

Information Space Dashboard

Theresa Krumbiegel, Albert Pritzkau, Hans-Christian Schmitz

Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE

Fraunhofer Str. 20, 53343 Wachtberg, Germany

{theresa.krumbiegel, albert.pritzkau, hans-christian.schmitz}@fkie.fraunhofer.de

Abstract

The information space, where information is generated, stored, exchanged and discussed, is not idyllic but a space where campaigns of disinformation and destabilization are conducted. Such campaigns are subsumed under the terms “hybrid warfare” and “information warfare” (Woolley and Howard, 2017). In order to enable awareness of them, we propose an information space dashboard comprising various components/apps for data collection, analysis and visualization. The aim of the dashboard is to support an analyst in generating a common operational picture of the information space, link it with an operational picture of the physical space and, thus, contribute to overarching situational awareness. The dashboard is work in progress. However, a first prototype with components for exploiting elementary language statistics, keyword and metadata analysis, text classification and network analysis has been implemented. Further components, in particular, for event extraction and sentiment analysis are under development. As a demonstration case, we briefly discuss the analysis of historical data regarding violent anti-migrant protests and respective counter-protests that took place in Chemnitz in 2018.

Keywords: Information Space, Hybrid Warfare, Machine Learning

1. Introduction

In civilian emergency response and disaster management as well as in military operations, situation awareness is based on observation and orientation, and it is a necessary precondition for decision and action (Boyd, 1975). In order to reach situational awareness, information has to be collected and fed into a common operational picture. The picture comprises a representation of the actual situation as well as planned or predicted future events. Based on the operational picture, the situation can be accessed, the possibilities of own actions can be estimated, respective plans can be developed and decisions can be taken. Last but not least, own activities can be monitored and controlled.

With the advent of the once-so-called “new” media, in particular social media, the information space has become an additional domain, in which situational awareness must be reached. The information space is linked to the physical world, as it contributes to the creation of (the common picture of) reality: negotiations take place on which information counts as factual and which doesn’t, which events are considered to be “real” and how they are to be assessed, which prognoses shall be believed and which not, whom one can trust and whom one cannot. Moreover, planning and preparation of events in the physical world take place: individuals and groups are mobilized for actions, which can be as diverse as demonstrations, riots or spontaneous help in emergency situations.

It is therefore not surprising, that the information space has become a theatre of operations in its own right: on the one hand, actors can try to destabilize a society and mobilize the population for their own purposes by means of propaganda and disinformation campaigns. For such aggressive activities, the term “hybrid warfare” has been coined. On the other hand, activities within the information space can be necessary to approach a population, inform it on the situation and the measures to be taken and, thus, stabilize a society. Such activities come under the umbrella of an “integrated approach” as it is followed by the European Union, among others (Schmitz et al., 2019).

The concepts of situational awareness and a common operational picture are not as elaborated and clear for the information space as they are for the physical domain. Our aim is to contribute to the sharpening of the concepts. We do so with a bottom-up approach, namely by developing means for creating various views on aspects of the information space and thus providing elements of an operational picture. To this end, we propose an information space dashboard as an analyst’s working environment, which comes with a toolbox for information analysis. An operational picture of the information space will be different from one of the physical space. Therefore, systems for supporting situational awareness will substantially differ: an information space dashboard cannot just be an adapted command and control information system (C2IS) with a map-view as its core element.

Within this paper, we will firstly refer to related work, then outline a prototype of an information space dashboard and its components, describe a demonstration case, and, finally, conclude by summing up the results, name open issues and discuss obvious ethical questions.

2. Related Work

Bergh (2019) discusses the need of the detection of influence operations on social media: “over the past few years national defence organisations have received a wakeup call with regard to social media and their use to attempt to manipulate opinion, whether in hybrid conflicts [...] or in low-level societal manipulation”. This had already been emphasized by Franke (2011) who pointed out that Social Media campaigns do not derive from one specific source, but can be conducted by a number of various actors, including foreign and domestic governments as well as activist groups. The targets can be as diverse and are not limited to one specific group or individual. Examples that bear witness to this fact are cases of computational propaganda during the 2016 US Presidential Election, the influence political bots had on the 2016 Brexit referendum and bot networks as well as computational propaganda that played a role in the 2014

presidential election, the constitutional crisis and impeachment process in Brazil (Woolley and Howard, 2017).

Bergh (2019) poses requirements and challenges regarding the development of a software solution to support the detection of influence operations. Requirements are in particular flexibility, interoperability for information sharing between different operators and the ability to respond rapidly to changing situations. Challenges are the organisational and technical patchwork, limitations of resources like computing power and storage, heterogeneity of data and (entitled) privacy concerns. These are to be considered in the process of developing an information space dashboard.

Beside this, the analysis of social media data has become a constant field of research. An important topic, among others, is the detection of threats and hate speech in online conversations with the help of machine learning algorithms. Colbaugh and Glass (2012) and Lerman and Hogg (2012) are to be named as examples for works in this domain.

To capture the information space, Information Extraction (IE) from vast amounts of unstructured data needs to be implemented. Li et al. (2019) developed a system for multilingual knowledge extraction that is able to perform entity discovery and linking, time expression extraction and normalization, relation extraction, event extraction and event coreference. Such a tool is invaluable in the context of comprehending the information space and should be considered during the development of the dashboard. We further refer to this in chapter 5. Liu et al. (2019) introduce a technique to manage data accumulations by means of synthesis. Even though the described approach focuses on the processing of Chinese text data, methods such as the finding of subtopics and the synthesis of news articles could be transferred to data written in other languages, for example German.

Another aim of the information space dashboard is the identification of disinformation (campaigns). As can be seen in subchapter 5.2., we already conducted research in the field of fake news detection. However, we can not disregard insights from other researches. Nadeem et al. (2019), for example, present an automatic end-to-end fact checking system (FAKTA). FAKTA incorporates document retrieval, stance detection, evidence extraction and linguistic analysis in order to predict the factuality of claims.

3. Information Space

In the general understanding, the information space includes all technology-enhanced communication, coordination, and collaboration services that facilitate the creation, sharing, and exchange of information and ideas within communities of interest. The creation of such communities is a fundamental characteristic of social media, they form quickly and enable effective communication. Though these communities are only virtual, they are usually no less robust than the physical communities in which we live. In many ways they are even more robust as spatial and temporal boundaries are removed. Social media services, in particular, promote the exchange of information between members of a community in a way that encourages contributions of content improving collaboration, knowledge-sharing and engagement. The information space is increasingly being used as an information source, including infor-

mation related to national and global security.

Exploitation of the information space generally has three fundamental objectives: information discovery, situational awareness enhancement and predictive analysis. Capabilities in addressing these objectives provide essential estimates of potential risks faced by communities, economies and the environment. When exploiting media sources, analysis is commonly limited to two aspects: users as basic units of the network and content as basic elements of communication (Kwak et al., 2010). These two aspects themselves are already invaluable sources of information. However, social networks additionally offer the context of communication and interaction represented by the network itself, namely, the network topology in the form of entities and relations. In addition to content, a given network structure promotes the derivation of activity and process patterns which can significantly improve situational awareness (Helbing et al., 2014).

4. Social Sensing Capabilities

In social networks, humans are central in the sensing process. Social media services, in particular, provide a rich and flexible platform for performing mining processes with different kinds of data such as text documents, images, audio and video files. In the context of this paper, we consider online sites and applications which consist of users, social links, and interactive communications as data sources. These social media services can be seen as a subset of social media that includes a social network of some kind. Social networks are transforming into inherently multi-modal data sources. In recent years, sensor data collection techniques and services have been integrated into many kinds of social networks and have increased the richness of the data collection process in the context of the network structure. Furthermore, it renewed interest in the study of collective dynamics, and in particular the study of individual mobility patterns in addition to social relations. We envision that the whole phenomenon of social networks will continue to evolve quickly as digital technology increasingly penetrates the realm of the physical world, providing new research challenges for information systems, and especially for our dashboard approach. Since most current social network services usually implement only simple models of a social network, it should be noted that these models cannot mirror the richness of real world complexity. But even abstract representations of social dynamics have proved to be useful in acquiring knowledge for decision making and in supporting pro-active intervention before critical events occur.

5. Analysis & Visualization

Analysis involves reviewing and assessing large collections of information by means of complex processes of analytical reasoning, hypothesis formulation, and decision making. The analysis process itself is inherently iterative, involving alternating narrowing and broadening of focus, and is often performed as an exploratory search for relevant information. With the dashboard approach, we attempt to analyze the information space by emphasizing different data representations. Diverse representations of data support the

exploratory search beyond predictable fact retrieval by enabling various levels of abstraction that can be applied to different problems, questions, tasks or stages of the analysis.

Both content and interactions, as introduced in chapter 3., are considered for the purpose of discovering actionable patterns and understanding human behaviour. To understand some characteristics of typical accounts, or of the overall network and its potential reach, the most basic metrics – e.g. the number of followers and following and patterns of tweeting – serve as a starting point. These metrics, as the result of elementary language statistics, already provide appropriate controls for specifying the data and views of interest. Controls enable analysts to selectively represent the data, to filter out unrelated information, and to sort information to expose patterns. Quantitative information derived from the input data such as normalized values, statistical summaries, and aggregates, serve as additional descriptive features to support the analysis.

Traditional information discovery methods are based on content: documents, terms, and the relationships between them (Leskovec and Lang, 2008). Emergent social network services, however, allow for a range of extended features for aggregating content, attributes, and social graphs and take advantage of this newly formed environment of user-generated content. Complex relationships between content and people represented in social applications must be leveraged in order to recognize activities, events, groups and trends. Indeed, it has been observed that the use of a combination of social structure and different kinds of data can be a very powerful tool for mining purposes (Qi et al., 2012). Beyond quantitative features we can rely on a number of methods for IE which take natural language texts as input and produce structured information specified by certain criteria. Various sub-tasks of IE such as Named Entity Recognition, Coreference Resolution, Named Entity Linking, Relation Extraction and Knowledge Base reasoning form the building blocks of a complex language understanding task. Many of these methods such as text categorization referred to in 5.2. or sentiment prediction referred to in 5.5. usually reframe this complex language understanding task as a simpler sequence or token classification problem.

The analysis process is usually based on various pre-processing steps preceding the presentation in the dashboard.

5.1. Elementary Language Statistics

Quantitative analysis of texts can serve the exploratory investigation of online media as they can reveal trends and topics under discussion. To these statistical means belong frequency distributions of content words plotted in various ways, extraction of key words and key phrases, and analysis of metadata, including hash tags, among others, as “the hashtags used by ‘ordinary’ Twitter members construct their position as commentators on cultural events produced by others” (Page, 2012).

5.2. Text Clustering and Classification

Text can be clustered and/or classified according to stylistic surface phenomena which are significantly correlated with

semantic or pragmatic properties of interest. Text clustering can give rise to the topics under discussion.

We successfully conducted experiments with a so-called “fake news” filter. This filter is actually a classifier that exploits specific syntactic constructions, word choices and elements of hate speech which have been proven significant by an exploratory investigation into disinformation campaigns. These features can be used to recognize potential (!) fake news articles (Schade et al., 2018; Pritzkau, 2019). The methodology does not come without questions, however, as the extraction of language-related features can establish a bias against specific types of authors. If the features are, e.g., significant for usage of the German language by native Russian speakers – cf. (Böttger, 2008; Gladrow, 1998) – the system will automatically assign a higher “fake news” probability to articles written by Russians.

5.3. Network Analysis

To identify the communicative and interactive behavior of a user in a social network, and to detect which behaviors are unusual and might therefore hint at a bot-like behavior, it is worthwhile to analyze associated metadata. A bot becomes noticeable by the controlled character of its activity in the information space. Through metadata, it is possible to represent structures at the micro and macro level that reveal such controlled activities.

The micro level is concerned with the identification of nodes and their connection in the network. Nodes represent individual Twitter accounts whose interaction with each other can be shown by edges connecting them. Patterns that emerge in the micro structure of a social network give insights into its prevailing macro structure.

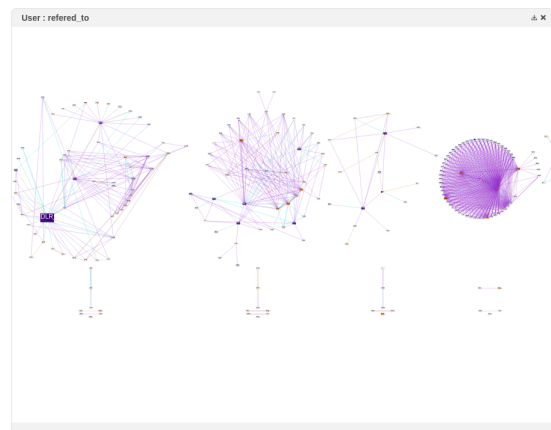


Figure 1: Reference Behavior in a Social Network

Figure 1 shows an example of referencing behavior in a social network. We see the user accounts as nodes and references (@-mention) between them as edges. The graph on the right-hand side differs notably from the other two representations. We can assume that in the case of the right-hand graph, a bot network is present, because the referential behavior of the nodes (i.e. of individual Twitter accounts) is uncommon; it seems as if the users are making references to unusually high numbers of other users. Additionally, the referenced users in such a network are often those that have

a great number of followers, e.g. celebrities. This strategy is most probably used to reach the greatest possible dispersion of a given content. It stands to reason that contents that are spread in this manner are highly likely to contain some form of misinformation. The analysis of the structures in a social network, therefore, not only reveals bots but can also point to instances of information campaigns. Thus, network analysis can support the identification of fake news as described in the previous paragraph.

5.4. Event Detection and Extraction

The information space can be considered as a source of information on the physical world. Within the information space, messages about actual and planned events are exchanged. These are to be extracted in order to enhance the situational awareness in the physical world.

We are currently carrying out work on an event detector. The aim of the event detector is to extract event information from a news stream and use this information to create entries in an event data base or to update existing event representations. An important challenge for an event detector is identity management: care must be taken to ascertain what different sources are actually reporting about. The question whether distinct events are under discussion or various messages rather report on the same event, albeit in different manners, has to be answered.

5.5. Sentiment Analysis

Sentiment analysis is a widely known approach in linguistics to determine an individual’s attitude towards an entity, e.g., an object or an event. To achieve this, text written by an individual, may it be a long statement or merely a tweet about a given topic, is analysed with regard to its polarity. The content can then be classified as positive, negative or neutral. We did not yet integrate sentiment analysis into the dashboard prototype.

Sentiment analysis can be applied to the output of the other tools introduced in this chapter. Regarding language statistics, it can for example be examined in which contexts popular hashtags are used and how they are perceived by social media users.

Events that were detected in the information space are also of interest for sentiment analysis. We assume that by determining opinions towards events, specific interest groups can be identified. These groups can have opposing standpoints. Information about the forming and existence of these groups is relevant, as it may be that conflicts in the information space might also be carried out in the physical world.

Finally, sentiment analysis can reveal the attitudes towards institutions and organisations within the theatre of operations. This comes out as crucial in external missions like humanitarian missions or peace-keeping missions where the partners depend on cooperation with the local population.

6. Prototype

An information space dashboard is a management tool which comes with a collection of various components/apps for (i) accessing the information space, i.e. collecting data,

(ii) analyzing data and (iii) visualizing analysis results. Data can be in diverse modalities, including texts in various languages, images, audio and video. At present, we only handle text data. Data analysis components should contribute to answering questions on what is happening, what is being reported (and what is not) and what will (supposedly) happen. Rather than giving an answer on one of these questions, the tools are to support the analyst in finding an answer. Visualization components serve both comprehensibility of analysis results and information exploration.

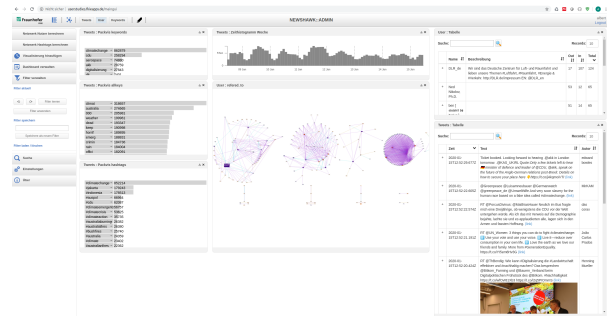


Figure 2: Dashboard Prototype

Figure 2 shows our prototype of an information space dashboard. In the preceding chapter, we have described worthwhile components, of which some have been fully implemented while others are still in a developmental state.

7. Demonstration Case

In the following, we will describe a demonstration case that deploys the above mentioned tools by reference to two self-compiled text corpora. Both corpora comprise reactions to the incident in Chemnitz, Germany at the 26th of August 2018, where a conflict at a city festival lead to the stabbing and in consequence death of one individual involved. Reports about the nationality and possible migrant status of the alleged attacker subsequently gave rise to a demonstration by right wing extremists that was accompanied by assaults against immigrants, the police and opposing demonstrators. In opposition to these developments, counter-activities took place, among them a concert against racism under the motto “Wir sind mehr” (“We are more”) a few days later. The entire situation was complex and confusing. We assume that with reference to the information space a clearer overview and better situational awareness could have been reached.

In the demonstration case, we take the view of an operator who has to elicit what is happening at present, what is reported about recent events and what is likely to happen next. As she cannot solve this task manually, she falls back on automatized processes and applies the information space dashboard. The tools of the information space dashboard should be used to analyse live data. The present demonstration case, however, exploits historical data for the purpose of illustration.

7.1. Data

One of our corpora is a Twitter corpus, the other consists of articles by the German Press Agency (dpa), cf. Table 1.

Both corpora were compiled by using the keyword “Chemnitz” as a search term. The Twitter Corpus was compiled during the time of the incidents in Chemnitz. The dpa corpus includes articles from the whole year 2018; we focus only on the ones that are concerned with the mentioned events.

We use two distinct corpora because we assume that the situation can be grasped better if different sources are factored in. It can be supposed that the language of the Twitter Corpus is far more informal and displays subjective stances, while dpa releases are written in an objective and formal style. We deemed this to not be problematic but rather an accurate depiction of the domain in which an information space dashboard would be used.

Corpus	Tokens	Types
Twitter	9413464	170073
dpa	257637	28839

Table 1: Corpus Overview

7.2. Exemplary Application of Tools

The operator firstly has to define a starting point for her task. To quickly gain initial insights, the analysis of word frequencies – both from the texts and their metadata – is applied. During this analysis, stop words are ignored in order to focus only on the tokens that are semantically relevant.

At this point, the operator detects two hashtags on Twitter that are widespread. These are #wirsindmehr (“we are more”) and #afd (“Alternative for Germany”, a German far-right political party). While the meaning of the abbreviation “afd” is known, “wirsindmehr” needs further inspection. Therefore, the operator consults an additional data pool, namely dpa data. In several articles she finds that the hashtag #wirsindmehr is the motto of a free concert that takes a stand against racism (dpa, 2018). Due to the popularity of the hashtag, and consequently of the concert, she expects that the event will be well-attended.

To get a deeper understanding of the current mindset of the population, the operator turns back to Twitter and analyses tweets that include the hashtags #wirdsindmehr and #afd with regard to their sentiment/polarity. She selects data from social media as it conveys more subjectivity. She finds that while #wirsindmehr is supposed to stand for something positive, negative sentiments are connected to it, too. This can be seen as combinations of hashtags occur like, e.g., #wirsindmehr #ihrseiderbärmlich (“we are more you are pathetic”), #wirsindmehr #ihrseidscheiße (“we are more you are shit”) and #ihrseidesnichtwert #wirsindmehr (“you are not worth it we are more”). The fact that the hashtag is used in a very positive and simultaneously an aggressive way is in indicator that (at least) two factions with greatly different opinions are forming in the population.

In a next step, the operator searches tweets about the concert that transfer negative sentiment to detect if users call for criminal acts or spread misinformation, both consciously or unconsciously, in order to substantiate their stance. A network analysis can reveal such behavior further and is applied subsequently.

The operator comes to the assessment that polarisation is rather increasing. Aggressive language lets violent incidents around the concert appear probable. Police forces are to be prepared accordingly.

8. Conclusions

We introduced the concept of an information space dashboard as a tool comprising components for data collection, data analysis and visualization. The aim of the dashboard is to support analysts in creating a common operation picture of the information space and, thus, contributing to overall situational awareness, including both the physical world and the information space. A first prototype including components for quantitative text analysis, text clustering and classification and network analysis has been created. Further components, namely for event detection and sentiment analysis, are under development. The information space dashboard is, thus, work in progress. As it will have to be adapted to changing tasks and conditions it will inherently be always work in progress: additional data source will have to be included and further analyses will have to be enabled.

Beside the development of additional components, next steps include evaluations with (potential) operators and other subject matter experts. User groups of diverse domains are to be considered: EU external civilian missions and UN missions are dependant on awareness of the situation in their theatres of operations. The same is true for the military which discusses information space operations in the context of defense against hybrid warfare. Of course, police forces have to be aware of activities in the information space – e.g., to be able to prevent hate crimes which are often announced in advance (Nagle, 2017) – but also emergency forces in order to get a better view on the situation and urgent needs, e.g., during disasters like floods or wild fires.

Observation and analysis of the information space can cause a bad taste as it is associatively linked with surveillance, censorship and suppression. Naturally, technology like an information space dashboard can be used for such ends. One means to prevent that is to make sure existing laws regarding privacy and freedom of speech are obeyed. The protection of individuals and their right to express themselves openly without fear of unwarranted consequences has a high priority, which means that not any arbitrary data source may be exploited. Furthermore, to avoid both misuse and the misunderstanding of actual, proper usage, it might be an adequate measure to make analytics transparent and provide a public overview on the information space. How this can be reached best, is still an open issue for us.

9. Bibliographical References

- Bergh, A. (2019). Message the message: Modularising software for influence operation detection in social media. In *Proceedings of the 24th International Command and Control Research Technology Symposium*, pages 1–14, Laurel, Maryland USA.
- Boyd, J. (1975). The Essence of Winning and Loosing. <https://www.danford.net/boyd/essence.htm>.

- Böttger, K. (2008). *Die häufigsten Fehler russischer Deutschlerner: Ein Handbuch für Lehrende*. Peter Lang Verlag, Münster.
- Colbaugh, R. and Glass, K. (2012). Early warning analysis for social diffusion events. *Security Informatics*, 1(1):18.
- dpa. (2018). Wie Chemnitz gegen sein Image als hässliche Stadt kämpft. <https://www.suedkurier.de/ueberregional/politik/Wie-Chemnitz-gegen-sein-Image-als-haesslich-e-Stadt-kaempft;art410924,9873202>.
- Franke, T. (2011). Social media: the frontline of cyberdefence? *NATO Review*.
- Gladrow, W. (1998). *Russisch im Spiegel des Deutschen: Eine Einführung in den russisch-deutschen und deutsch-russischen Sprachvergleich*. Peter Lang Verlag, Frankfurt am Main.
- Helbing, D., Brockmann, D., Chadefaux, T., Donnay, K., Blanke, U., Woolley-Meza, O., Moussaid, M., Johansson, A., Krause, J., Schutte, S., and Perc, M. (2014). How to Save Human Lives with Complexity Science. *European Journal for Security Research*, 4(1):51–71.
- Kwak, H., Lee, C., Park, H., and Moon, S. (2010). What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web - WWW '10*, pages 11–13, New York, New York USA.
- Lerman, K. and Hogg, T. (2012). Using stochastic models to describe and predict social dynamics of web users. *ACM Transactions on Intelligent Systems and Technologies*.
- Leskovec, J. and Lang, K. (2008). Statistical properties of community structure in large social and information networks. *Proceedings of the 17th international conference on World Wide Web. ACM*, pages 695–704.
- Li, M., Lin, Y., Hoover, J., Whitehead, S., Voss, C. R., Dehghani, M., and Ji, H. (2019). Multilingual Entity, Relation, Event and Human Value Extraction. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 110–115, Minneapolis, Minnesota USA.
- Liu, H., Qin, W., and Wan, X. (2019). An Interactive Chinese News Synthesis System. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 18–23, Minneapolis, Minnesota USA.
- Nadeem, M., Fang, W., Xu, B., Mohtarami, M., and Glass, J. (2019). FAKTA: An Automatic End-to-End Fact Checking System. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 78–83, Minneapolis, Minnesota USA.
- Nagle, A. (2017). *Kill All Normies: Online Culture Wars from 4chan and Tumblr to Trump and the Alt-Right*. Zero Books.
- Page, R. (2012). The linguistics of self-branding and micro-celebrity in twitter: The role of hashtags. *Discourse Communication*, 6(2):181–201.
- Pritzkau, A. (2019). Vertrauenswürdiger Meinungs-austausch im Kontext von Polarisierung, Desinformation.
- Qi, G.-J., Aggarwal, C. C., and Huang, T. (2012). Community detection with edge content in social media networks. In *Data Engineering (ICDE), 2012 IEEE 28th International Conference on. IEEE*, pages 534–545, Washington, DC USA.
- Schade, U., Pritzkau, A., Claeser, D., Dembach, M., and Kent, S. (2018). “Fake News” und ihre Identifikation.
- Schmitz, H.-C., Deneckere, M., de Zan, T., and Gräther, W. (2019). Situational Awareness, Information Exchange and operational Control for Civilian EU Missions. *European Journal for Security Research*, 4(1):51–71.
- Woolley, S. and Howard, P. N., (2017). *Computational Propaganda Worldwide: Executive Summary*, pages 1–14. S.C. Woolley and Philip N. Howard, Oxford, UK.