

Towards Conflictual Narrative Mechanics

Laura Spillner¹, Carlo R.M.A. Santagiustina^{2,3}, Thomas Mildner¹ and Robert Porzel¹

¹University of Bremen

²Ca'Foscari University

³Venice International University

{laura.spillner, mildner, porzel}@uni-bremen.de, carlo.santagiustina@unive.it

Abstract

We propose a five steps methodology to retrieve, reconstruct and analyse conflict related narratives in a standardized and automated way. Our methodology combines AI and network analysis techniques to build a visual representation of key agents and entities involved in a conflict and to characterize their relations. Unlike the majority of existing methods, ours can be applied to any type of conflict, as, through two data downloading phases, it first generates a bird's-eye representation and then a fine-grained map of any conflict. Given the broad applicability of the proposed methodology, we believe that this work moves the first steps towards a better understanding of conflictual narrative mechanics.

1 Introduction

Online social media have started as tools for people to connect with others all over the globe [Whiting and Williams, 2013; Given *et al.*, 2017; Mildner and Savino, 2021]. Meanwhile, platforms like Facebook, Reddit, and Twitter have been labeled 'game-changers' for entertainment [Whiting and Williams, 2013; Cunningham and Craig, 2017] while offering novel opportunities to advertisers and businesses [Forbes and Forbes, 2013; Ellwein and Noller, 2015; Mathur *et al.*, 2018]. Increasing the impact of social media even further, platforms take an important part of the everyday political discourse [Hermida *et al.*, 2012; Holt *et al.*, 2013; Rho and Mazmanian, 2020]. Twitter in particular is used not only by people who share their opinions and engage in political discussions, the platform is also used by official agencies as well as people working in governments to disseminate information more rapidly than via traditional news media.

When it comes to political and social conflicts, the surrounding online conversation is often characterized by a range of opinions. Discourse may evolve around conflict, parties involved therein, or actions undertaken while different 'causal' relationships between events and actors may be asserted in naïve manners. By being constantly exposed to these conversational dynamics, users' opinions may be influenced by the narrative that is most popular based on general interest of their personal social media bubble. This effect is further amplified by tailored algorithms that elevate content

predicted to be aligned with prior interest. This results in a concerning problem for users as well as public discourse to experience alternative points of views.

In a recent crisis, this problematic tendency became blatantly real when Russia started its invasive war against Ukraine on February 14th, 2022, escalating a fiery situation that began with Russia's annexation of the Ukrainian peninsula Crimea in 2014. Soon after the war's outbreak, on the 28th of February, 2022, the social network Twitter decided to expand labeling policies for content as a counter to the spread of misinformation on behalf of Russia, adding the 'Russia state - affiliated media' tag to related posts. Showing an impact of the spread of Russia's media outlets [Aguerra *et al.*, 2022], Twitter's decision moved Russia to block the platform in its nation, as well as Facebook who followed a similar strategy. Both interventions, of Twitter and of Russia, highlight how vulnerable public speech and information can become. In this particular case, Russia's government decided to cut off their country from major western media outlets, allowing them to precisely control news available to Russian's citizens.

Resulting conflictual narratives, occurring during crises like the one mentioned above, present an urgent need to understand underlying mechanics. In this paper, we present a methodology, as a work in progress, to study and investigate the online narratives surrounding conflicts and crisis. The methodology itself is not necessarily limited to conflict alone, but aims to discover different perspectives on social media while limiting any introduced researcher bias when constructing the corpus itself. This paper describes the idea behind the methodology. We aim to use the methodology to analyse Twitter conversations of Russia's invasive war against Ukraine, and we present here the results of the first (preliminary) part of this analysis.

Our hypothesis is that conflict narratives are strategically designed around recurrent story-telling patterns and frames that assign a set of (asymmetric and stereotyped) roles to involved parties. We plan to ultimately utilize the methodology presented herein to investigate this hypothesis on the case-study of the Russo-Ukraine war; the results of this investigation will be presented in future work.

When trying to capture the context and viewpoint of narratives, one of the main difficulties is not to introduce researcher bias to sampled data: By collecting, for example, only posts written in English, by defining which terms to

query for, which conflict-related narrative frames are being looked up, or which actors in the conflict are of interest apart from the main parties, any resulting corpus represents only a certain fraction of the discourse.

Therefore, we define a two-phase approach to investigate conflict narratives based on online conversations on Twitter:

1. A dataset is collected of tweets that explicitly mention the conflict at hand through the names of the main parties or neutral conflict-related terms. From this, the entities involved in the conflict, such as actors, events, and locations, either actively or passively, are discovered and the most important entities and their co-occurrence are identified.
2. More fine-grained data is then collected by searching for all the main entities and capturing frames used to characterize these entities as well as possible relationships between them.

So far, we have conducted the first phase of the analysis on the conflict between Russia and Ukraine, as characterized by a sample of the Twitter conversation in the spring of 2022. In order to explain the proposed methodology, we present the results of this first step, detail the steps planned for the second phase analysis, and discuss which challenges remain to be solved in order to expand the methodology to a greater scope and other case studies.

2 Related Work

The following section presents a brief summary of recent literature on conflict-related narratives with a particular focus on how these narratives are constructed and diffused through online social media.

2.1 Conflict narratives

Conflicts and their escalation in the physical and narrative space are generally the by-product of heterogeneous beliefs [Millner and Ollivier, 2020], asymmetric information [Smith and Stam, 2004], cognitive biases [Erb *et al.*, 1998], like the availability heuristic and the confirmation bias, as well as complex entanglements [Acton, 2018; Austin, 2003] of competing interests, strategies, and objectives, which are often opaque being difficult to elicit and model.

Conflict narratives, and more in general debates about conflicts through which these narratives evolve and spread, have long been studied in the social and political sciences under many different approaches and perspectives [Ober, 2000; Uluğ *et al.*, 2021; Anderlini *et al.*, 2001]. For example, the study of conflict framing [van der Goot *et al.*, 2022], factional discourse design [Wiant, 2002], and other polarizing communication strategies [Somer *et al.*, 2021] is key for understanding any consensus-building and group-mobilization process when there are competing views or interests at stake.

The narrative dimensions of conflicts are generally more visible in non-authoritarian countries, where partisan narratives [Polsky, 2012] are constructed and employed in relation to public support-building and policy justification. Conflict narratives operate at all levels of national and international governance processes. Especially democratic coun-

tries aim to mitigate and resolve potential or actual conflicts through transparent, informed, and participatory deliberative processes. From local debates about public policies to be implemented in response to a pandemic, like COVID-19, to the renegotiation of trade agreements between parties that compete for the control on strategic natural resources and technologies, like conflicts for rare earths or for the control over semi-conductor industry technologies, the emergence of conflict narratives is an ubiquitous phenomenon in contemporary times. This is particularly evident in decentralized and multi-directional online communication mediums such as social media [Abdel-Fadil, 2019; de Graaf and Meijer, 2019], which have become the default propagation medium for (popular) narratives, including conflict-related ones.

2.2 Conflict modeling

Recent attempts to model conflicts, like [Charap and Colton, 2016], have shown that conflicts are not necessarily the outcome of diverging material interests among individuals and groups, as modeled in early game-theoretic works [Dresher, 1956]. They can also be the product of: (i) differently-biased or competing world views used to decipher events and to comprehend the intentions underlying specific actions or communications by other individuals and groups [Camerer, 1997]; or (ii) to the deterioration of inter-group trust [Rohner *et al.*, 2013]. As a result, conflicts may emerge and exacerbate even when the material interests of the different parties converge from a rational (i.e., utilitarian) perspective. For this reason, the presence of extrinsic incentives [Hoover Green, 2016] may not suffice to mitigate or resolve ongoing conflicts grounded on incompatible belief-systems or on the lack of trust, like during the cold-war [Deighton, 1993]. For example, this is the case in attrition wars, which may be represented as negative-sum games. Rather than immediately negotiating a mutually advantageous agreement, two parties are ready to bear the material and humanitarian costs of a long-lasting conflict if that gives them the opportunity to punish the other, by reciprocating harm and causing a similar or larger damage to their opponents.

2.3 Conflict-related narratives extraction and events re-construction

The incessant growth of online social media and their communities [Can and Alatas, 2019], together with the increasing availability of computational power and advanced linguistic analysis methods [Willaert *et al.*, 2021b; Beuls *et al.*, 2021; Asghar, 2016] for big textual datasets offers an opportunity to capture and model conflict narrative dynamics on an unprecedented scale.

Recent works [van Schalkwyk *et al.*, 2020; Santagiustina and Warglien, 2022 forthcoming; Willaert *et al.*, 2021a] have used posting activities on social media and online newspaper articles, including comment sections, for capturing and mapping partisan or faction-specific arguments and their dynamics across time through a combination of AI and network analysis methods. A further step towards the automated mapping of conflicts and of their key actors and events has been done by [Gottschalk and Demidova, 2019] and [Gottschalk *et*

al., 2021] who combined large-scale knowledge graphs with semi-structured sources in an event KG RDF-representation.

The former branches of research open the way to a new AI-augmented research field on all kinds of conflicts. Such research could combine the potentials of NLP, network analysis, and computational linguistic methods with the semantic web, serving as interfaces for the real-time observation and understanding of conflict-related narratives.

3 Methodology

To analyze the narratives that surround certain conflicts, we developed a methodology consisting of two phases which include two cycles of data collection. The reasons behind this are twofold: Firstly, it allows the discovering of actors and entities of interest that are involved in the conflict dynamically from the online conversation surrounding it, instead of defining them a-priori. Thus, we hope to reduce the bias that would otherwise be introduced during the corpus construction. Secondly, this also makes it possible to employ the same methodology to investigate extremely different topics and types of conflicts.

Although in this paper, we focus on the Russian war against Ukraine, the aim for this approach is to be independent of any specific topic in order to enable other works to also reconstruct narratives from different crisis, whether violent or non-violent, where the parties involved are not clearly defined, may change over time, or where the definition of parties depends on specific viewpoints.

In the following sub-sections, we present how we are currently using this two phase approach to identify which narratives are shaping the online conversation on Titter concerning the ongoing war in Ukraine. Phase 1 describes the analysis that has been conducted so far as well as the results of this, while Phase 2 discusses the steps that we plan to conduct next.

3.1 Phase 1 - Discovery

In this first discovery phase, the goal is to collect tweets from a specified time interval which explicitly mention the conflict in question. We thereby collect any tweet that meets two conditions: Tweets that (1) either contain the name of at least one of the directly involved parties or the name of a specific event and that (2) contain a generic term like conflict, tension, or crises, denoting that the tweet refers to a conflict related to the selected event or involved parties. From these tweets, we identify all actors and entities involved in the conflict which are mentioned most often in this context. We are then able to construct a network based on their co-occurrences.

Step 1: Corpus construction

To construct the first corpus on the current war in Ukraine, we collected all tweets

- between the first of January and the first of May 2022
- which contain the words: “Ukraine”, “Russia”, and any of “conflict(s)”, “tension(s)”, “crises” or “crisis”
- which are not retweets.

As of now, we have only conducted a preliminary trial in order to test and validate our methodology. For this, we

have only searched for these terms in English. Hence, the collected corpus only represents a perspective from English-speaking Twitter users. However, we want to also collect tweets containing these key terms in other languages, including Ukrainian and Russian, in order to generate a multilingual corpus that represents a broader view and contains differing narratives that include those from both of the opposing parties.

An advantage of this two phase approach is that at the beginning, only a small number of terms needs to be defined, which are compatible with any type of conflict. We plan to employ the help of native speakers of other languages to translate these terms and validate the results of the first phase. However, we hope that we will not require the help of experts in the domain which are also native speakers of the additional languages. Apart from the definition of this small set of neutral conflict-related starting terms, the collection of the data (across both phases) is entirely automatic. We chose to go with only these three terms (conflict, tension and crisis) firstly in order to limit the scope of the twitter query, and secondly because we believe that other synonyms of those terms generally connote either violent or non-violent conflicts (e.g. synonyms of “conflict” on ConceptNet include terms such as “battle” or “disagreement”).

For this preliminary analysis, we collected a total of 724400 tweets. In a first pre-processing step, we removed special UTF-8 characters, like emojis, emoticons, and URL links.

Step 2: Entity Recognition

After generating this dataset, we use this broad-coverage overview of the conflict at hand to discover from it which actors or other entities are relevant to the conflict. These will constitute the terms which we will explicitly teach for in the second phase. Of interest are

- named entities such as known persons, organizations, countries or peoples, e.g. in this case “Putin”, “Ukraine”, “EU”, “Russians”.
- noun phrases which include these same terms as signifiers, such as “Ukrainian president”, “russian army”, “Ukraine war”, “russia-ukraine conflict”, etc.

Therefore, instead of defining a term for the conflict and thus characterizing it ourselves – based on the different connotations of “war”, “unrest”, “invasion”, “operation”, etc. – we are able to discover which terms are used in online conversations. The same holds for other entities involved, e.g. a person being referred to as a “president” or “dictator”. Either variant will be collected if it appears in a noun phrase together with one of the named entities.

In this step, we use the open source NLP library spaCy for part-of-speech tagging, dependency parsing, and named entity recognition. From this, we identify all noun phrases which include a token that was classified as a named entity. Such identified entities would include only geopolitical entities such as countries, nationalities or religious/political groups, organizations, persons, events, and locations. However, these entities would exclude instances such as cardinals.



Figure 1: Entities network for the Ukraine-Russia conflict. Network filtered by node weighted degree: 0.999 percentile; and then by edge weight: 0.99 percentile.

Identified noun phrases can then be consolidated in entities, by linking them to named entities in a known knowledge graph such as WikiData (consolidating e.g. “the dictator Putin”, “Vladimir Putin” and “the russian president”, etc.). For the preliminary analysis that follows, the consolidation step has yet to be implemented. In future versions of this work this will be done through name matching and name substring matching, without disambiguation, or through more advanced distributional semantics and ML methods, like OpenTapioca [Delpeuch, 2019].

Step 3: Co-occurrence network

In the third step, identified entities are mapped through an undirected network based on their co-occurrences in the corpus. Each node represents an entity, with edge weights denoting how often two entities occur together in one tweet in the dataset. By so doing we obtain a weighted and undirected network containing 933081 nodes and 10871807 edges. With an average degree of nodes equal to 23.303 and an average weighted degree of nodes equal to 46.162.

Network nodes are then filtered based on their weighted degree centrality metric; alternative centrality measures, like pagerank, betweenness or eigen-centrality may also be used for this purpose. This filtration is done in order to remove entities that are less influential (more peripheral) and have a marginal role in the entity network for the selected conflict. As shown in Figure 2, this filtration step removes those entities which are mentioned rarely in the dataset, like entities weekly related to the conflict and other irrelevant noun phrases. This might include entities such as sports teams or their fans who might be involved in a metaphorical conflict,

misspellings, or other noun phrases that do not play a relevant role in online conflict-related narratives.

3.2 Phase 2 - Analysis of Narratives

In the second phase, we plan to retrieve thinner grained data about the conflict on the selected social media, in order to (re-)construct and analyze the narratives surrounding specific relational blocks. As a starting point for this will serve the key actors/entities and relations (dyads of actors), which were identified in Phase 1 based on the chosen centrality metric, together with the largest/heaviest cliques of order 3+.

Step 4: Second corpus construction

At this point, we have identified a number of key entities and relations, as well as different reference terms which are used to refer to them. These will then be used to construct a new set of queries from which a second corpus of tweets is collected. In this second phase, we are going to search explicitly for the entities that the first phase identified as being perceived as important actors in the conflict, using the terms that were discovered to be used to refer to them, by twitter users.

While the first corpus included only tweets that reference the conflict itself, thus providing a broader view of the on-line conversation surrounding it, this second corpus will include tweets referring specifically to one or more of the relevant entities and allow for a more fine-grained analysis. By collecting these tweets, we aim to identify how the entities themselves, as well as the relationships between them, are characterized by different people online, for example through adjectives and verbs qualifying the relation between two key entities.

Our hypothesis is that the narratives surrounding the actors involved in the conflict are based on recurring phrase fragment patterns and frames that assign specific roles and attributes to the involved parties. In many cases, these roles are a-symmetric and mirror certain stereotypes that are common to (almost) all conflict narratives, like the role of the victim and that of the perpetrator. We plan to use the methodology we describe herein to investigate this hypothesis on the corpus we are currently collecting about the war in Ukraine. Using the tweets from this Phase 2 corpus, we aim to collect and assign a number of different frames to the actors and relations, which are essential constituents of the conflict narrative. For example, we expect that we will find

- verbs related to asymmetric roles in the conflict, like: aggression / protection, offense / defense, attack / counter-attack, ownership claims, “deserviness” claims, resisting/ surrendering, etc.
- conflict related nouns and adjectives, like: aggressor / aggressed, invader / invaded, liberator / liberated, oppressor / oppressed, strong / weak, winning / losing, perpetrator / victim, etc.
- characterizations of the conflict or its escalation, like: justified / unjustified, legitimate / illegitimate, necessary / unnecessary, explainable / unexplainable, expected / unexpected, hot / cold, violent / non-violent, verbal / physical, ideological, political, economic, financial, military, etc.

- equivalently, characterizations of
 - the peoples or populations involved in the conflict
 - the leaders of the factions involved in the conflict
 - the countries involved in the conflict
 - the factions, armies or soldiers involved in the conflict
 - etc.

Step 5 - Analysis

From the entities, relations and frames used to characterize them, which are collected in the previous step, a second, fine-grained and dynamic network will be constructed. This network, which could be based of RDF-star or a labelled property graph, is an in-depth representation of the narratives describing the relations between the key actors and their characteristics. For this step, we plan to again employ the part-of-speech recognition and dependency tree identified using the spaCy toolkit to reconstruct this network.

4 Discussion

The network we have already constructed from the Phase 1 dataset shows an interesting bird's-eye perspective on the on-line narratives surrounding the selected conflict, that is, the war between Ukraine and Russia. However, this pipeline for automatically reconstructing conflict narratives is still at a very early development stage. Open questions remain concerning the advantages and disadvantages of the centrality measures and AI methods to be used at each stage of the process, for instance, filtering the actors network. A future aim is to explore and benchmark these alternative procedures and metrics, as well as their impact on the results, throughout future works, employing multiple conflict datasets. Here follows a brief discussion about the criticalities that we have identified at the current state of the work and that we will address as a priority in the next development stages.

Firstly, we plan to evaluate the most useful way to construct the co-occurrence network. Currently, the network is based on the counts of co-occurrences of identified entities and noun phrases in tweets. More precisely, the more often two entities appear in the same tweet, the higher the weight of that edge is. An alternative approach is to take the frequency of any two connected nodes into consideration when having to weigh the relation between two entities (as more frequently occurring entities are also more likely to appear together by mere chance). This metric would also allow us to differentiate between pairs of nodes that appear together more often than random, and pairs of nodes that appear together less often than random. Similar approaches have been used in other fields of studies, for example, for species probabilistic co-occurrence analysis. In connection with this, we are also considering a number of different options to calculate the centrality of a node, which is currently based on its degree.

The second part which we will further investigate is the representation of conflict dynamics over time. The current network is based on the entire set of tweets collected from a defined time frame. During the next iteration, we plan to firstly collect tweets from a larger time interval and secondly

to slice the network into smaller time-chunks. This will enable us to visualize and analyze how the network changes over time. We expect an interesting perspective to be added to both the co-occurrence network of the first phase as well as to the more fine-grained network.

Thirdly and finally, we plan for the next iteration of this dataset to analyze more of the context data surrounding the tweets themselves: Many tweets include metadata about the country of origin, time zone, location of users, and their language, signifying emojis, flags and hashtags, as well as the metadata connected to the user's account. Considering users' biographies or geolocation might make it possible to look at which narratives are more prominent among which groups of users or locations. This is also connected to our previous goal of adding more languages to the corpus.

Acknowledgments

This work was funded by the by the FET-Open Project #951846 "MUHAI – Meaning and Understanding for Human-centric AI" by the EU Pathfinder and Horizon 2020 Program and the "Empowering Digital Media" grant from the Klaus Tschira Foundation.

References

- [Abdel-Fadil, 2019] Mona Abdel-Fadil. The politics of affect: The glue of religious and identity conflicts in social media. *Journal of Religion, Media and Digital Culture*, 8(1):11–34, 2019.
- [Acton, 2018] James M Acton. Escalation through entanglement: How the vulnerability of command-and-control systems raises the risks of an inadvertent nuclear war. *International security*, 43(1):56–99, 2018.
- [Aguerri *et al.*, 2022] Jesús Aguerri, Mario Santisteban, and Fernando Miró-Llinares. The fight against disinformation and its consequences: Measuring the impact of "russia state-affiliated media" on twitter. *OSF Preprints*, 2022.
- [Anderlini *et al.*, 2001] Sanam Naraghi Anderlini, Vasu Gounden, Ed Garcia, Barbara Harff, Maha Khan, Deepa Khosla, Michael S Lund, David Nyheim, Luis Alberto Padilla, Thomas Schmalberger, et al. *Journeys through conflict: Narratives and lessons*. Rowman & Littlefield, 2001.
- [Asghar, 2016] Nabihha Asghar. Automatic extraction of causal relations from natural language texts: a comprehensive survey. *arXiv preprint arXiv:1605.07895*, 2016.
- [Austin, 2003] Greg Austin. Unwanted entanglement: the philippines' spratly policy as a case study in conflict enhancement? *Security Dialogue*, 34(1):41–54, 2003.
- [Beuls *et al.*, 2021] Katrien Beuls, Paul Van Eecke, and Vanja Sophie Cangalovic. A computational construction grammar approach to semantic frame extraction. *Linguistics Vanguard*, 7(1), 2021.
- [Camerer, 1997] Colin F Camerer. Progress in behavioral game theory. *Journal of economic perspectives*, 11(4):167–188, 1997.

- [Can and Alatas, 2019] Umit Can and Bilal Alatas. A new direction in social network analysis: Online social network analysis problems and applications. *Physica A: Statistical Mechanics and its Applications*, 535:122372, 2019.
- [Charap and Colton, 2016] Samuel Charap and Timothy J Colton. The negative-sum game and how to move past it. *Adelphi Series*, 56(460):151–184, 2016.
- [Cunningham and Craig, 2017] Stuart Cunningham and David Craig. Being ‘really real’ on YouTube: authenticity, community and brand culture in social media entertainment. *Media International Australia*, 164(1):71–81, August 2017. Publisher: SAGE Publications Ltd.
- [de Graaf and Meijer, 2019] Gjalt de Graaf and Albert Meijer. Social media and value conflicts: An explorative study of the dutch police. *Public Administration Review*, 79(1):82–92, 2019.
- [Deighton, 1993] Anne Deighton. *The Impossible Peace: Britain, the Division of Germany, and the Origins of the Cold War*. Oxford University Press, 1993.
- [Delpeuch, 2019] Antonin Delpeuch. Opentapioca: Lightweight entity linking for wikidata, 2019.
- [Dresher, 1956] Melvin Dresher. Theory of games of strategy. Technical report, RAND CORP SANTA MONICA CA, 1956.
- [Ellwein and Noller, 2015] Carsten Ellwein and Benedikt Noller. Social Media Mining: Impact of the Business Model and Privacy Settings. In *Proceedings of the 1st ACM Workshop on Social Media World Sensors*, Side-WayS ’15, pages 3–8, New York, NY, USA, September 2015. Association for Computing Machinery.
- [Erb *et al.*, 1998] Hans-Peter Erb, Gerd Bohner, Knut Schmilzle, and Susanne Rank. Beyond conflict and discrepancy: Cognitive bias in minority and majority influence. *Personality and Social Psychology Bulletin*, 24(6):620–633, 1998.
- [Forbes and Forbes, 2013] Lukas P. Forbes and Lukas P. Forbes. Does Social Media Influence Consumer Buying Behavior? An Investigation Of Recommendations And Purchases. *Journal of Business & Economics Research (JBER)*, 11(2):107–112, January 2013. Number: 2.
- [Given *et al.*, 2017] Lisa M. Given, Denise Cantrell Winkler, and Kathryn Hopps-Wallis. Social Media for Social Good: A Study of Experiences and Opportunities in Rural Australia. In *Proceedings of the 8th International Conference on Social Media & Society*, #SMSociety17, pages 1–7, New York, NY, USA, July 2017. Association for Computing Machinery.
- [Gottschalk and Demidova, 2019] Simon Gottschalk and Elena Demidova. Eventkg—the hub of event knowledge on the web—and biographical timeline generation. *Semantic Web*, 10(6):1039–1070, 2019.
- [Gottschalk *et al.*, 2021] Simon Gottschalk, Endri Kacupaj, Sara Abdollahi, Diego Alves, Gabriel Amaral, Elisavet Koutsiana, Tin Kuculo, Daniela Major, Caio Mello, Gul-lal S Cheema, et al. Oekg: The open event knowledge graph. In *CLEOPATRA@ WWW*, pages 61–75, 2021.
- [Hermida *et al.*, 2012] Alfred Hermida, Fred Fletcher, Darryl Korell, and Donna Logan. Share, Like, Recommend. *Journalism Studies*, 13(5-6):815–824, October 2012. Publisher: Routledge eprint: <https://doi.org/10.1080/1461670X.2012.664430>.
- [Holt *et al.*, 2013] Kristoffer Holt, Adam Shehata, Jesper Strömbäck, and Elisabet Ljungberg. Age and the effects of news media attention and social media use on political interest and participation: Do social media function as leveller? *European Journal of Communication*, 28(1):19–34, February 2013. Publisher: SAGE Publications Ltd.
- [Hoover Green, 2016] Amelia Hoover Green. The commander’s dilemma: Creating and controlling armed group violence. *Journal of Peace Research*, 53(5):619–632, 2016.
- [Mathur *et al.*, 2018] Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–26, November 2018. arXiv: 1809.00620.
- [Mildner and Savino, 2021] Thomas Mildner and Gian-Luca Savino. Ethical User Interfaces: Exploring the Effects of Dark Patterns on Facebook. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7. Association for Computing Machinery, New York, NY, USA, May 2021.
- [Millner and Ollivier, 2020] Antony Millner and Hélène Ollivier. Beliefs, politics, and environmental policy. *Review of Environmental Economics and Policy*, 2020.
- [Ober, 2000] Josiah Ober. Political conflicts, political debates, and political thought. *The Shorter Oxford History of Europe I: Classical Greece*, pages 111–38, 2000.
- [Polsky, 2012] Andrew J Polsky. Partisan regimes in american politics. *Polity*, 44(1):51–80, 2012.
- [Rho and Mazmanian, 2020] Eugenia Ha Rim Rho and Melissa Mazmanian. Political Hashtags & the Lost Art of Democratic Discourse. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, pages 1–13, New York, NY, USA, April 2020. Association for Computing Machinery.
- [Rohner *et al.*, 2013] Dominic Rohner, Mathias Thoenig, and Fabrizio Zilibotti. War signals: A theory of trade, trust, and conflict. *Review of Economic Studies*, 80(3):1114–1147, 2013.
- [Santagiustina and Warglien, 2022 forthcoming] Carlo Romano Marcello Alessandro Santagiustina and Massimo Warglien. The architecture of partisan debates: The online controversy on the no-deal brexit. *PLoS one*, 2022 (forthcoming).
- [Smith and Stam, 2004] Alastair Smith and Allan C Stam. Bargaining and the nature of war. *Journal of Conflict Resolution*, 48(6):783–813, 2004.

- [Somer *et al.*, 2021] Murat Somer, Jennifer L McCoy, and Russell E Luke. Pernicious polarization, autocratization and opposition strategies. *Democratization*, 28(5):929–948, 2021.
- [Uluđ *et al.*, 2021] Özden Melis Uluđ, Brian Lickel, Bernhard Leidner, and Gilad Hirschberger. How do conflict narratives shape conflict-and peace-related outcomes among majority group members? the role of competitive victimhood in intractable conflicts. *Group Processes & Intergroup Relations*, 24(5):797–814, 2021.
- [van der Goot *et al.*, 2022] Emma van der Goot, Sanne Kruikemeier, Jeroen de Ridder, and Rens Vliegthart. Online and offline battles: Usage of different political conflict frames. *The International Journal of Press/Politics*, page 19401612221096633, 2022.
- [van Schalkwyk *et al.*, 2020] Francois van Schalkwyk, Jonathan Dudek, and Rodrigo Costas. Communities of shared interests and cognitive bridges: The case of the anti-vaccination movement on twitter. *Scientometrics*, 125(2):1499–1516, 2020.
- [Whiting and Williams, 2013] Anita Whiting and David Williams. Why people use social media: a uses and gratifications approach. *Qualitative Market Research: An International Journal*, 16(4):362–369, August 2013.
- [Wiant, 2002] Fredel M Wiant. Exploiting factional discourse: Wedge issues in contemporary american political campaigns. *Southern Journal of Communication*, 67(3):276–289, 2002.
- [Willaert *et al.*, 2021a] Tom Willaert, Sven Banisch, Paul Van Eecke, and Katrien Beuls. Tracking causal relations in the news: data, tools, and models for the analysis of argumentative statements in online media. *Digital Scholarship in the Humanities*, 2021.
- [Willaert *et al.*, 2021b] Tom Willaert, Paul Van Eecke, Jeroen Van Soest, and Katrien Beuls. An opinion facilitator for online news media. *Frontiers in big Data*, 4:46, 2021.