

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Tracking causal relations in the news

Willaert, Tom; Banisch, Sven; Van Eecke, Paul; Beuls, Katrien

Published in: Digital Scholarship in the Humanities

DOI: 10.1093/IIc/fqab107

Publication date: 2022

Document Version Other version

Link to publication

Citation for pulished version (HARVARD):

Willaert, T, Banisch, S, Van Eecke, P & Beuls, K 2022, 'Tracking causal relations in the news: data, tools, and models for the analysis of argumentative statements in online media', *Digital Scholarship in the Humanities*, vol. 37, no. 4, pp. 1358-1375. https://doi.org/10.1093/llc/fqab107

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
 You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Tracking causal relations in the news: data, tools, and models for the analysis of argumentative statements in online media

Tom Willaert D Vrije Universiteit Brussel, Brussels, Belgium

Sven Banisch

Max Planck Institute for Mathematics in the Sciences (MPI-MiS), Leipzig, Germany and Karlsruhe Institute of Technology, Institute of Technology Futures (KIT-ITZ), Karlsruhe, Germany

Paul Van Eecke

Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Brussels, Belgium

Katrien Beuls

Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Brussels, Belgium

Abstract

Online debates and debate spheres challenge our assumptions about democracy, politics, journalism, trust, and truth in ways that make them a necessary object of study. In the present article, we argue that the study of online arguments can benefit from an interdisciplinary approach that combines computational methods for text analysis with conceptual models of opinion dynamics. The article thereby seeks to make a conceptual and methodological contribution to the field by highlighting the role of domain-crossing causal statements in debates of societal interest, and by providing a method for automatically mining such statements from textual corpora on the web. The article illustrates the relevance of this approach for the study of online debates by means of a case study in which we analyse cross-cutting statements on climate change and energy technologies from the comment section of the online newspaper *The Guardian*. In support of this case study, we use data and methods that are made openly available through the Penelope ecosystem of tools and techniques for computational social science.

Correspondence: Tom Willaert, Vrije

Universiteit Brussel, Brussels, Belgium. **E-mail**: tom.willaert@vub.be

1 Background

Over the past decades, the rise of social media and the digitization of news and discussion platforms have radically transformed how individuals and groups communicate, organize themselves, and express their beliefs and opinions (Singer and Brooking, 2018; Sunstein, 2018). As Alan Rusbridger, former editorin-chief of the newspaper The Guardian has it, emerging technologies and platforms have 'empower[ed] people who were never heard, creating a new form of politics and turning traditional news corporations inside out' (Rusbridger, 2018, pp. xx-xxi). Online debates and debate spheres thus challenge our assumptions about democracy, politics, journalism, trust, and truth in a way that make them a necessary object of study. Yet precisely because they subvert existing assumptions and epistemologies, the study of online opinion dynamics is a challenging endeavour, both on a conceptual and a practical level. Indeed, any effort to empirically analyse online debates raises interconnected questions of how and where to obtain relevant data, how to discern the different voices and argumentative statements expressed in those data, how to understand the dynamics that govern opinions, and, not in the least, which methods and instruments might be applied to the task of mining these communicative phenomena.

Depending on the affordances of the platform on which they figure, opinions or beliefs can be explicitly or implicitly expressed through texts, video, sound, pictures, upvotes, downvotes, shares, likes, dislikes, emojis, or any combination of these. However, as written words still fill most of the comment sections of news websites, blogs and discussion fora, the present will be primarily concerned with the analysis of textual data, focusing mainly on texts in English. The value of such texts for the scientific study of the social world has become readily apparent (Rogers 2013, 2019; Watts, 2013), resulting in a rich ecosystem of research instruments and infrastructures through which online (social) media texts can be mined and analysed. Examples include APIs, software libraries and other instruments for accessing web content (e.g. Reddit API, 2020; Twitter API, 2020), research tools for capturing, filtering and analysing web data (e.g. Borra and Rieder, 2014; Peeters and Hagen, 2018;

DMI tools, 2020), thematic social media observatories on topics of societal relevance (e.g. Projet Politoscope, 2017; Chavalarias and Panahi, 2018), and databases and archives of social media data (e.g. Baumgartner *et al.*, 2020). These efforts face the challenge of making insightful those web data that accumulate at volumes and speeds exceeding human faculties for reading, thus foregrounding a persistent need for methods and technologies that facilitate interpretative approaches to large collections of online texts.

In the present article, we argue that the study of online arguments can benefit from an interdisciplinary approach that combines computational methods for text analysis with conceptual models of opinion dynamics, notably cognitive maps. The article thereby seeks to make a conceptual and methodological contribution to the field by highlighting the role of domain-crossing causal statements in debates of societal interest (i.e. statements that bridge multiple discursive spheres), and by providing a method for automatically mining such statements from textual corpora on the web. The article illustrates the relevance of this approach for the study of online debates by means of a case study in which we analyse crosscutting statements on climate change and energy technologies from the comment section of the online newspaper The Guardian. This results in an exploratory, mesoscopic view of the debate that prioritizes the overall position, range, and relatedness of argumentative statements, and allows us to gauge the public opinion space on technologies including nuclear power, coal, and renewables such as solar power and wind energy. In support of this case study, the article uses data and methods that are made openly available through the Penelope ecosystem of tools and techniques for computational social science (Penelope, 2020; Willaert et al., 2020), to which we provide links for each step in our method.

Within the scope of this work, we identify two main fields of research as particularly relevant to the task of automatically mapping argumentative statements in online debates: natural language processing (NLP), notably its subfields text mining and information extraction, and concepts from opinion dynamics, notably the creation of cognitive maps and insights from argument communication theory. In the remainder of this background section, we present a succinct (and therefore inevitably limited) introduction to both fields for the purpose of highlighting the open challenges that need to be addressed if we aim to operationalize and combine aspects of these fields for the study of online opinion landscapes.

1.1 Challenges related to text mining and information extraction

One active area of research that is of particular interest to those seeking to interpret bodies of textual data, such as news website comments or other (social) media texts, concerns the analysis of texts using methods from computational linguistics and NLP. Advances in this field have, for instance, yielded instruments and methods for the automatic extraction and analysis of sentiments (understood as a person's emotional stance towards a statement) (Bakshi et al., 2016; Pak and Paroubek, 2010; Wang and Manning, 2012), argumentative claims (Farzindar et al., 2017; Stede et al., 2018), moral rhetoric (Sagi and Dehghani, 2014), topics (Maier et al., 2018), and narratives (e.g. Mani, 2013). Computational techniques for the identification and statistical analysis of patterns in texts have moved to the centre of emerging reading paradigms in the social sciences and (digital) humanities (Burdick et al., 2012). This includes 'culturomics' or the quantitative analysis of cultural trends based on large textual corpora such as Google Books, and 'distant reading' or the machineguided interpretation of literary and historical texts (Eve, 2019; Underwood, 2019). Similar data-driven paradigms for studying texts at scale are also taking up an important position in the field of journalism, as exemplified by recent investigations into the Paradise and Panama papers (ICIJ, 2020). Arguably, in order to explore argumentative statements and online opinion landscapes expressed in texts on the web, two open challenges related to text mining should be addressed. For one thing, we need a conceptual understanding and definition of what constitutes an expression of an argumentative statement, opinion or belief. For another, we need to associate this definition with a linguistic pattern that might be extracted from texts. Furthermore, the outcomes of any such analysis hinge on additional methods for, among others, counting, clustering, filtering, and comparing the retrieved textual patterns. These challenges cannot be met in a vacuum, as they require extensive domain knowledge of the debates and media at hand, explorations of the

actual texts, as well as a dialogue with theoretical frameworks of opinion dynamics.

1.2 Challenges related to cognitive maps and argument communication theory

Opinions and beliefs mined from texts can be rendered insightful through models and (visual) representation schemes that map the systemic relations between these opinions and beliefs. In this article, we will focus on mining expressions of causations from texts in order to build up 'cognitive maps', or visual representations of the cognitive organization of beliefs. Within the social sciences, the study of expressions of causation as a method for accessing and assessing belief systems and opinions has been formalized and streamlined by political scientist Robert Axelrod and others since the 1970s. Axelrod's method of 'causal mapping' (also referred to as 'cognitive mapping') was thereby introduced as a means of reconstructing and evaluating administrative and political decision-making processes, based on the principle that choices are often evaluated based on their potential consequences (Axelrod, 2016, p. 5). Concretely, Axelrod's causal mapping method comprises a set of conventions to graphically represent networks of causes and effects (the nodes in a network) as expressed in texts, as well as the qualitative aspects of these causal relations (the networks directed edges, notably assertions of whether the causal linkage is positive or negative). As we will demonstrate for the case of the climate change debate, the resulting graphs or maps can thus provide a visual, structural overview of the relations among causal assertions.

The effort of mapping online debates and opinion landscapes, however, extends well beyond the social sciences and humanities. Physicists and mathematicians working within the field of complex systems science have elaborated conceptual models of the mechanisms that govern the formation of beliefs and opinions (Banisch and Olbrich, 2019; Castellano *et al.*, 2009; Friedkin and Johnsen, 2011; Hegselmann and Krause, 2002). This article therefore also engages with the concept of cognitive maps as fleshed out in the area of theoretical modelling of opinion dynamics (Banisch and Olbrich, 2021; Friedkin *et al.*, 2016; Schröder and Wolf, 2017; Van Overwalle and Heylighen, 2006). Taking into account the cognitive

Tracking Causal Relations in the News

organization of beliefs and attitudes, these models allow us to incorporate meaningful statements and to seek a correspondence with actual debates such as the evolution of attitudes on the Iraq war (Friedkin et al., 2016) or on sustainable urban transportation (Wolf et al., 2015). In particular, this article interacts with the framework of argument communication theory (Banisch and Olbrich, 2021; Mäs et al., 2013; Mäs and Flache, 2013). Argument communication theory seeks to explain collective processes of opinion formation based on the general idea that individuals exchange arguments about certain issues under discussion. Arguments can support a positive attitude towards an issue (pro-argument) or a negative one (contra-argument), and the adoption of an argument shifts the opinion into the respective direction. Such a mapping of argumentative statements to attitudes on different issues of a debate can be seen as a specific type of cognitive map and has been referred to as a cognitive-evaluative map in Banisch and Olbrich (2021). Argument communication theory has generated useful insights on how argument exchange may shape general opinion trends such as polarization (Mäs and Flache, 2013) and ideological alignment of opinions on multiple issues (Banisch and Olbrich, 2021). Recent experimental work indicates that argument communication theory has the potential to provide a useful perspective on actual debates, such as public opinion on different technologies for energy production (Shamon et al., 2019).

It should be emphasized that linking theoretical frameworks such as causal mapping and argumentcommunication theory to authentic textual data is a non-trivial task. This is definitely the case when one aims to achieve this using openly available and reusable research instruments. In the case study that follows, we propose an approach to this challenging problem by applying different methods to the study of online debates on climate change expressed in the comment section of The Guardian. We thereby scrape opinion data from comments, from which we then automatically extract causation frames using a method from computational linguistics. We then proceed to provide a basic frequency analysis of common causes and effects in the debate, and we visualize these as a causal map, thus providing a first macroscopic view of the debate. Next, we move beyond causal maps and add a degree of complexity by focusing on crossdomain statements about energy technologies, which allows us to establish a connection with the theoretical model of argument-communication theory and gain insight into societal issues.

2 Mapping Causal Statements on Energy Technologies from the Comment Section of *The Guardian*

In the contexts of computational social science and digital methods, the study of climate change opinion dynamics is an active area of research that covers many aspects of the (online) climate debate, including the position of various energy technologies as possible causes or solutions to the current climate crisis (see for instance Rogers and Marres, 2000), and sentiments in news coverage about nuclear power (Burscher et al., 2016). Likewise, recent modelling efforts in the field of opinion dynamics have focused on attitude formation on climate-related topics, such as smart mobility (Wolf et al., 2015) and energy sources and technologies (Banisch and Olbrich, 2021). Adding to this research, the present case study combines data mining, computational text analysis, and theoretical insights from conceptual models in order to map the overall position and mutual relationship between causal statements about energy technologies as figuring in The Guardian's online comment section on climate change articles.

The choice for studying the comment section of *The Guardian* is motivated by this news website's prominent position in the media landscape (Reid, 2018), as well as by its communicative setting, which is geared towards user engagement. Through this interaction with readers, the news platform embodies many of the recent shifts that characterize our present-day media ecology (Rusbridger, 2018, Chap. 11). As we are interested in general debate dynamics, and notably the argumentative statements that underpin these dynamics, it should be remarked at this point that we will focus primarily on comments, which we will aggregate for the entire debate sphere on articles about climate change. Comments will thus not be aggregated on the level of users or individual articles.

2.1 Collecting online opinion data from news websites

The communicative setting of *The Guardian* resembles that of typical online newspapers. Registered users can post comments on content that is open to commenting, and these comments are moderated according to community standards and participation guidelines (*The Guardian*, 2009, 2019c). What constitutes these commenters' belief systems might be inferred with varying degrees of precision. On the most general level, generic profiles of the 'progressive' readership of *The Guardian* (2019a) can be moderately informative. In order to truly assess opinion land-scapes, however, we need to turn to the actual comment texts and the statements these contain.

It should be noted that at the time of writing this article, readers' comments were not yet accessible through the API of The Guardian. For the scientific and educational purposes of this case study, we thus compiled our own dataset of climate-change-related news articles and news website comments using a dedicated web scraper. Concretely, articles from the 'climate change' (Guardian, 2019b) subsection of the news website of The Guardian dated from September 2008 up to and including April 2019 (N=13,064) were collected, along with up to 200 comments and associated metadata for articles where commenting was enabled at the time of scraping. The cut-off for the number of comments is a direct result of the technical restrictions of the data collection procedure. Since some articles received less than 200 comments over the period under discussion, this yielded a dataset of N = 985,823 comments (see Penelope Climate Data, 2020).¹

2.2 Constructing causal maps of the climate change debate by means of semantic frame extraction

Our web scraping efforts yielded a substantive collection of web texts from which argumentative statements can be extracted using text mining methods. We thereby focus specifically on expressions of causation, as it can be argued that expressions of causation (arguments of the form 'X causes Y') are central to the climate change debate. Clashes of opinions on climate change might for instance concern diverging assessments of whether global warming is man-made or not (for a sample of arguments in favor of or against anthropogenic global warming, see ProCon, 2019). Similarly, contradictory statements might be observed when it comes to the potential (beneficial or adversarial) ecological and environmental effects of energy sources and technologies, such as fossil fuels and renewables (Procon, 2019). Based on examples of this kind, it can be stated that argumentative expressions of causation are closely associated with opinions or beliefs, and that as such, these expressions can be considered a valuable indicator for the range and diversity of the opinions and beliefs that constitute the climate change debate.

Following Robert Axelrod's causal mapping method introduced earlier, expressions of causation can be represented visually as a 'causal map':

The concepts a person uses are represented as points, and the causal links between these concepts are represented as arrows between these points. This gives a pictorial representation of the causal assertions of a person as a graph of points and arrows. (Axelrod, 2016, p. 5).

In order to build up these graphs, these cause and effect statements are to be extracted from relevant sources by means of a series of heuristics and an encoding scheme (it should be noted that for this task Axelrod had human readers in mind). The assertion 'Our present topic is the militarism of Germany, which is maintaining a state of tension in the Baltic Area' might for instance be encoded as follows: 'the militarism of Germany' (cause concept), /+/(a positive relationship), 'maintaining a state of tension in the Baltic area' (effect concept) (Tucker Wrightson, 2016, pp. 296–97). Emphasizing the role of human interpretation, it is acknowledged that due to the complexities of the English language, no strict set of rules can capture the entire spectrum of causal assertions, but that such guidelines can nonetheless lead to outcomes that meet scientific standards in terms of validity and reliability (Tucker Wrightson, 2016, p. 332).

To facilitate the task of encoders, the causal mapping method has gone through various iterations since its original inception, all the while preserving its original premises. Recent software packages have for instance been devised to support the data encoding and drawing process (see for instance Laukkanen and Wang, 2015). Causal or cognitive mapping has thus become an established opinion and decision mining method within political science, business and management, and other domains. It has notably proven to be a valuable method for the study of recent societal and cultural conflicts. Homer-Dixon *et al.* (2014), for instance, rely on cognitive-affective maps created from survey data to analyze interpretations of the housing crisis in Germany, Israeli attitudes toward the Western Wall, and moderate *versus* sceptical positions on climate change. Similarly, Shaw *et al.* (2017) venture to answer the question of 'Why did Brexit happen?' by building causal maps of nine televised debates that were broadcast during the four weeks leading up to the Brexit referendum.

In order to apply any procedure for cognitive mapping to the domain of online argumentative communication, the method needs to be expanded from applications at the scale of human readers and relatively small corpora of archival documents and survey answers, to the realm of 'big' textual data and larger quantities of information. This necessarily involves a degree of automation, which we provide though semantic frame extraction.

2.2.1 Semantic frame extraction

Causal mapping is based on the extraction of so-called cause concepts, (causal) relations, and effect concepts from texts. The complexity of each of these concepts can range from the relatively simple (as illustrated by the easily-identifiable cause and effect relation in the example of 'German militarism' cited earlier), to more complex assertions such as 'The development of international cooperation in all fields across the ideological frontiers will gradually remove the hostility and fear that poison international relations', which contains two effect concepts (viz. 'the hostility that poisons international relations' and 'the fear that poisons international relations'). As such, this statement would have to be encoded as a double relationship (Tucker Wrightson, 2016, pp. 297–98).

The coding guidelines in Tucker Wrightson (2016) further reflect that extracting cause and effect concepts from texts is an operation that works on both the syntactic and semantic levels of assertions. This can be illustrated by means of the guidelines for analysing the aforementioned causal assertion on German militarism: 1. The first step is the realization of the relationship. Does a subject affect an object? 2. Having recognized that it does, the isolation of the cause and effect concepts is the second step. As the sentence structure indicates, 'the militarism of Germany' is the causal concept, because it is the initiator of the action, while the direct object clause, 'a state of tension in the Baltic area,' constitutes that which is somehow influenced, the effect concept (Tucker Wrightson, 2016, p. 296).

In the field of computational linguistics, this procedure for extracting information related to causal assertions from texts can be considered an instance of an operation called semantic frame extraction (for the concept of semantic frames, see Fillmore, 1982). A semantic frame captures a coherent part of the meaning of a sentence in a structured way. As documented in the FrameNet project (Baker *et al.*, 1998), the 'causation' frame is defined as follows:

A Cause causes an Effect. Alternatively, an Actor, a participant of a (implicit) Cause, may stand in for the Cause. The entity Affected by the Causation may stand in for the overall Effect situation or event (Framenet, 2001).

In a linguistic utterance such as a statement in a news website comment, the causation frame can be evoked by a series of lexical units, such as 'cause', 'bring on', etc. In the example 'If such a small earthquake causes problems, just imagine a big one!', the causation frame is triggered by the verb 'causes', which therefore is called the frame evoking element. The cause slot is filled by 'a small earthquake', the effect slot by 'problems' (Framenet, 2001).

In order to automatically mine cause and effect concepts from the corpus of comments on *The Guardian*, the present article uses the Penelope semantic frame extractor: a tool that exploits the fact that semantic frames can be expressed as form-meaning mappings called constructions (Penelope semantic frame extractor, 2020; for an in-depth discussion and evaluation of this tool see Beuls *et al.*, 2021). Notably, frames were extracted from *Guardian* comments by focusing on the following lexical units (verbs, prepositions, and conjunctions), listed in FrameNet as frame evoking elements of the causation frame: 'cause', 'due to', 'because', 'because of', 'give rise to', 'lead to', and 'result in'. The semantic frame extractor thus yields outputs such as the following:

'Has anyone totted up the extra pollution on London streets emanating from traffic jams [effect] caused by Extinction Rebellion [cause]?'

Excluding duplicates such as framesets in blockquotes, running the semantic frame extractor on the corpus of *Guardian* comments yields 123,957 complete cause-effect relations (framesets) on a range of topics.

2.2.2 Building up the causal map

A causal map diagram like the one proposed by Axelrod can automatically be generated from some of the most frequent causes and effect relations found in the comments. When we extract the 20 most frequent ngrams (bigrams) for the causes and effects in the retrieved framesets (after removing stopwords), 11 of the retrieved bigrams in the causes and effects overlap, which indicates that a certain effect of one phenomenon can in turn be mentioned as the cause of another. In Fig. 1, we present a causal map of a selection of the retrieved bigrams (viz. 'fossil fuel', 'global warming', 'sea ice', 'extreme weather', 'climate change', 'human activity', 'carbon emissions', 'CO2 emissions', 'atmospheric CO2', 'greenhouse gas', 'carbon dioxide', and 'greenhouse effects'), showing how these bigrams are interconnected based on the causal relations expressed in the comments. This graph was produced by looking up relations that had one of the



Fig 1. A causal map of a selection of the most frequent bigrams in cause frame elements and effect frame elements in statements mined from comments on climate articles on TheGuardian.com. This directed graph demonstrates how an effect (e.g. 'global warming' as an effect of 'fossil fuel') can in turn function as the cause of another phenomenon (e.g. 'extreme weather'). The thickness of the edges represents the number of occurrences of the expressed relation in the data. The causal map only shows relations with a minimum of 10 occurrences.

shared most frequent bigrams in both the cause frame element and effect frame element, and storing these relations as edges in a directed graph.

This automatically generated causal map offers a macroscopic view of the belief system expressed in the comments. Tracing some of the paths through this causal map for instance reveals how fossil fuels are causally related to carbon dioxide, which is in turn associated with the greenhouse effect. From there, we can further trace a path towards global warming, which is an effect of this carbon dioxide, and which in turn causes extreme weather. Furthermore, the causal map and the bigram analysis also reflect the prominent position of energy technologies and sources (e.g. 'fossil fuels' and 'nuclear power') in the climate change debate, which we will investigate further in the next section. As will follow, we can obtain detailed insights into the dynamics that govern The Guardian's opinion landscape concerning these energy technologies from establishing a dialogue between the extracted causal frames and the theoretical framework of argument communication theory.

2.3 Causal statements on energy technologies and argument communication theory

Opinions regarding different issues within a thematic complex such as energy production and climate change are usually not independent but correlated over a population. A person favouring windmills will for instance tend to think positively about solar power as well. On the other hand, persons in favour of coal are more likely to oppose renewables altogether. Argument communication theory seeks to explain such correlations on the basis of argumentative associations that cut across different topics. The identification of such cross-cutting arguments is therefore essential for an application to specific debates.

Argument communication models describe processes of collective deliberation as a repeated exchange of arguments. In the context of energy technologies, arguments may signal favor or disfavor regarding a specific technology, but arguments can also carry attitudinal information regarding more than one issue. Consider for instance, the following statement from one of the mined *Guardian* comments: Unfortunately efficiency comes from scale, and because the energy density of renewables is so much lower than for chemical or nuclear fuels [cause] you need a big area [effect].

For a human inspector it is quite obvious that this statement challenges renewables and supports 'chemical or nuclear fuels', as the big area required by renewables because of their lower energy density is presented as unfavourable. The role of such cross-cutting arguments in argument communication theory is illustrated in Fig. 2, which models the argumentative associations between renewables and nuclear power by a cognitive-evaluative structure of seven pro and con arguments. While six arguments target a single technology (three for each), C1, the statement corresponding with the example cited above, manifests a negative association between the two topics.

According to argument communication theory, the attitude on an issue is defined by the ratio of pro and con arguments someone holds. Therefore, someone who believes C1 will tend negatively to renewables and positively towards nuclear power, because the argument carries opposing evaluative contributions. In turn, a person with a positive attitude towards renewables will likely object to this statement or render it irrelevant while a person with a positive stance on 'chemical or nuclear fuels' will evaluate it in a more favourable way (Shamon et al., 2019). This cognitive mechanism is known as defensive processing (Wood et al., 1995) or attitude congruence bias (Taber et al., 2009) and has recently been integrated into argument communication theory (Banisch and Shamon, 2021). When attitudes are related by cross-cutting arguments as C1, certain combinations of arguments become incongruent and associated with higher cognitive dissonance (Festinger, 1957). For instance, a person who rejects C1 (top, right) will be biased to favour other con arguments on renewables over pro arguments, and vice versa for nuclear fuels, because these arguments are congruent with the attitude supported by C1.

At the collective level, even a small strive for cognitive coherence makes extreme opinions the most likely outcome (Banisch and Shamon, 2021). When individuals repeatedly exchange arguments, this induces a process of constraint satisfaction (Thagard and Verbeurgt, 1998) which converges to belief configurations that are more congruent than others because



coherent configurations



Fig. 2. Opinions on renewables and nuclear power modelled as a cognitive-evaluative map comprising seven pro and con arguments. The argument C1 contains a negative association between the two technologies, and therefore it increases the cognitive coherence of configurations that are strongly positive towards one and strongly negative towards the other technology.

incongruent configurations are slightly more difficult to maintain for each individual. In such a process, cross-cutting arguments may strongly constrain opinions in a population favouring specific configurations of beliefs at the extremes of the opinion scale. The most congruent configurations in our example, shown on the right-hand side of Fig. 2, are those in which the small negative correlation between opinions on renewables and nuclear power induced by C1 is amplified.

As follows from the above, the identification of cross-cutting arguments is particularly important to gain insight into the complex organization of opinions and beliefs on a thematic complex such as coal *versus* nuclear *versus* renewable energy sources. In the next section, we demonstrate that the semantic frame extractor can function as a tool for identifying real-world arguments that bridge argumentative domains related to different technologies for energy production.

2.3.1 *Frequency analysis of causal statements on energy technologies*

Statements on energy technologies can be identified among the extracted causal relations by filtering for a list of keywords derived from a recent overview of world electricity generation by source (IEA, 2017): 'coal', 'nuclear', 'biofuel', 'biomass', 'hydropower', 'hydroelectric', 'natural gas', 'solar energy', 'solar power', 'wind energy', 'wind power', and 'oil'. Out of the total number of causal relations that were extracted from the comment dataset, 3,097 thus contain references to one or more of the aforementioned energy technologies in the 'cause' frame element of the relation. Querying the causal part of the statements is motivated by argument communication theory, from which we expect that the issues in the theory (different technologies) are more likely to appear as a cause whose effects carry evaluations.

When we consider the distribution of statements by keyword or combination of keywords, it becomes apparent that the majority of retrieved statements matches for a single keyword (95.38%), and that statements containing two (4.59%) or three (0.03%) matches are less frequent. This suggests that overall, statements discussing one energy domain are more frequent than 'domain-crossing' statements that mention multiple energy technologies. Overall, the most frequent matches are the single terms 'coal' (37.16%), 'oil' (25.90%), and 'nuclear' (23.60%), suggesting that the domains of fossil fuels and nuclear energy take up more of the debate than renewables such as wind power (1.97%). The predominance of statements on oil, coal and nuclear energy is also reflected in the statements domain-crossing statements, where

matching on both 'oil' and 'coal' (39.8%), 'coal' and 'nuclear' (32.87%), and 'natural gas' and 'coal' (10.49%) are prominent in the distribution. Domain-crossing statements that combine terms associated with fossil fuels and renewables, such as 'coal' and 'wind power' (2.10%) are seemingly less frequent.

While a keyword-based approach thus points towards some initial trends in the causal statements on the energy debate, including the presence of domaincrossing arguments, cross-cutting arguments on this level (the 'microscopic level' of a single statement) seem rather sparse. Therefore, a more interesting and efficient method for retrieving domain-crossing instances that also offers a more 'mesoscopic' perspective on the debate, is to examine whether certain causes are associated with similar effects. This can be achieved by semantically clustering these effect statements.

2.3.2 Statement graphs and spatial graph embedding

In order to semantically cluster the extracted statements (strings), we make use of the 'statement graph generator' that is openly available in the Penelope ecosystem of tools (Penelope Network Tools, 2020). This instrument constructs semantic relatedness networks on the basis of shared lemmata (i.e. the dictionary form of words) between statements. The tool thereby first lemmatizes the words in the statements and filters them for nouns, verbs, and adjectives as main carriers of meaning. An edge between two statements is generated whenever the two statements share a lemma and the weight associated with this edge is the number of shared items. We acknowledge that more sophisticated measures of semantic relatedness exist (see Budanitsky and Hirst, 2006; Mehler et al., 2016), and that these might be worth closer examination for mapping short statements. However, the simple definition based on counting shared words has the advantage that derived relations are easy to interpret and understand. The networks of statements created in this way can be used to interactively explore the space of statements with network visualization tools such as Gephi or the Interactive Visualiser on the Penelope platform (Penelope network tools, 2020). For this article, Gephi has been used for it provides plenty of methods for spatial graph embedding (i.e.

force-based methods) as well as opportunities to interactively handle networks with annotations.²

Using the statement graph generator, we have clustered the above mentioned set of 3,097 causal statements about energy technologies based on the effect part of these statements. A force-based layout of this network in a 2D space enables us to view the principal patterns of semantic similarity between the effect portions, thus representing the main semantic clusters to which arguments in favour or against different technologies might refer. In Fig. 3, we present a selection of these effect clusters. The full graph containing all the statements on energy technologies is provided in the Appendix (Fig. A1). In this graph, the smaller nodes represent the clustered effect statements, coloured by MFW (most frequent word). A link is drawn from a statement in this semantic space to an energy technology when that technology is mentioned in the cause portion of that respective statement.

As statements thus cluster around semantic domains such as the technologies' effects on society ('people'), the economy ('price', 'cost'), and the environment ('damage', 'emission', 'pollution'), crosscutting argumentative relations become visible. It can for instance be seen that argumentative domains on nuclear energy and coal overlap to some extent, as both are concerned with the effects of these technologies on emissions. The domains on oil and coal and to some extent also nuclear and natural gas intersect in terms of their association with pollution. The domains on oil and coal, and some statements on biomass, solar, and nuclear share associations with cost and pricing. Finally, the argumentative domains of biomass, nuclear, oil, and coal share a focus on damage. Having identified these semantic domains, specific domains can be examined in more detail.

The graph in Fig. 4 offers a closer inspection of the domain of 'damage'. In the graph, all of the effect portions of the statements contain the word 'damage', and the cause parts of the statements contain one or more energy technologies. When the cause part of a statement contains two energy sources, its node is shared between the two energy domains.

The majority of statements in the graph evaluate the potential damaging effects of (variations of) a single energy technology. For oil, this for instance concerns the impact of oil companies, oil extraction, and oil spills (such as the 2010 Deepwater Horizon



Fig. 3. Visual rendition of a selection of the argumentative domains of causal statements that contain a reference to oil, biofuel, solar, biomass, nuclear, hydro, wind, coal, or natural gas in the cause portion of the statements, semantically clustered by effects. The nodes in the semantic clusters are coloured by the most frequent word (MFW). This graph was generated using the 'statement graph generator' that is available from the Penelope ecosystem of tools (Penelope Network Tools, 2020).

disaster). In terms of biomass, the extracted statement addresses the effects of its burning. For the nuclear domain, statements assess the impact of the use of (spent) nuclear fuel, nuclear weaponry, and nuclear processes. Finally, statements on coal concern the emissions of coal plants, coal mining, and coal-fired stations. The graph also reveals two arguments that are explicitly associated with two technologies. One of these addresses the damage ascribed to the combined effects of the oil and gas industries, another evaluates the effects of replacing polluting coal with cleaner gas.

Even though we limit our scope to the portion of causal statements related to energy technologies, two important observations can be made with regards to the analyses that were conducted in this section. For one thing, spatial renditions of causal statements based on a measure of semantic relatedness can provide some interesting insights into the dynamics, complexities, and diversity of statements that figure in online debate spheres on climate change. Spatial renditions for instance provide a mesoscopic view of the proximity and overlap between argumentative domains that might get lost in a regular causal map. For another, and from a methodological point of view, the real-world opinion data that was mined from online resources can inform the further construction of more abstract theoretical models for opinion dynamics, in this case argument communication theory. The sample of statements on energy technologies that we have examined for instance suggests that these



Fig. 4. Visual rendition of the argumentative domains of causal statements for which the effect portions contain the word 'damage', and the cause parts of the statements contain one or more energy technologies, viz. biomass, oil, nuclear, or coal.

statements are often confined to a single domain (a single energy source). However, we also found that similar arguments can be used to discuss these different domains. In the comments, fossil fuels and renewables are for instance pitted against each other in terms of cost. In addition, our analysis also allowed us to identify some domain-crossing statements that discussed the combined effects of energy sources, or that put two energy sources in opposition with each other (evaluating for instance the effect of replacing one technology with another). Such observations can inform further explorations of opinion dynamics in empirically-grounded scenarios.

3 Discussion and Conclusion

In this article, we have raised the question of how to explore the opinion landscapes expressed in large corpora of online data, such as the comment sections of news websites. We have argued that this challenge can be addressed through an alignment of computational methods for text analysis, and theoretical models for opinion dynamics, and we have stressed how the construction of tools and instruments can help bridge gaps between these fields. We have illustrated a possible configuration of different approaches by implementing a machine-guided pipeline for collecting, extracting, and modelling opinions related to energy technologies in the comment sections of The Guardian using modules from the Penelope ecosystem of tools for computational social science. In the remainder of this section, we highlight some of the main conceptual and methodological findings that result from this exploration, and discuss their implications for computational social science, and the use of digital methods within the SSH-disciplines in general. We start by addressing some limitations and avenues for future research in the domains of data, text mining tools, and models.

In terms of data and sources, a first series of remarks that can be made concerns the

generalizability of our proposed approach, that is, its suitability for the analysis of opinion landscapes beyond the rather specific case of energy technologies discussed in news website comments. To this end, it can be stated that while the causation frame has indeed proved to be a productive way of approaching this particular case, the causal frame extractor has also been integrated into prototypes for research instruments that have a wider focus, such as the Penelope climate change opinion observatory (Penelope climate change opinion observatory, 2020) and Penelope opinion facilitator (Penelope opinion facilitator, 2020; Willaert et al., 2021). These software projects address the climate change debate in more general terms, and include other sources, such as social media posts, parliamentary speeches and newspaper articles. To this it can be added that the extraction of causal frames can serve case studies well beyond the climate change debate, as exemplified by some of the causal mapping projects on Brexit and other societal issues cited in Section 2.2.

In the area of text mining methods, and semantic frame extraction in particular, it should be acknowledged that there are many ways in which causation can be conveyed in the English language, and not all of these are captured in our current implementation. Some of these expressions are explicit (e.g. 'Global warming causes rising sea levels'), whereas others (or, in fact, most) are implicit (e.g. 'Cutting CO2 emissions might reverse global warming'). Given the large diversity of possible expressions of causation in English, the computational method proposed in our showcase thus only captures a limited subset of the targeted opinion sphere, namely those expressions from a sample of maximum 200 comments per article that contain the frame evoking elements 'cause', 'due to', 'because', 'because of', 'give rise to', 'lead to' and 'result in'. Moreover, the method refrains from capturing certain qualitative aspects of the causal statements, such as the orientation of the relation (i.e. is the relation between cause and effect positive or negative?), and the strength of the relation (i.e. is there a strong causal association or not?). Advances in the area of semantic frame extraction, including the constructions of frame extractors that capture expressions of negation or modality, are thus a pathway for further research.

In this context, one of the main contributions of this article consisted in highlighting the role of arguments that cut across different domains and topics within the thematic complex related to energy production. Argument communication theory suggests that cross-cutting arguments constrain the opinion space by enforcing alignment of opinions regarding different technologies (Banisch and Olbrich, 2021). The outputs of text mining methods can be aligned with these argument-based models for collective deliberation because they provide insight into the types of arguments that are characteristic of certain debates. This is of help in the design of empirically informed models of opinion dynamics. While we have focused on the identification of arguments that address different topics, further advances in precision language processing are needed for the computational inference of the directionality (positive or negative) of arguments and hence for a more seamless integration of empirical data and models.

In addition to these observations, three more general concluding remarks can be made about the status of machine-supported methods in the social sciences and humanities. First, it can be stated that both 'distant' and 'close' readings and related methods from the humanities remain of the greatest value for dealing with the intricacies and subtleties of online information. As the importance of social media and digital media as sites of debate and opinion formation is becoming increasingly clear, (computational) linguistic analyses and lessons drawn from handling the complexities of various kinds of texts (e.g. political, literary or historical) prove increasingly relevant. In this sense, Axelrod's causal mapping method is but one of the many discourse-analytical frameworks awaiting to be implemented computationally. Furthermore, as the present article has illustrated for the case of argument communication theory, text mining methods can be used to empirically support theoretical frameworks for the analysis of opinions and beliefs. This linking of models to actual debates might in turn yield more precise and applicable theories.

Next, the technical implementation of text mining tools for online (social) media texts can be a catalyst for important conceptual innovations. Indeed, online communication and discourse pose numerous challenges to existing tools for the automated processing of language. Online discourse might for instance at once be very rich (e.g. large collections of news website comments), and very sparse (e.g. the scarce content of a single comment). Furthermore, online communication can be marked by completely innovative language (e.g. the use of neologisms), and an overall 'messiness' (e.g. in terms of the complexity of expressed frames, as well as more basic aspects of the written language, such as spelling and punctuation). The development of technologies that can deal with the diversity, innovative nature, and 'noisiness' of online discourse thus incites us to reconsider some commonly held assumptions about language and the processing of language. To this point, and as has been explored in this article, the conception and construction of experimental tools can be a productive practice for aligning the conceptual advances from different disciplines that are required for dealing with the intricacies of discourse on the web.

Finally, and crucially, the pursuit of technologically facilitated methods for opinion mining provides opportunities to re-assess and value assumptions about how to assign and convey meaning. Social media and discussion platforms impact how we communicate and express ourselves, but also how we synthesize information. A key question lifted by attempts to observe or render interpretable online communicative phenomena (such as opinion dynamics), is how to convey these observations. Interdisciplinary explorations across numerical, textual, and visual modes of expressions are required to achieve the most meaningful representations of online data.

Data and software information

The data and tools cited in this article are available from the Penelope ecosystem of tools and techniques for computational social science (see Penelope, 2020; Penelope Components, 2020).

Funding

This research project was funded by the European Union's Horizon 2020 research and innovation programme under Grant agreement No. 732942.

References

- Axelrod, R. (2016). The cognitive mapping approach to decision making. In Axelrod, R. (ed.), *Structure of Decision: The Cognitive Maps of Political Elites.* Princeton: Princeton University Press, pp. 3–17.
- Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). 'The Berkeley Framenet Project', in Proceedings of the 17th International Conference on Computational Linguistics. Association for Computational Linguistics, Vol. 1, Université de Montréal, Montreal, Quebec, Canada, pp. 86–90.
- Bakshi, R. K., Kaur, N., Kaur, R., and Kaur, G. (2016). Opinion mining and sentiment analysis. Proceedings of the 10th INDIACom; 3rd International Conference on Computing for Sustainable Global Development, INDIACom, New Delhi, pp. 452–455.
- Banisch, S. and Olbrich, E. (2019). Opinion polarization by learning from social feedback. *The Journal of Mathematical Sociology*, 43(2): 76–103.
- Banisch, S. and Olbrich, E. (2021). An argument communication model of polarization and ideological alignment. *Journal of Artificial Societies and Social Simulation*, **24**(1): 1.
- Banisch, S. and Shamon, H. (2021). Biased processing and opinion polarisation: experimental refinement of argument communication theory in the context of the energy debate. https://ssrn.com/abstract=3895117, http://dx.doi. org/10.2139/ssrn.3895117 (accessed 3 September 2021).
- Baumgartner, J., Zannettou, S., Keegan, B., Squire, M., and Blackburn, J. (2020). 'The Pushshift Reddit Dataset', in *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14, Atlanta, Georgia, USA, pp. 830–39.
- Beuls, K., Van Eecke, P., and Cangalovic, V. S. (2021). A computational construction grammar approach to semantic frame extraction. *Linguistics Vanguard*, 7(1), 20180015.
- **Borra, E. and Rieder, B.** (2014). Programmed method: developing a toolset for capturing and analyzing tweets. *Aslib Journal of Information Management*, **66**(3): 262–78.
- **Budanitsky, A. and Hirst, G.** (2006). Evaluating wordnet-based measures of lexical semantic relatedness. *Computational Linguistics*, **32**(1): 13–47.
- Burdick, A., Drucker, J., Lunenfeld, P., Presner, T., and Schnapp, J. (2012). *Digital Humanities*. Cambridge: MIT Press.
- Burscher, B., Vliegenthart, R., and Vreese, C. H. D. (2016). Frames beyond words: applying cluster and sentiment analysis to news coverage of the nuclear power issue. *Social Science Computer Review*, **34**(5): 530–45.

- **Castellano, C., Fortunato, S., and Loreto, V.** (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, **81**(2): 591.
- **Chavalarias, D. and Panahi, M.** (2018). Climate Tweetoscope. https://iscpif.fr/chavalarias/?portfolio=cli mate-tweetoscope (accessed 13 January 2021).
- **DMI Tools** (2020). DMI Tools. Digital Methods Initiative. https://wiki.digitalmethods.net/Dmi/ToolDatabase (accessed 13 January 2021).
- **Eve, M.** (2019). Close Reading with Computers: Textual Scholarship, Computational Formalism, and David Mitchell's 'Cloud Atlas'. Stanford: Stanford University Press.
- **Farzindar, A., Inkpen, D., and Hirst, G.** (2017). Natural Language Processing for Social Media: Second Edition. Synthesis Lectures on Human Language Technologies. San Rafael: Morgan & Claypool Publishers.
- **Festinger, L.** (1957). *A theory of cognitive dissonance.* Stanford: Stanford University Press.
- **Fillmore, C.** (1982). Frame semantics. In The Linguistic Society of Korea (ed.), *Linguistics in the Morning Calm*, pp. 111–37. Seoul, Korea: Hanshin Publishing Co.
- Framenet (2001). Causation. https://framenet2.icsi.berke ley.edu/fnReports/data/frameIndex.xml?frame=Causa tion (accessed 13 January 2021).
- Friedkin, N., Proskurnikov, A., Tempo, R., and Parsegov,
 S. (2016). Network science on belief system dynamics under logic constraints. *Science*, 354(6310): 321–26.
- Friedkin, N. and Johnsen, E. (2011). Structural analysis in the social sciences: Vol. 33. Social Influence Network Theory: A Sociological Examination of Small Group Dynamics. Cambridge: Cambridge University Press.
- *Guardian* (2009). Frequently asked questions about community on the Guardian website. https://www.theguar dian.com/community-faqs (accessed 20 December 2021).
- *Guardian* (2019a). The Guardian Advertising: Modern Advertising. https://advertising.theguardian.com/advertising (accessed 13 January 2021).
- *Guardian* (2019b). The Guardian Climate Change. https:// www.theguardian.com/environment/climate-change (accessed 13 January 2021).
- *Guardian*. (2019c). Sign in FAQs. https://www.theguar dian.com/help/identity-faq (accessed 20 December 2021).
- Hegselmann, R. and Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3). https://jasss.soc.surrey.ac.uk/5/3/2.html.

- Homer-Dixon, T., Milkoreit, M., Mock, S. J., Schröder, T., and Thagard, P. (2014). The conceptual structure of social disputes: cognitive-affective maps as a tool for conflict analysis and resolution. *SAGE Open*, 4(1). https: //doi.org/10.1177/2158244014526210.
- **ICIJ.** (2020). International Consortium of Investigative Journalists. Offshore Leaks Database. https://offshore leaks.icij.org/ (accessed 13 January 2021).
- IEA. (2017). World Electricity Generation By Source. International Energy Agency. Numbers retrieved via https://en.wikipedia.org/wiki/Renewable_energy#cite_ note-IEA-data-1 (accessed 13 January 2021).
- Laukkanen, M. and Wang, M. (2015). Comparative Causal Mapping: The CMAP3 Method. Farnham: Gower Publishing, Ltd.
- Maier, D., Waldherr, A., Miltner, P., Jähnichen, P., and Pfetsch, B. (2018). Exploring issues in a networked public sphere: combining hyperlink network analysis and topic modelling. *Social Science Computer Review*, 36(1): 3–20.
- Mani, I. (2013). Computational narratology. In Hühn, P. *et al.* (eds), *The Living Handbook of Narratology*. Hamburg: Hamburg University. http://www.lhn.uni-hamburg.de/article/computational-narratology (accessed 13 January 2021).
- Mäs, M. and Flache, A. (2013). Differentiation without distancing: explaining bi-polarization of opinions without negative influence. *PloS One*, 8(11). https: //doi.org/10.1371/journal.pone.0074516.
- Mäs, M., Flache, A., Takács, K., and Jehn, K. A. (2013). In the short term we divide, in the long term we unite: demographic crisscrossing and the effects of faultlines on subgroup polarization. *Organization Science*, **24**(3): 716–36.
- Mehler, A., Lücking, A., Banisch, S., Blanchard, P., and Job, B. (2016). Towards a theoretical framework for analyzing complex linguistic networks. Berlin: Springer.
- Pak, A. and Paroubek, P. (2010). Twitter as a corpus for sentiment analysis and opinion mining. *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valetta, Malta. pp. 1320–1326. http://lrecconf.org/proceedings/lrec2010/ pdf/385_Paper.pdf
- Peeters, S. and Hagen, S. (2018). 4CAT: Capture and Analysis Toolkit. Computer software. Vers. 1.0. https:// 4cat.oilab.nl/ (accessed 13 January 2021).
- **Penelope** (2020). The Penelope Platform: Tools and Techniques for Computational Social Science. https://penelope.vub.be/ (accessed 13 January 2021).

- Penelope Climate Change Opinion Observatory. (2020). The Penelope Climate Change Opinion Observatory. https://penelope.vub.be/observatories/climate-changeopinion-observatory/(accessed 13 January 2021).
- Penelope Climate Data. (2020). Guardian Climate Change Data API. https://app.swaggerhub.com/apis/EHAI/Guardi anClimateChangeData/1.0.0 (accessed 13 January 2021).
- Penelope Network Tools. (2020). Network Components API. https://app.swaggerhub.com/apis/pournaki/ PenelopeNetworkComponents/1.0.0 (accessed 13 January 2021).
- **Penelope Opinion Facilitator.** (2020). The Penelope Opinion Facilitator. https://penelope.vub.be/opinion-fa cilitator/(accessed 13 January 2021).
- Penelope Components. (2020). Components. https://penel ope.vub.be/components/(accessed 13 January 2021).
- **Penelope Semantic Frame Extractor**. (2020). Semantic Frame Extractor. https://app.swaggerhub.com/apis/ EHAI/Semantic-Frame-Extractor-API/1.0.0 (accessed 13 January 2021).
- **ProCon.** (2019). Is Human Activity Primarily Responsible for Global Climate Change? https://climatechange.pro con.org (accessed 13 January 2021).
- **Projet Politoscope.** (2017). Projet Politoscope. CNRS Institut des Systèmes Complexes Paris Ile-de-France (ISC-PIF). http://politoscope.org (accessed 13 January 2021).
- Reddit API. (2020). Reddit API Documentation. https:// www.reddit.com/dev/api/ (accessed 13 January 2021).
- Reid, A. (2018). Guardian.co.uk Most Read Newspaper Site in UK in March. https://www.journalism.co.uk/news/ nrs-guardian-co-uk-is-uk-s-top-monthly-news-site/s2/ a553108/ (accessed 13 January 2021).
- **Rogers, R. and Marres, N.** (2000). Landscaping climate change: a mapping technique for understanding science and technology debates on the world wide web. *Public Understanding of Science*, **9**(2): 141–63.
- Rogers, R. (2013). *Digital Methods*. Cambridge and London: MIT press.
- **Rogers, R.** (2019). *Doing Digital Methods*. Los Angeles: SAGE Publications.
- **Rusbridger, A.** (2018). Breaking News: The Remaking of Journalism and Why It Matters Now. Edinburgh: Canongate Books Limited.
- Sagi, E. and Dehghani, M. (2014). Measuring moral rhetoric in text. Social Science Computer Review, 32(2): 132–44.
- Schröder, T. and Wolf, I. (2017). Modeling multi-level mechanisms of environmental attitudes and nehaviours:

the example of carsharing in Berlin. *Journal of Environmental Psychology*, **52**: 136–48.

- Shamon, H., Schumann, D., Fischer, W., Vögele, S., Heinrichs, H. U., and Kuckshinrichs, W. (2019). Changing attitudes and conflicting arguments: reviewing stakeholder communication on electricity technologies in Germany. *Energy Research & Social Science*, 55: 106–21.
- Shaw, D., Smith, C. M., and Scully, J. (2017). Why did Brexit happen? Using causal mapping to analyse secondary, longitudinal data. *European Journal of Operational Research*, **263**(3): 1019–32.
- Singer, P. and Brooking, E. (2018). *LikeWar: The Weaponization of Social Media*. New York: Houghton Mifflin Harcourt Publishing Co.
- Stede, M., Schneider, J., and Hirst, G. (2018). Argumentation Mining: Synthesis Lectures on Human Language Technologies. San Rafael: Morgan & Claypool Publishers.
- Sunstein, C. (2018). #Republic: Divided Democracy in the Age of Social Media. Princeton: Princeton University Press.
- Taber, C. S., Cann, D., and Kucsova, S. (2009). The motivated processing of political arguments. *Political Behavior*, **31**(2):137–55.
- Thagard, P. and Verbeurgt, K. (1998). Coherence as constraint satisfaction. *Cognitive Science*, **22**(1):1–24.
- Tucker Wrightson, M. (2016). The documentary coding method. In Axelrod R. (ed.), *Structure of Decision: The Cognitive Maps of Political Elites.* Princeton: Princeton University Press, pp. 291–332.
- **Twitter API.** (2020). Programmatically Analyze, Learn from, and Engage with the Conversation on Twitter. https://developer.twitter.com/en/docs/twitter-api (accessed 13 January 2021).
- Underwood, T. (2019). Distant Horizons: Digital Evidence and Literary Change. Chicago: University of Chicago Press.
- Van Overwalle, F. and Heylighen, F. (2006). Talking nets: a multiagent connectionist approach to communication and trust between individuals. *Psychological Review*, 113 (3): 606.
- Wang, S. and Manning, C.D. (2012). Baselines and Bigrams: Simple, Good Sentiment and Topic Classification. *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics.* Jeju, Republic of Korea, pp. 90–94.
- Watts, D. J. (2013). Computational social science: exciting progress and future directions. *The Bridge on Frontiers of Engineering*, **43**(4): 5–10.
- Willaert, T., Van Eecke, P., Beuls, K., and Steels, L. (2020). Building social media observatories for monitoring

online opinion dynamics. *Social Media and Society*, April 2020. https://doi.org/10.1177/2056305119898778

- Willaert, T., Van Eecke, P., Van Soest, J., and Beuls, K. (2021). An opinion facilitator for online news media. *Frontiers in Big Data* **4**, July 2021. https://doi.org/10. 3389/fdata.2021.695667
- Wolf, I., Schröder, T., Neumann, J., and de Haan, G. (2015). Changing minds about electric cars: an empirically grounded agent-based modelling approach. *Technology Forecasting and Social Change*, **94**: 269–85.
- Wood, W., Rhodes, N., and Biek, M. (1995). Working knowledge and attitude strength: an information-processing analysis. *Attitude strength: Antecedents and Consequences*, **4**: 189–202.

Notes

1 The articles' text and metadata were scraped from *The Guardian*'s website using the open source Scrapy

framework. Causal frames were then extracted from each article's text and comments, and added as an additional field to the article's database entry. The database (MongoDB) allows for easy querying based on keyword(s) or timestamps. This database is available via a web service through requests.

2 For the exploratory analysis in Gephi that were also used in this article, the Statement Graph Generator generates two node labels. The first one is the full statement and the second one corresponds to the word (noun, verb, or adjective) in the statement that is most frequent in the entire corpus (MFW, for most frequent word). While showing the full statements is informative for small graphs (<100 statements), it is not appropriate if graphs contain a larger number of statements. For a large set of statements, the label based on overall word frequency (MFIC) is a more appropriate indication of the topic addressed by a statement.

Appendix



Fig. A1. Visual rendition of all argumentative domains of causal statements that contain a reference to oil, biofuel, solar, biomass, nuclear, hydro, wind, coal, or natural gas in the cause portion of the statements, semantically clustered by effects. The nodes in the semantic clusters are coloured by the most frequent word (MFW). This graph was generated using the 'statement graph generator' that is available from the Penelope ecosystem of tools (Penelope Network Tools, 2020).