages 6 and 13.]

[12] Stephen P Borgatti. Centrality and network flow. *Social networks*, 27(1):55–71, 2005. [Cited on page 7.]

[13] Stephen P Borgatti. Identifying sets of key players in a social network. *Computational & Mathematical Organization Theory*, 12(1):21–34, 2006. [Cited on page 7.]

[14] Stephen P Borgatti, Kathleen M Carley, and David Krackhardt. On the robustness of centrality measures under conditions of imperfect data. *Social networks*, 28(2):124–136, 2006. [Cited on page 3.]

[15] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*, 2021. [Cited on page 6.]

[16] Gian Maria Campedelli, Mihovil Bartulovic, and Kathleen M Carley. Learning future terrorist targets through temporal meta-graphs. *Scientific reports*, 11(1):1–15, 2021. [Cited on pages 5 and 15.]

[17] Shaosheng Cao, Wei Lu, and Qiongkai Xu. Deep Neural Networks for Learning Graph Representations. In *AAAI*, 2016. [Cited on page 6.]

[18] Kathleen M Carley. Social cybersecurity: an emerging science. *Computational and mathematical organization theory*, 26(4):365–381, 2020. [Cited on page 3.]

[19] Kathleen M Carley, Guido Cervone, Nitin Agarwal, and Huan Liu. Social cyber-security. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pages 389–394. Springer, 2018. [Cited on page 3.]

[20] Haochen Chen, Bryan Perozzi, Yifan Hu, and Steven Skiena. HARP: Hierarchical Representation Learning for Networks. *arXiv:1706.07845 [cs]*, June 2017. URL http://arxiv.org/abs/1706.07845. arXiv: 1706.07845. [Cited on page 6.]

[21] Xueqi Cheng, Xiaohui Yan, Yanyan Lan, and Jiafeng Guo. Btm: Topic modeling over short texts. *IEEE Transactions on Knowledge and Data Engineering*, 26(12):2928–2941, 2014. [Cited on page 5.]

[22] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*, 2011. [Cited on page 3.]

[23] Iain Cruickshank. *Multi-view Clustering of Social-based Data*. PhD thesis, Carnegie Mellon University, Oct 2020. [Cited on pages 4 and 19.]

[24] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [Cited on page 6.]

[25] Peter Sheridan Dodds, Joshua R Minot, Michael V Arnold, Thayer Alshaabi, Jane Lydia Adams, Andrew J Reagan, and Christopher M Danforth. Computational timeline reconstruction of the stories surrounding trump: Story turbulence, narrative control, and collective chronopathy. *PloS one*, 16(12):e0260592, 2021. [Cited on page 5.]

[26] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press, 1996. [Cited on page 4.]

[27] Paul Expert, Tim S Evans, Vincent D Blondel, and Renaud Lambiotte. Uncovering space-independent communities in spatial networks. *Proceedings of the National Academy of Sciences*, 108(19):7663–7668, 2011. [Cited on page 4.]

[28] Wei Feng, Chao Zhang, Wei Zhang, Jiawei Han, Jianyong Wang, Charu Aggarwal, and Jianbin Huang. Streamcube: Hierarchical spatio-temporal hashtag clustering for event exploration over the twitter stream. In *2015 IEEE 31st international conference on data engineering*, pages 1561–1572. IEEE, 2015. [Cited on page 5.]

[29] Junhao Gan and Yufei Tao. Dbscan revisited: Mis-claim, un-fixability, and approximation. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*, pages 519–530, 2015. [Cited on pages 4 and 14.]

[30] Venkata Rama Kiran Garimella and Ingmar Weber. A long-term analysis of polarization on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11, 2017. [Cited on page 3.]

[31] Zakariya Ghalmane, Mohammed El Hassouni, Chantal Cherifi, and Hocine Cherifi. Centrality in modular networks. *EPJ Data Science*, 8(1):1–27, December 2019. ISSN 2193-1127. doi: 10.1140/epjds/s13688-019-0195-7. URL https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-019-0195-7. Number: 1 Publisher: SpringerOpen. [Cited on page 7.]

[32] Zakariya Ghalmane, Mohammed El Hassouni, and Hocine Cherifi. Immunization of networks with non-overlapping community structure. *Social Network Analysis and Mining*, 9 (1):45, 2019. [Cited on page 7.]

[33] Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. Fake news on twitter during the 2016 us presidential election. *Science*, 363(6425):374–378, 2019. [Cited on page 3.]

[34] Aditya Grover and Jure Leskovec. node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, pages 855–864, San Francisco, California, USA, 2016. ACM Press. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939754. URL http://dl.acm.org/citation.cfm?doid=2939672.2939754. [Cited on page 6.]

[35] Naveen Gupta, Anurag Singh, and Hocine Cherifi. Centrality measures for networks with community structure. *Physica A: Statistical Mechanics and its Applications*, 452:46–59, 2016. [Cited on page 7.]

[36] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017. [Cited on pages 6 and 13.]

[37] Claire Hardaker and Mark McGlashan. "real men don't hate women": Twitter rape threats

and group identity. *Journal of Pragmatics*, 91:80–93, 2016. [Cited on page 6.]

[38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [Cited on page 14.]

[39] AmaÇ HerdaĞdelen, Wenyun Zuo, Alexander Gard-Murray, and Yaneer Bar-Yam. An exploration of social identity: The geography and politics of news-sharing communities in twitter. *Complexity*, 19(2):10–20, 2013. [Cited on page 6.]

[40] Michael A Hogg. Social identity theory. In *Understanding peace and conflict through social identity theory*, pages 3–17. Springer, 2016. [Cited on page 6.]

[41] Michael A Hogg. Social identity, self-categorization, and the small group. In *Understanding group behavior*, pages 227–253. Psychology Press, 2018. [Cited on page 6.]

[42] Michael A Hogg, John C Turner, and Barbara Davidson. Polarized norms and social frames of reference: A test of the self-categorization theory of group polarization. *Basic and Applied Social Psychology*, 11(1):77–100, 1990. [Cited on page 6.]

[43] Petter Holme. Modern temporal network theory: A colloquium. *The European Physical Journal B*, 88(9):234, September 2015. ISSN 1434-6028, 1434-6036. doi: 10.1140/epjb/e2015-60657-4. URL http://arxiv.org/abs/1508.01303. arXiv: 1508.01303. [Cited on page 5.]

[44] Petter Holme and Jari Saramäki. Temporal networks. *Physics reports*, 519(3):97–125, 2012. [Cited on page 5.]

[45] Liangjie Hong and Brian D Davison. Empirical study of topic modeling in twitter. In *Proceedings of the first workshop on social media analytics*, pages 80–88, 2010. [Cited on page 5.]

[46] Matthew J Hornsey. Social identity theory and self-categorization theory: A historical review. *Social and personality psychology compass*, 2(1):204–222, 2008. [Cited on page 6.]

[47] Xinyu Huang, Dongming Chen, Tao Ren, and Dongqi Wang. A survey of community detection methods in multilayer networks. *Data Mining and Knowledge Discovery*, 35(1): 1–45, 2021. [Cited on page 4.]

[48] Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018. [Cited on pages 6 and 13.]

[49] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [Cited on page 6.]

[50] Sumeet Kumar, Ramon Villa Cox, Matthew Babcock, and Kathleen M Carley. A weakly supervised approach for classifying stance in twitter replies. *arXiv preprint arXiv:2103.07098*, 2021. [Cited on page 4.]

[51] David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Roth-

schild, et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018. [Cited on page 3.]

[52] Thomas Magelinski and Kathleen M Carley. Community-based time segmentation from network snapshots. *Applied Network Science*, 4(1):25, 2019. [Cited on pages 16 and 23.]

[53] Thomas Magelinski, Jialin Hou, Tymofiy Mylovanov, and Kathleen M Carley. Detecting disruption: Identifying structural changes in the verkhovna rada. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pages 194–203. Springer, 2019. [Cited on pages 22 and 23.]

[54] Thomas Magelinski, Mihovil Bartulovic, and Kathleen M Carley. Canadian federal election and hashtags that do not belong. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pages 161–170. Springer, 2020. [Cited on pages 5 and 23.]

[55] Thomas Magelinski, Mihovil Bartulovic, and Kathleen M Carley. Measuring node contribution to community structure with modularity vitality. *IEEE Transactions on Network Science and Engineering*, 8(1):707–723, 2021. [Cited on pages 7, 18, and 23.]

[56] Thomas Magelinski, Lynnette Hui Xian Ng, and Kathleen M Carley. A synchronized action framework for responsible detection of coordination on social media. *arXiv preprint arXiv:2105.07454*, 2021. [Cited on page 17.]

[57] Naoki Masuda and Petter Holme. Detecting sequences of system states in temporal networks. *Scientific reports*, 9(1):1–11, 2019. [Cited on pages 5 and 16.]

[58] Leland McInnes and John Healy. Accelerated hierarchical density based clustering. In *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on*, pages 33–42. IEEE, 2017. [Cited on page 4.]

[59] Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11):205, 2017. [Cited on pages 4 and 14.]

[60] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. [Cited on page 6.]

[61] Peter J Mucha and Mason A Porter. Communities in multislice voting networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(4):041108, 2010. [Cited on page 4.]

[62] Peter J Mucha, Thomas Richardson, Kevin Macon, Mason A Porter, and Jukka-Pekka Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *science*, 328(5980):876–878, 2010. [Cited on page 4.]

[63] Annamalai Narayanan, Mahinthan Chandramohan, Rajasekar Venkatesan, Lihui Chen, Yang Liu, and Shantanu Jaiswal. graph2vec: Learning Distributed Representations of Graphs. *arXiv:1707.05005 [cs]*, July 2017. URL http://arxiv.org/abs/1707.05005. arXiv: 1707.05005. [Cited on page 6.]

[64] Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006. [Cited on page 4.]

[65] Nynke MD Niezink, Tom AB Snijders, and Marijtje AJ van Duijn. No longer discrete:

Modeling the dynamics of social networks and continuous behavior. *Sociological Methodology*, 49(1):295–340, 2019. [Cited on pages 5 and 15.]

[66] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. A tale of many cities: universal patterns in human urban mobility. *PloS one*, 7 (5):e37027, 2012. [Cited on page 4.]

[67] Miles Osborne, Saša Petrovic, Richard McCreadie, Craig Macdonald, and Iadh Ounis. Bieber no more: First story detection using twitter and wikipedia. In *Sigir 2012 workshop on time-aware information access*, pages 16–76. Citeseer, 2012. [Cited on page 5.]

[68] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999. [Cited on page 7.]

[69] Arjunil Pathak, Navid Madani, and Kenneth Joseph. A method to analyze multiple social identities in twitter bios. *Proceedings of the ACM on Human-Computer Interaction*, 5 (CSCW2):1–35, 2021. [Cited on page 6.]

[70] Tiago P Peixoto and Laetitia Gauvin. Change points, memory and epidemic spreading in temporal networks. *Scientific reports*, 8(1):1–10, 2018. [Cited on pages 5 and 15.]

[71] Tiago P Peixoto and Martin Rosvall. Modelling sequences and temporal networks with dynamic community structures. *Nature communications*, 8(1):1–12, 2017. [Cited on pages 5 and 15.]

[72] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. DeepWalk: online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14*, pages 701–710, New York, New York, USA, 2014. ACM Press. ISBN 978-1-4503-2956-9. doi: 10.1145/2623330.2623732. URL http://dl.acm.org/citation.cfm?doid=2623330.2623732. [Cited on page 6.]

[73] Stephany Rajeh, Marinette Savonnet, Eric Leclercq, and Hocine Cherifi. Characterizing the interactions between classical and community-aware centrality measures in complex networks. *Scientific Reports*, 11(1):1–15, 2021. [Cited on pages 7 and 18.]

[74] Stephany Rajeh, Ali Yassin, Ali Jaber, and Hocine Cherifi. Analyzing community-aware centrality measures using the linear threshold model. In *International Conference on Complex Networks and Their Applications*, pages 342–353. Springer, 2021. [Cited on pages 7 and 22.]

[75] Yuxiang Ren, Bo Liu, Chao Huang, Peng Dai, Liefeng Bo, and Jiawei Zhang. Heterogeneous deep graph infomax. *arXiv preprint arXiv:1911.08538*, 2019. [Cited on page 6.]

[76] Emanuele Rossi, Henry Kenlay, Maria I Gorinova, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael Bronstein. On the unreasonable effectiveness of feature propagation in learning on graphs with missing node features. *arXiv preprint arXiv:2111.12128*, 2021. [Cited on page 13.]

[77] Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. Dbscan revisited, revisited: why and how you should (still) use dbscan. *ACM Transactions on*

Database Systems (TODS), 42(3):1–21, 2017. [Cited on page 4.]

[78] Tom A. B. Snijders. The Statistical Evaluation of Social Network Dynamics. *Sociological Methodology*, 31(1):361–395, January 2001. ISSN 0081-1750, 1467-9531. doi: 10.1111/0081-1750.00099. URL http://doi.wiley.com/10.1111/0081-1750.00099. [Cited on pages 5 and 15.]

[79] Tom AB Snijders, Gerhard G Van de Bunt, and Christian EG Steglich. Introduction to stochastic actor-based models for network dynamics. *Social networks*, 32(1):44–60, 2010. [Cited on pages 5 and 15.]

[80] PK Srijith, Mark Hepple, Kalina Bontcheva, and Daniel Preotiuc-Pietro. Sub-story detection in twitter with hierarchical dirichlet processes. *Information Processing & Management*, 53(4):989–1003, 2017. [Cited on page 5.]

[81] Leo G Stewart, Ahmer Arif, and Kate Starbird. Examining trolls and polarization with a retweet network. In *Proc. ACM WSDM, workshop on misinformation and misbehavior mining on the web*, volume 70, 2018. [Cited on page 3.]

[82] Lei Tang, Huan Liu, Jianping Zhang, and Zohreh Nazeri. Community evolution in dynamic multi-mode networks. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD 08*, page 677, Las Vegas, Nevada, USA, 2008. ACM Press. ISBN 978-1-60558-193-4. doi: 10.1145/1401890.1401972. URL http://dl.acm.org/citation.cfm?doid=1401890.1401972. [Cited on pages 5 and 15.]

[83] Vincent A Traag, Ludo Waltman, and Nees Jan van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):1–12, 2019. [Cited on page 4.]

[84] John C Turner and Katherine J Reynolds. Self-categorization theory. *Handbook of theories in social psychology*, 2(1):399–417, 2011. [Cited on page 6.]

[85] John C Turner, Michael A Hogg, Penelope J Oakes, Stephen D Reicher, and Margaret S Wetherell. *Rediscovering the social group: A self-categorization theory.* basil Blackwell, 1987. [Cited on page 6.]

[86] Joshua Uyheng, Thomas Magelinski, Ramon Villa-Cox, Christine Sowa, and Kathleen M Carley. Interoperable pipelines for social cyber-security: assessing twitter information operations during nato trident juncture 2018. *Computational and Mathematical Organization Theory*, pages 1–19, 2019. [Cited on page 3.]

[87] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. [Cited on page 6.]

[88] Petar Velickovic, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. *ICLR (Poster)*, 2(3):4, 2019. [Cited on pages 6 and 14.]

[89] Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*, 2015. [Cited on page 14.]

[90] Maximilian Walther and Michael Kaisser. Geo-spatial event detection in the twitter stream.

In *European conference on information retrieval*, pages 356–367. Springer, 2013. [Cited on page 5.]

[91] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural Deep Network Embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, pages 1225–1234, San Francisco, California, USA, 2016. ACM Press. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939753. URL http://dl.acm.org/citation.cfm?doid=2939672.2939753. [Cited on page 6.]

[92] Yuan Wang, Jie Liu, Jishi Qu, Yalou Huang, Jimeng Chen, and Xia Feng. Hashtag graph based topic model for tweet mining. In *2014 IEEE International Conference on Data Mining*, pages 1025–1030. IEEE, 2014. [Cited on page 5.]

[93] Stanley Wasserman, Katherine Faust, et al. Social network analysis: Methods and applications. 1994. [Cited on pages 4 and 7.]

[94] Michael Miller Yoder, Qinlan Shen, Yansen Wang, Alex Coda, Yunseok Jang, Yale Song, Kapil Thadani, and Carolyn P Rosé. Phans, stans and cishets: Self-presentation effects on content propagation in tumblr. In *12th ACM Conference on Web Science*, pages 39–48, 2020. [Cited on page 7.]

[95] Yong Zhuang and Osman Yağan. Information propagation in clustered multilayer networks. *IEEE Transactions on Network Science and Engineering*, 3(4):211–224, 2016. [Cited on page 4.]

[96] Yong Zhuang and Osman Yagan. A vector threshold model for the simultaneous spread of correlated influence. In *ICC 2019-2019 IEEE International Conference on Communications (ICC)*, pages 1–7. IEEE, 2019. [Cited on page 6.]

[97] Yong Zhuang, Alex Arenas, and Osman Yağan. Clustering determines the dynamics of complex contagions in multiplex networks. *Physical Review E*, 95(1):012312, 2017. [Cited on page 4.]

[98] Yuan Zuo, Junjie Wu, Hui Zhang, Hao Lin, Fei Wang, Ke Xu, and Hui Xiong. Topic modeling of short texts: A pseudo-document view. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 2105–2114, 2016. [Cited on page 5.]