Media analysis nuclear energy

Jakob Rauch
Vrije Universiteit Amsterdam
Department of Spatial Economics

October 13, 2021

Abstract

We analyze 3,005 unique articles about nuclear energy from Dutch newspapers, published between 2011 and 2021. The goal is to get an understanding of the key developments in the public discourse on nuclear energy; to find out which topics, persons, and organizations are important in Dutch news media, and how this changes over time. We use anchored topic modeling (Gallagher et al., 2017) to separate nine clearly defined topics: Climate impact, waste and storage, geopolitics, safety, cost, technology, ethics, politics, and choice of site. We show that the safety topic has become much less prominent since Fukushima, while climate impact is discussed considerably more. We also see that the discussion about expanding nuclear energy in the Netherlands has picked up since last year, with articles mentioning Dutch politics and possible site choices becoming more frequent. We apply sentiment analysis (Hutto and Gilbert, 2014) and show that the safety and waste disposal topics are associated with the most negative terms, while climate impact is associated with more positive terms. Generally, nuclear energy is discussed in much more positive terms today than it was after Fukushima and the average sentiment surrounding the topic has been in a steady upward trend over the last years. We finally use named-entity recognition based on Schweter and Akbik (2020) and Conneau et al. (2019) to identify the key persons and organizations in the Dutch discourse. The resulting network can be visualized and gives insight into how the national and international stakeholders of the Dutch discourse are connected.

.

The code for this project can be found at https://github.com/jakob-ra/rli-nuclear-energy

Contents

1 Introduction

This project was comissioned by the Dutch Council for the Environment and Infrastructure (Raad voor de leefomgeving en infrastructuur, henceforth Rli). The aim is to get an understanding of the key developments in the public discourse on nuclear energy; to find out which topics, persons, and organizations are important in Dutch news media, and how this changes over time. To this end, we analyze 3,005 unique news articles in Dutch language. The dataset and pre-processing steps are described in section 2. We employ various natural language processing methods, such as anchored topic modeling, sentiment analysis, and named entity recognition. These methods are described in section 3. The results are presented in section 4.

2 Data

The following section details how the extract of news articles was obtained from LexisNexis, how duplicate articles were dealt with, and what the resulting dataset looks like in terms of descriptive statistics.

2.1 LexisNexis extract

The data consists of 3,578 articles from Dutch news media. The articles are in Dutch language and were downloaded through the searchable news database of LexisNexis using the following search string:

> (kernenergie! OR "nucleaire energie" OR "nucleaire stroom" OR "nucleaire elektriciteit" OR atoomenergie OR atoomstroom OR nucleair) OR (kerncentrale OR atoomcentrale OR kernreactor OR "nucleaire centrale" OR atoomreactor) AND (klimaat! OR energietransitie OR energiesysteem) AND NOT ("nucleair! verdrag" OR "nucleair! akkoord" OR "nucleair! programma" OR "nucleair! wapen!" OR "nucleair! deal" OR "nucleair! activiteit!").

We limit the analysis to articles published between 11$^{\text{th}}$ of March 2011 (the day of the Fukushima disaster) and 2$^{\text{nd}}$ of June 2021. Only articles from the following sources were considered:

- National: NRC Handelsblad, de Volkskrant, Het Financieele Dagblad, De Telegraaf, Trouw, Algemeen Dagblad, Reformatorisch Dagblad, het Parool.
- Regional: Dagblad van het Noorden, Provinciale Zeeuwse Courant, Dagblad De Limburger, Noordhollands Dagblad, Leidsch dagblad, Brabants dagblad, de Gelderlander, De Stentor, Tubantia, Eindhovens Dagblad.
- Sector: Energeia.
- Weekly magazines: Elsevier Weekblad, De Groene Amsterdammer

2.2 Manual article selection

For the remanining articles, Rli confirmed manually that they are relevant to the analysis. Exclusion criteria included the following:

- Articles that do not have nuclear energy or nuclear power plants as a (sub)topic. For example articles that mention the terms in passing without elaboration.
- Articles about nuclear energy or nuclear power plants outside Europe, unless a connection is made with nuclear energy or nuclear power plants in The Netherlands. Exceptions are the nuclear accident in Fukushima, considering its effect on public opinion.
- Articles on nuclear accords and non-proliferation, unless a connection is made with nuclear energy or nuclear power plants in The Netherlands.
- Relevant identical articles in different regional papers are all included to enable regional analysis.
- Errors, being not a news article or a letter e.g. 'agenda's'

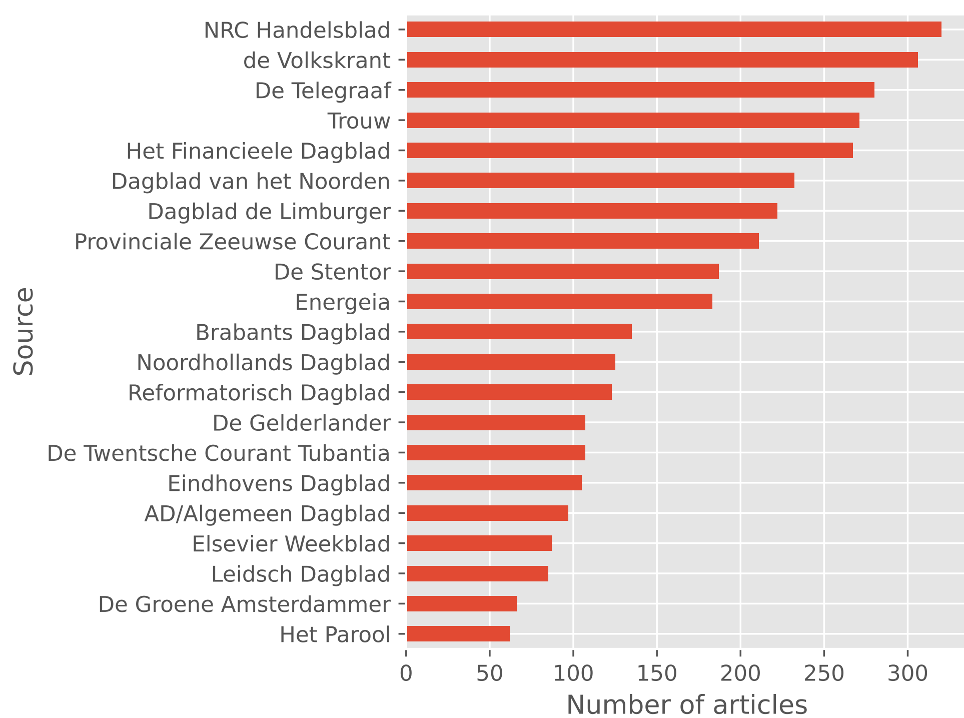

Figure 1. Number of articles by source.

2.3 Article deduplication

In the LexisNexis extract, we encountered duplicate sources, such as 'Dagblad De Limburger' and 'Dagblad de Limburger', which were merged. We also encountered nine different local version of 'De Stentor', such as 'De Stentor / Apeldoornse Courant', 'De Stentor / Dagblad Flevoland', and 'De Stentor / Gelders Dagblad'. These were also merged in order to keep the number of sources manageable.

In the next step, we dropped articles with exactly the same text and source as another article, since these articles do not provide us with additional information (162 articles). There are also 411 articles with the exact same text, which appear in multiple sources. Since we are interested in the differences between sources (e.g. between different regions), we simply group these articles and keep a list of the sourcs they appeared in. This leaves us with 3,005 unique articles.

2.4 Descriptive statistics

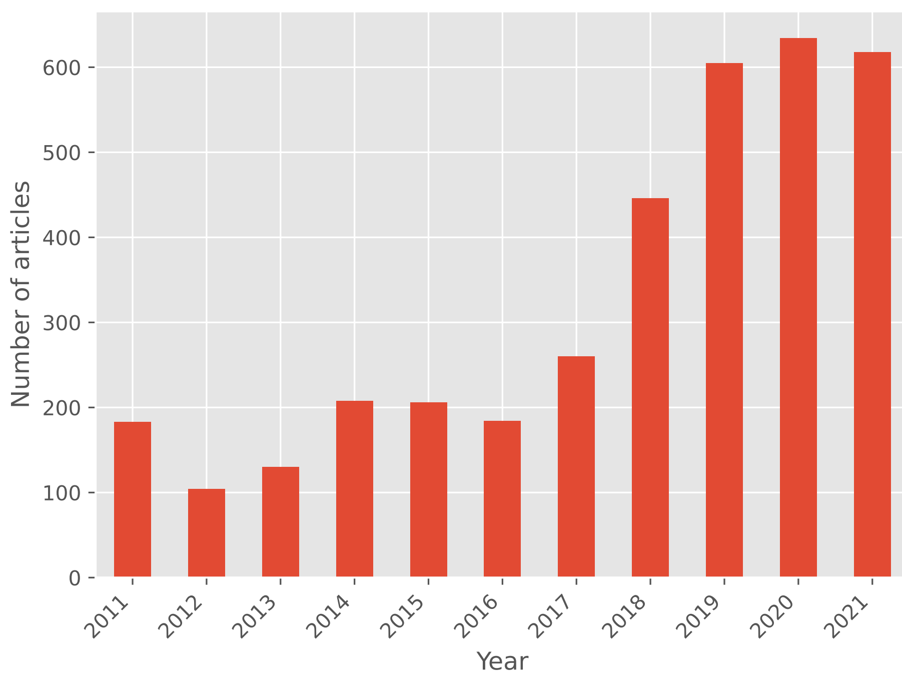

Figure 2. Number of articles over time.

Figure 1 gives an overview of the number of articles by source. It shows that the big national newspapers publish more articles on the topic of nuclear energy than smaller regional newspapers and weekly magazines.

Figure 2 shows the number of articles published each year. We see that 2011 has a spike in articles compared to the following years, most likely related to the Fukushima disaster. We also see that the bulk of articles is published in the last four years, indicating significantly increased interest in the topic in recent years.

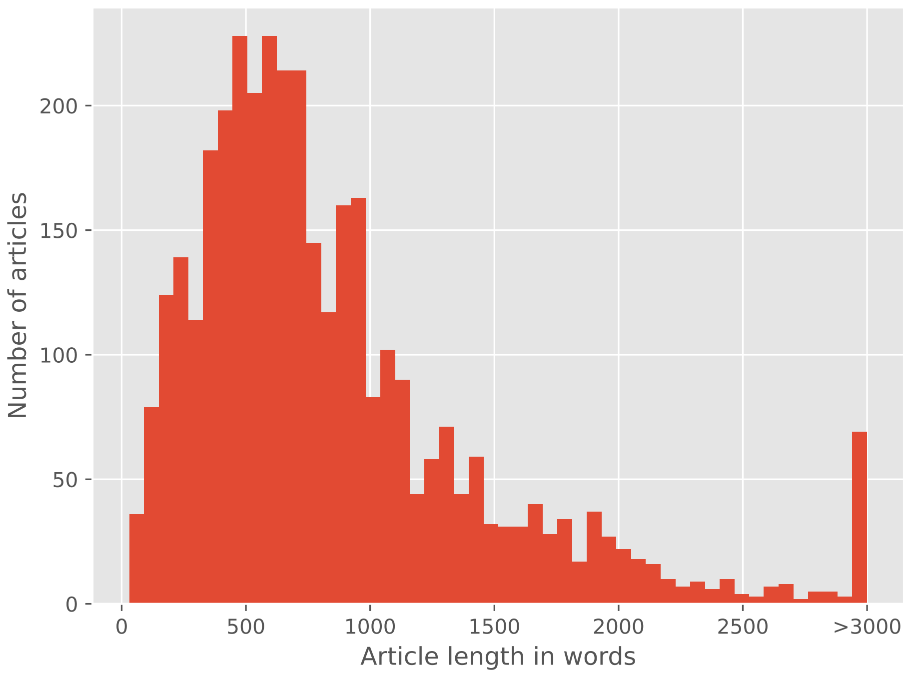

Figure 3. Histogram of article lengths.

Figure 3 shows the distribution of article lengths. We see that the bulk of articles is made up of articles up to 1000 words, while there are only few long-form articles. Figure 4 further shows that these long-form articles are concentrated in the two weekly magazines, and that the national newspapers tend to publish longer articles than the regional ones.

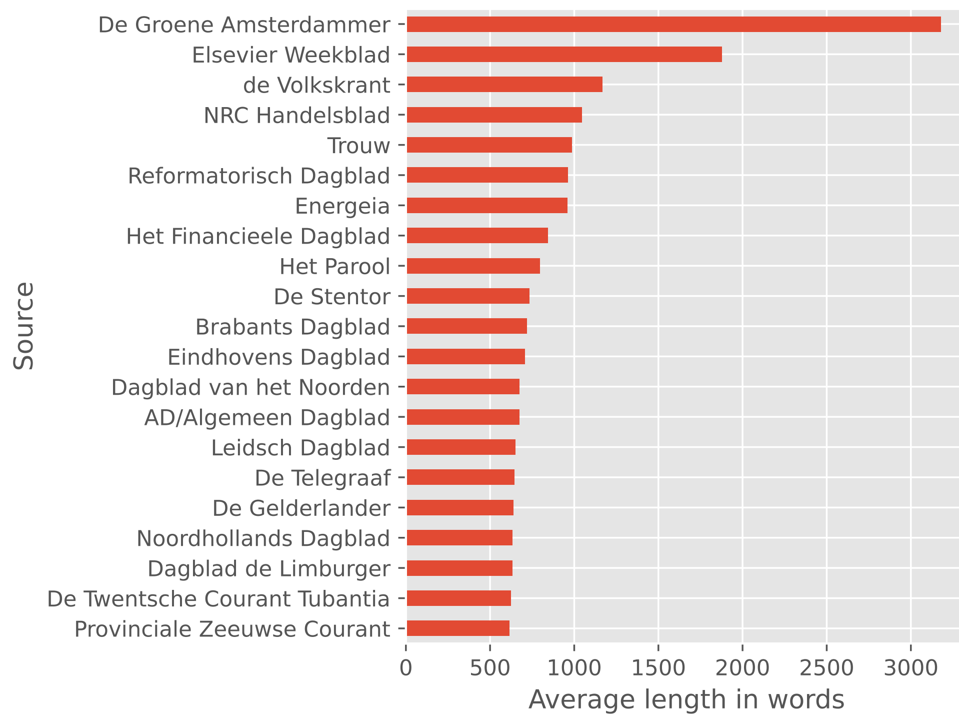

Figure 4. Average article length by source.

3 Methods

We use three methods to summarize the content of the articles: topic modeling, named-entity recognition, and sentiment analysis. The following section will describe each of them in detail.

3.1 Topic modeling

Topic models are a class of statistical models used to discover the abstract topics contained in a collection of documents. Words that tend to appear together within documents are clustered together. We thus end up with a number of such word clusters; looking at the words in the cluster we can determine an appropriate topic name.

Because topic models operate on the list of unique words found across all documents, it is advantageous to group different forms of the same word together (such as 'broeikasgassen' and 'broeikasgas'). This was done by lemmatizing all words in the documents (i.e. reducing back them back to their root form) using the natural language processing library spaCy (Honnibal et al., 2020).

To further improve the specificity of the topic model, we used a stop word list to remove the 2,000 most common words in the Dutch language (using the list of Speer et al., 2018). These stop words are usually too common to be informative about a narrowly defined topic. To make sure that we do not remove words that happen to be informative in our context (e.g. 'veiligheid', 'leeftijd'), the stop word list was manually reviewed. Finally, we also remove words that are uninformative because they were part of our initial search query, such as 'nucleaire energie', 'kernenergie', and 'atoomcentrale'; the idea being that we are looking for topics *within* the topic of nuclear energy.

The word list used as an input list for our topic models is then constructed as follows: we take the most 50,000 most common unigrams (single words) and bigrams (combinations of two words), with the restriction that they have to appear in at least five different articles.

We first take an entirely unsupervised approach and apply Latent Dirichlet Allocation (LDA, see Pritchard et al., 2000; Blei et al., 2003). This approach has the advantage of letting the data speak and potentially discovering previously unknown topic clusters. Based on prior research of the Rli we were in the fortunate position to have a lot of domain knowledge and a very good idea of the topics surrounding nuclear energy. However, the results from LDA led to the discovery of another topic present in the data: that of ethics, based around words such as 'mensheid', 'planeet', wetenschap', 'filosoof', 'ethiek', and 'moreel', as shown in figure 5.

Although useful for topic discovery, the unsupervised LDA approach lacks in terms of separating topics that we know to be distinct and might not detect underrepresented topic clusters that are of interest. In the next step, a set of topics for further analysis was decided upon in consultation with Rli, based on prior research and the previous results. The topic list features the nine most prominent topics that can be clearly separated. There were several potential other topics that were identified by Rli in previous research. However, these topics turned out to be mentioned too little (for instance, electric grid stability), or were not clearly separable from existing topics and therefore absorbed by superordinate topics (e.g. intergenerational fairness $\rightarrow$ ethics).

The final topic list is as follows: Climate impact, waste and storage, geopolitics, safety, cost, technology, ethics, politics, and choice of site. To ensure clear separation of these topics, we apply an anchored topic model (Gallagher et al., 2017). For each of the topics, we define a list of anchor words. These anchor words help guide the topic model into the direction of these words, while being flexible enough to integrate similar words into the same topic cluster.

The initial anchor words are chosen using both prior research and the results from unsupervised LDA. Based on the initial list, multiple methods are used to further extend the list of anchor words. First, we identify similar words by looking at the words most frequently

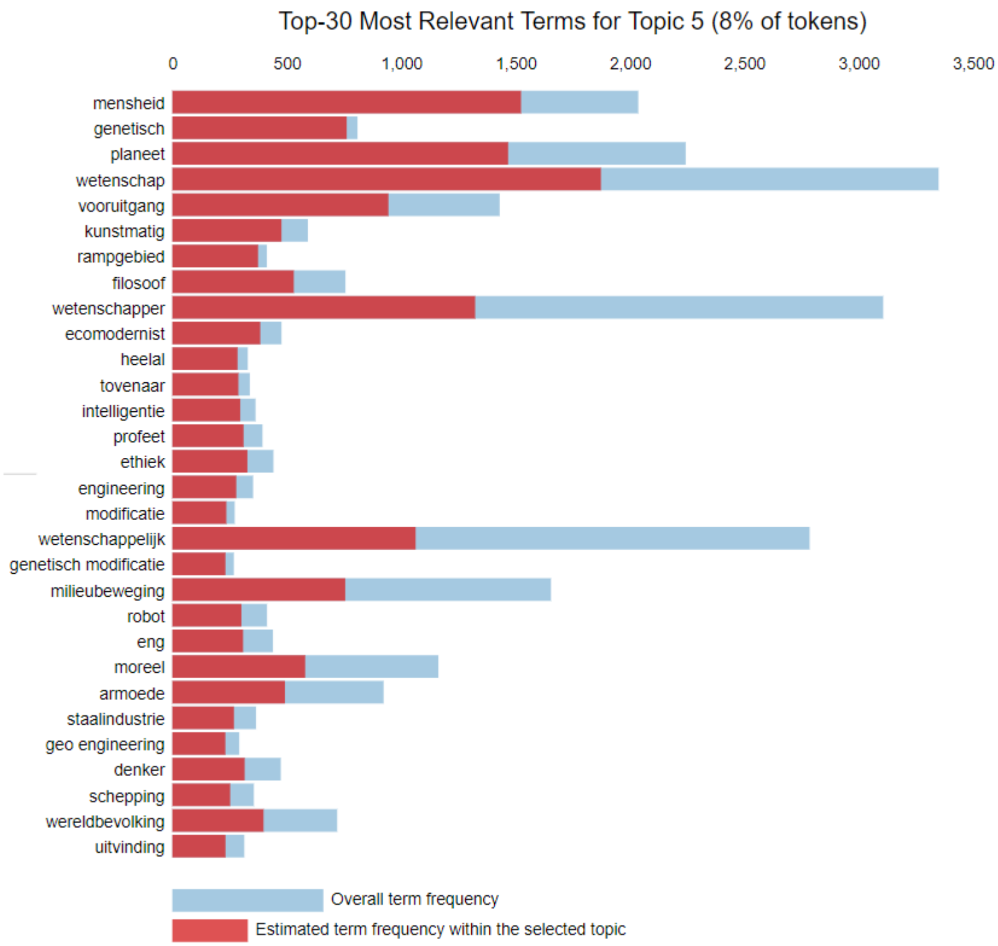

Figure 5. 'Ethics' topic cluster discovered using (unsupervised) LDA.

occurring together with the initial anchor word. Next, we find similar words via pre-trained word vectors, trained on Dutch wikipedia (Bojanowski et al., 2017). Finally, we find the most frequent phrases containing a sub-phrase that we already know to be part of a given topic (e.g. 'reactor' reveals 'zoutreactor' and 'reactorvat' as other, more specific topic keywords for the technology topic).

The anchored topic model is fit to the data using 15 topics. Besides the 9 anchored topics, 6 topics are used to 'soak up' remaining topics that we are less interested in. The anchor strength parameter, which controls the weight given to an anchor word relative to other words, is set to 10. This makes sure that we get output topics corresponding to our anchor words and that all anchor words for a predefined topic end up together in the output topics.

3.2 Sentiment analysis

To find out whether nuclear energy is reported about in a negative or positive context, we employ sentiment analysis. There is at the time of this writing no model for sentiment analysis in the Dutch language with satisfactory performance. Therefore, for the purpose of sentiment analysis, all articles were translated into English using Google Cloud Translate.

The translated articles were then used as input to the VADER algorithm (Hutto and Gilbert, 2014). The algorithm assigns a valence score between -1 (very negative) to 1 (very positive) to each word, based on a lexicon of sentiment-indicating words. The algorithm also uses a pattern-based rule set to detect, for instance, negation. As recommended by the authors, we use the compound score, a weighted composite score based on the valence of each input word that is also normalized to lie between -1 and 1.

It should be mentioned that the absolute levels of sentiment scores may not be very informative. Considering that words with negative sentiment are just generally less frequent, average scores tend to have a positive bias, and an average score above 0 does not necessarily mean that, e.g. a topic is reported about more positively than negatively. The focus in interpreting sentiment scores should instead be on relative comparisons (for instance, between topics or over time).

3.3 Named-entity Recognition

We use named-entity recognition to extract the people and organizations mentioned in the news articles. The goal is to identify the most important stakeholders in the debate about nuclear energy and how they are connected.

To this end, we employ an out-of-the-box solution provided by Schweter and Akbik (2020) based on the pre-trained XLM-RoBERTa transformer model introduced by Conneau et al. (2019). The model is fine-tuned on the Dutch CoNLL-03 named-entity recognition task, where it achieves state of the art performance.

4 Results

The following section lays out the results regarding topics, sentiment, and entities surrounding nuclear energy in the Dutch news media.

4.1 Topic prominence

The anchored topic model achieves an excellent coherence score for the anchored topics after being fit to the data. The most specific keywords associated with each topic are given in the

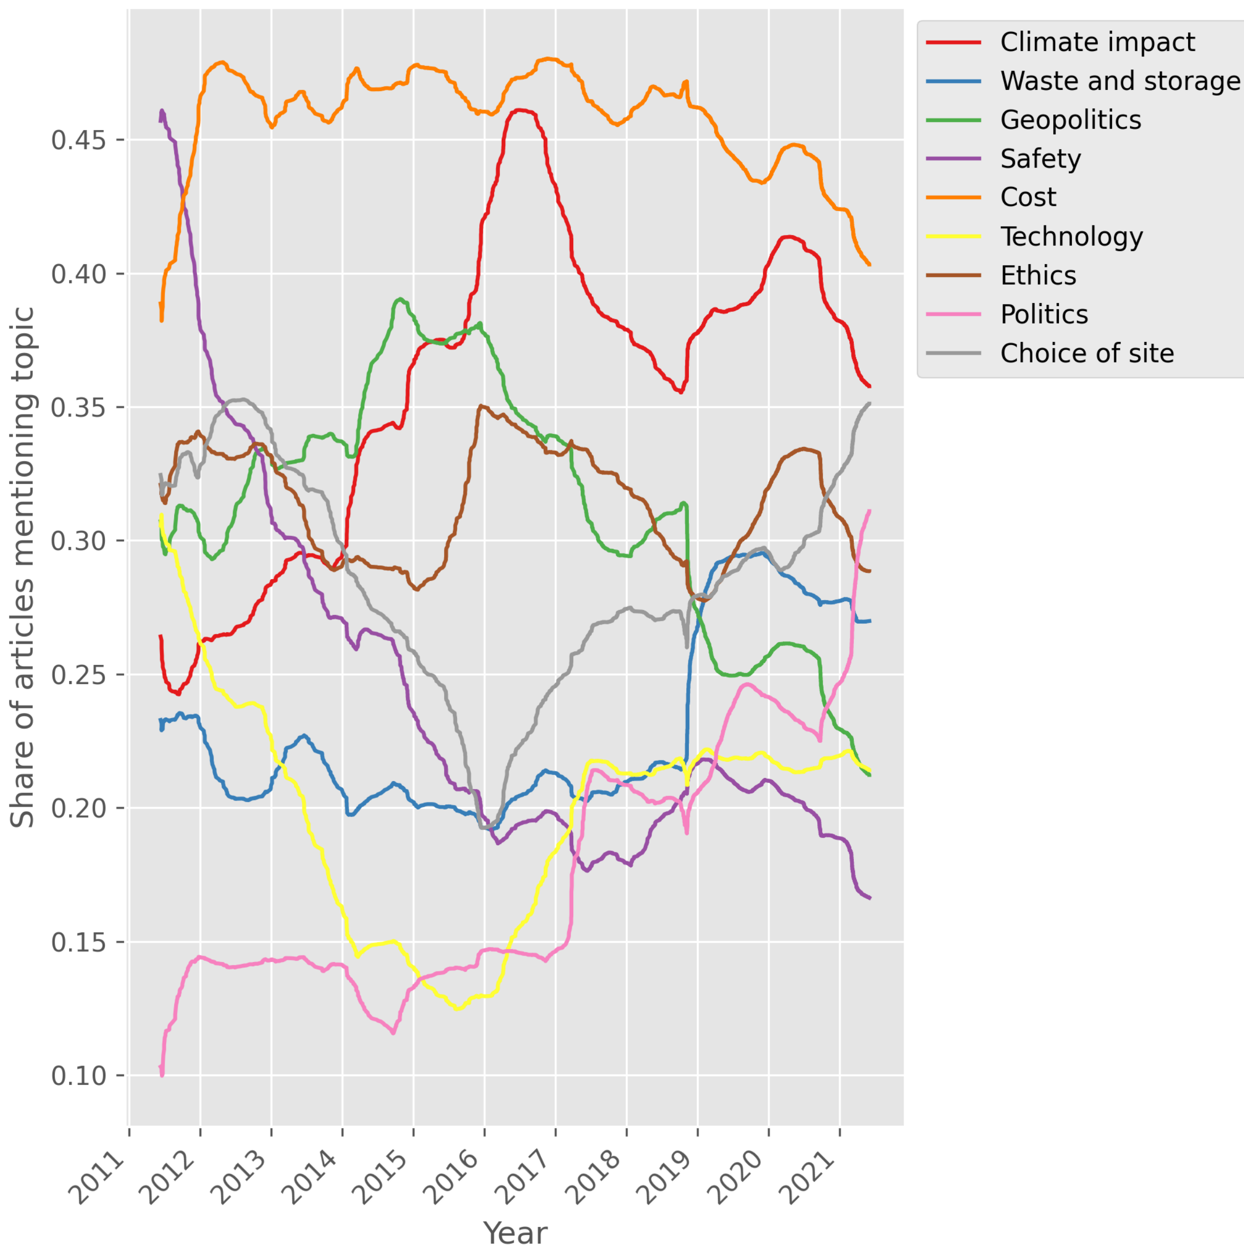

Figure 6. Topic prominence over time.

appendix. Using the fitted topic model, we now make predictions for whether an article mentions a given topic. The prediction is based on the words in the article and their specificity score for the topic.

Figure 6 shows the prominence of topics over time. Most strikingly, safety goes from being the most prominent topic to being the least prominent topic, while climate impact is increasingly prominent. Cost is consistently a very prominent topic. Politics and choice of site become increasingly prominent in the last few years, indicating a new spark in the debate of whether nuclear energy in the Netherlands should be expanded. The waste and storage topic also becomes much more prominent in 2019, connected to a series of proposals for disposal sites of nuclear waste.

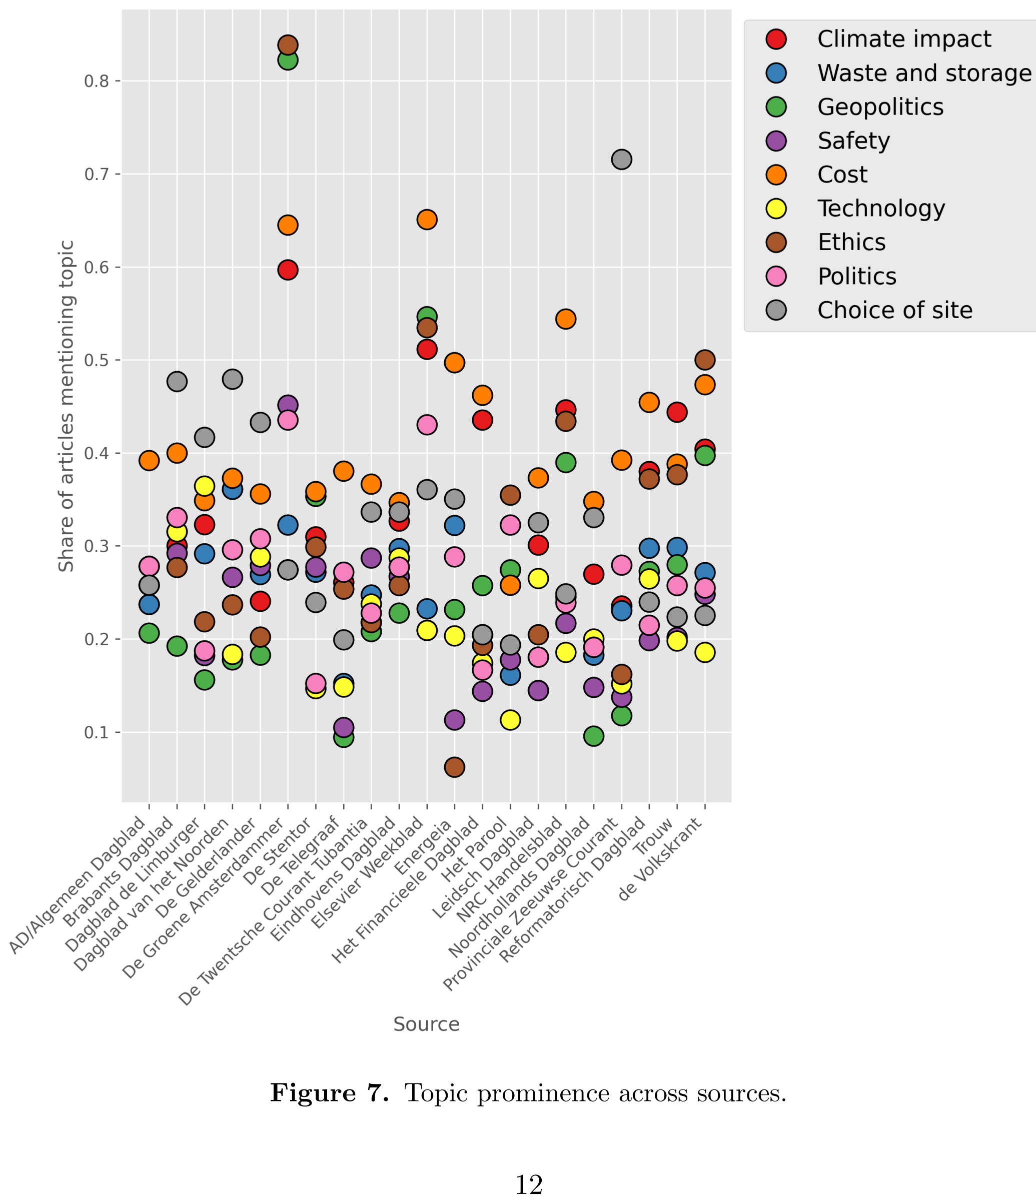

Figure 7. Topic prominence across sources.

Figure 7 also shows topic prominence differentiated by news outlet. We can see that the regional newspapers report more on the choice of site topic, namely those from the North, East, and South of the country (Brabants Dagblad, Dagblad de Limburger, Dagblad van het Noorden, De Gelderlander). Some of the corresponding regions contain designated sites for new nuclear power plants or have raised potential siting of their own accord at the provincial level. Provinciale Zeeuwse Courant, a newspaper from Zeeland, the only region that currently has a nuclear plant, writes twice as much about this topic than average. The cost topic tends to be discussed more in the long-form articles of De Groene Amsterdammer and Elsevier Weekblad, as well as the NRC Handelsblad.

4.2 Sentiment

Figure 8 shows that the sentiment with which nuclear energy is reported about has recovered from its all time low after Fukushima and has been on a relatively steady upward trend since then.

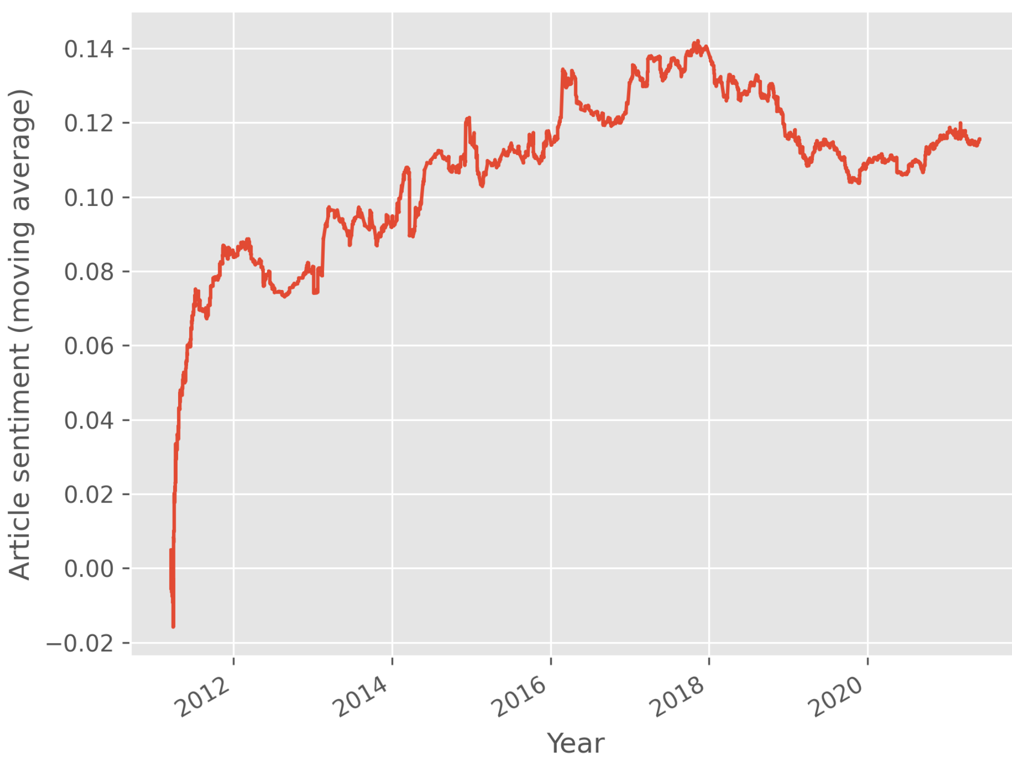

Figure 8. Overall sentiment over time.

Figure 9 shows the average sentiment with which each topic is discussed. Unsurprisingly, the two big disadvantages of nuclear energy, safety and waste disposal are discussed most negatively, while its potential to reduce emissions is discussed positively. The fact that the politics topic has the most positive sentiment here can be primarily interpreted as politics being discussed positively in the context of nuclear energy and probably to a lesser extent nuclear energy being discussed positively by political entities. Note that the sentiment scores of topics are purely relative measures. They only indicate whether the wording used around a certain topic is more or less negative/positive than the words used around another topic.

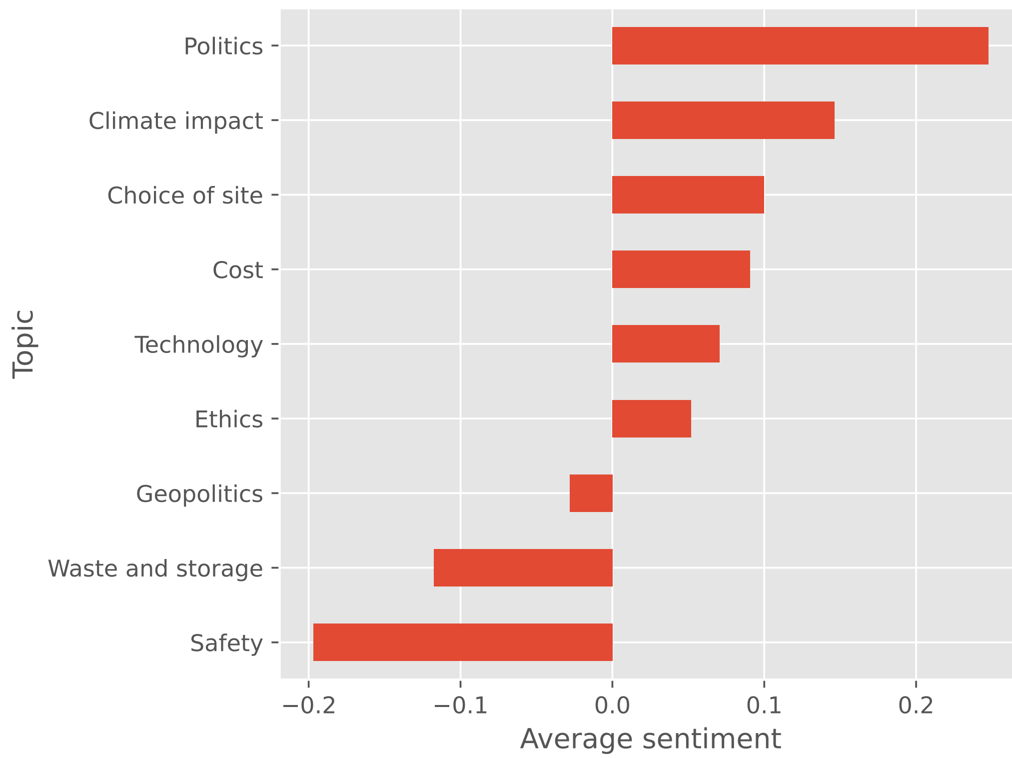

Figure 9. Topic sentiment.

4.3 People and organizations

Figure 10 shows the 50 most frequently mentioned organizations. The list includes Dutch (and some German) political parties, EU institutions, energy companies, international organizations, and NGOs.

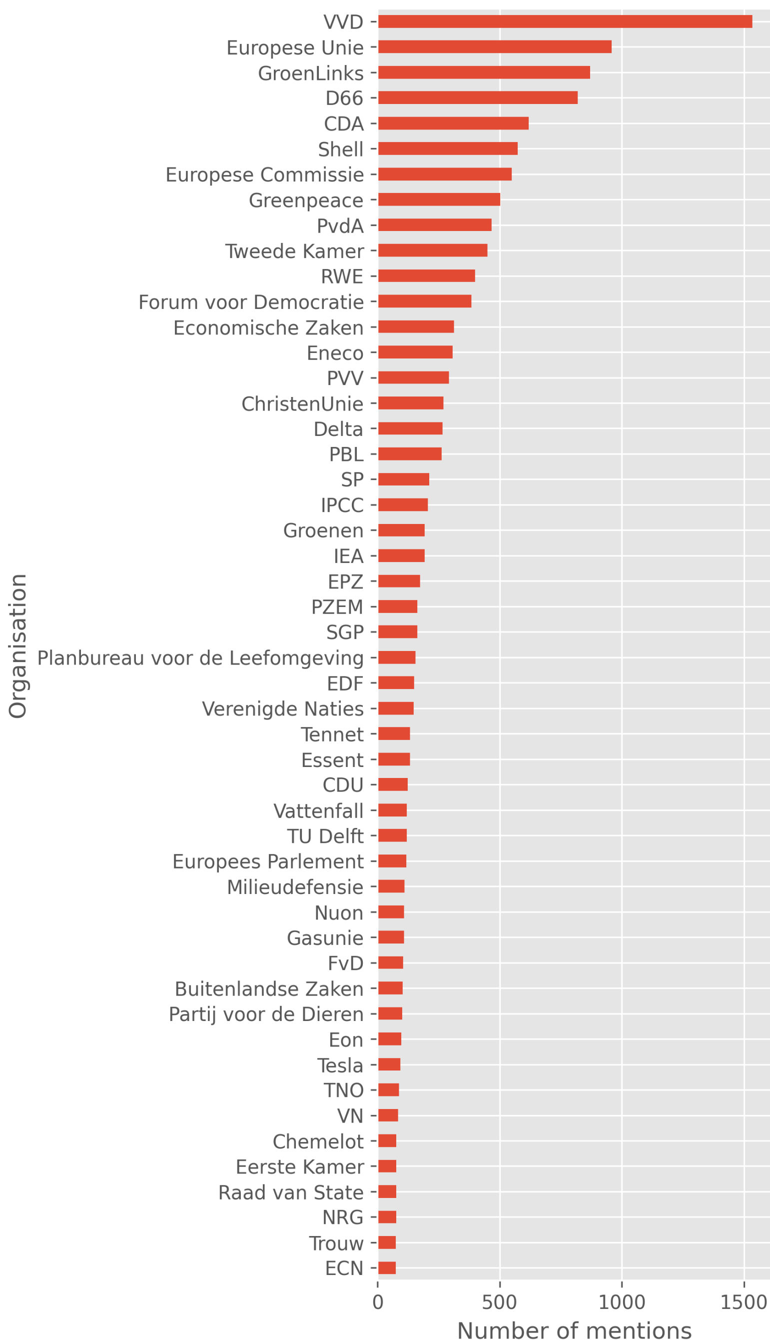

Figure 10. Top 50 mentioned organizations.

Figure 11 shows the 50 most frequently mentioned people. The list is heavily dominated by Dutch and international politicians. There are also researchers like Jan Leen Kloosterman and activists, such as Marjan Minnesma and Johan Vollenbroek. The prominence of some persons are quite informative, such as that of Loek Radix, executive director of Chemelot (the company's industrial park in Limburg is being discussed as a potential site for a nuclear power plant) or that of Andre Versteegh, nuclear engineer and chairman of Stichting KernVisie, a foundation committed to the advancement of nuclear energy.

Overall, a striking result is that Dutch politicians and political organizations make up the vast majority of recognized entities. To some extent, this is not surprising, as the largest section of most newspapers is dedicated to reporting on day-to-day politics. It is interesting, however, that the topic of nuclear energy is mentioned so much in this context of Dutch politics (rather than other contexts such as research or business news). As Figure 6 shows, the discussion of nuclear energy in Dutch politics has also intensified over time.

Going beyond mere mentions of persons and organizations, we can also show how they are inter-connected. For this, we look at the co-occurrence of entities within a sentence. A visualization of the resulting network can be found in figure 12. The position and color of the nodes in the network is determined by a clustering algorithm. We see that some of the clusters are immediately recognizable: Dutch politics, international politics, European politics, NGOs, international organizations, research, and energy companies.

Note that in this visualization only the larger nodes are shown. A full, interactive visualization can be found here. In the interactive visualization, we can find information on the topic loadings of each entity (i.e. how often the entity co-occurs with the mention of certain topics within a sentence), as well as the average sentiment associated with the entity (also measured on the sentence level).

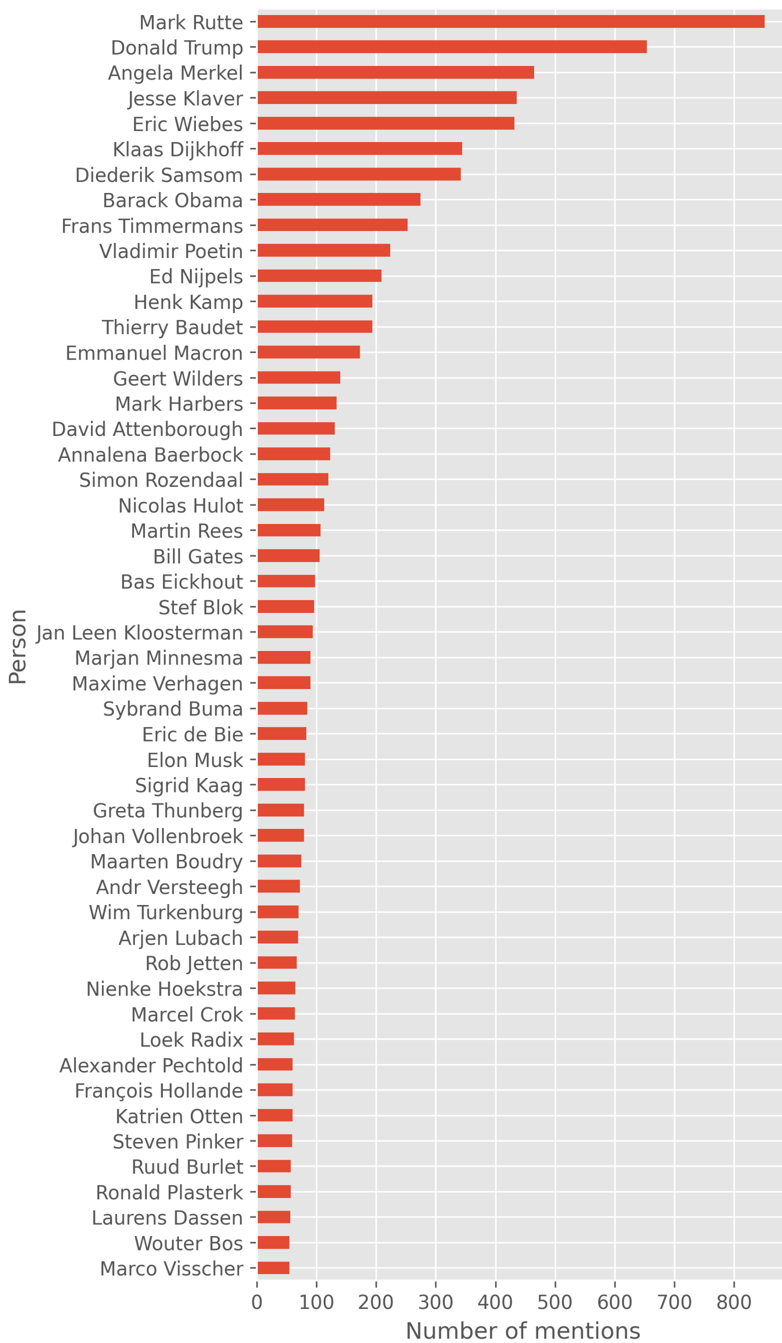

Figure 11. Top 50 mentioned people.

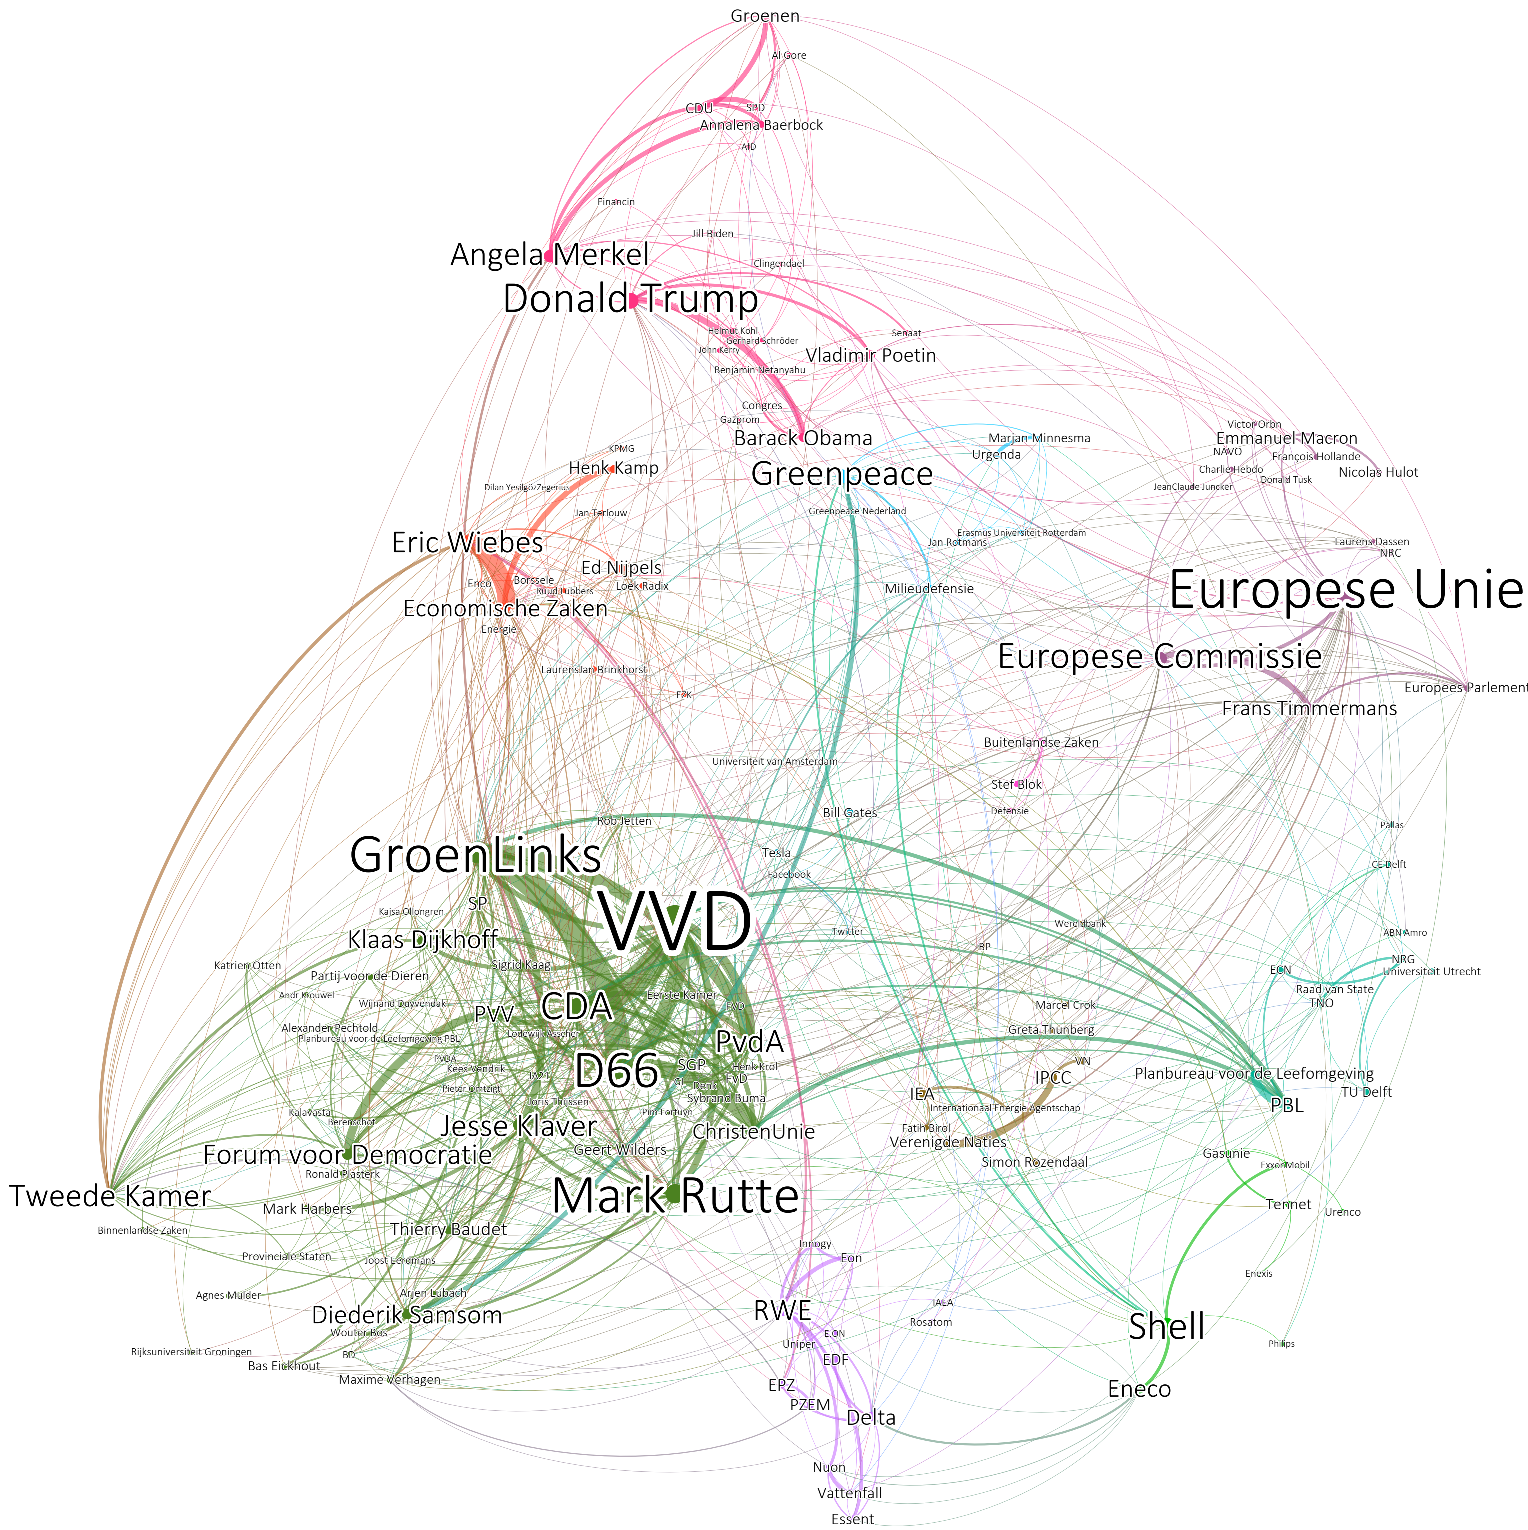

Figure 12. Partial network visualization of people and organizations found across all articles. The connection strength is given by the number of co-occurences within sentences. A full, interactive visualization can be found here.

References

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.

Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Gallagher, R. J., Reing, K., Kale, D., and Ver Steeg, G. (2017). Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*, 5:529–542.

Honnibal, M., Montani, I., Van Landeghem, S., and Boyd, A. (2020). spaCy: Industrial-strength Natural Language Processing in Python. https://doi.org/10.5281/zenodo.1212303.

Hutto, C. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8.

Pritchard, J. K., Stephens, M., and Donnelly, P. (2000). Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959.

Sang, E. F. and De Meulder, F. (2003). Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*.

Schweter, S. and Akbik, A. (2020). Flert: Document-level features for named entity recognition. arXiv preprint arXiv:2011.06993. https://huggingface.co/flair/ner-dutch-large.

Speer, R., Chin, J., Lin, A., Jewett, S., and Nathan, L. (2018). Luminosoinsight/wordfreq: v2.2. https://doi.org/10.5281/zenodo.1443582.

Appendix

Confidence in results

This section details potential sources of errors in our analysis. Firstly, the goal of this analysis is to find out how the Dutch news media reports about nuclear energy. Given our optimized search string and manual selection of articles, our dataset is not some subsample of the relevant news articles but actually the whole population, i.e. all relevant articles published on the topic in the given time frame. Because we have data for the whole population the results need not be qualified with confidence intervals but are exact.

The only potential error source is therefore the methodology – how well does a sentiment score capture the true sentiment, how well does the topic model capture each topic (in terms of our idea of it), and how well does the named entity recognition detect entities. For both sentiment and named entity recognition we use out-of-the-box tools, the performance of which was tested on benchmark datasets (including news datasets, which would be most relevant comparison in our case).

The named entity recognition we use was tested on the CoNLL-2002 shared task, where it correctly identified more than 95% of entities in the Dutch section of the task. The benchmark dataset is quite comparable to our application: The Dutch section consist of four editions of the Belgian newspaper "De Morgen", which were manually annotated (Sang and De Meulder, 2003).

Regarding sentiment analysis, the VADER algorithm performs as well as individual human raters, predicting the correct sentiment about 88% of the time, which was also confirmed on a news dataset of New York Times opinion editorials (Hutto and Gilbert, 2014). Of course, we first had to translate the Dutch texts to English, which is another potential error source. We cannot quantify the exact error rate of Google Translate, both because the algorithm is proprietary and because evaluating machine translation performance objectively is its own challenge. The state of the art in machine translation has become close to the performance of human translators. If the only goal is to use the translations for rule-based sentiment analysis, the required quality of translations is quite low - it is only important that the words found in the sentiment lexicon (e.g. dangerous, advantageous) are translated correctly, but not that the translation is grammatically correct. It is safe to say that the Google Translate algorithm would pass these requirements with a very low error rate.

Finally, the performance of the custom-built topic model is a bit more difficult to quantify since we do not have a labeled dataset to test on. We can only judge the performance qualitatively by looking at the topic keywords and checking them against our conception of the topics. The list of topic keywords can be found in the appendix. The keywords which were not used as anchor words (and therefore generated by the model) are highly consistent and specific to each topic, indicating that the topic model performs well and that predictions on which topic is mentioned in a given sentence will be highly accurate.

Table of topic keywords

The following is a list of the top keywords associated with each topic. The keywords are ordered by their mutual information score, a score that indicates how specific/informative a word is for a given topic. Only keywords above a mutual information threshold of 0.2 are shown, but never more than 100 keywords per topic. The number of mentions for a given keyword is given in parentheses, and keywords with an asterisk were used as anchor words for the topic.

Climate impact: uitstoot* (1338), co2* (1398), broeikasgassen* (568), opwarming* (717), duurzame* (2091), duurzaam* (701), broeikasgas* (192), leefomgeving* (259), ipcc* (275), klimaatbeleid* (603), klimaatakkoord* (826), verduurzaming* (162), klimaatverandering* (1519), verduurzamen* (147), atmosfeer* (188), temperatuurstijging* (81), klimaatneutraal* (183), broeikaseffect* (60), duurzaamheid* (412), klimaatprobleem* (263), emissie* (37), klimaat* (2000), uitstoot broeikasgassen (284), hernieuwbaar* (72), opwarmen* (26), broeikasgasuitstoot* (21), uitstoot co2 (187), fossiele (1393), 2030 (1152), duurzamer* (78), 2050 (946), energietransitie* (966), 1990 (319), fossiele brandstoffen (752), uitstoten (176), planbureau leefomgeving (179), planbureau (259), reductie (186), doelen (358), beperken (250), opzichte 1990 (107), klimaatpanel (90), uitstoot broeikasgas (63), wereldwijde (342), verminderen (222), emissies (129), co2 uitstoot (59), terugdringen (176), opzichte (306), climate (131), graad (160), 2020 (567), doelstellingen (238), doelstelling (220), broeikasgas co2 (45), decennia (410), transport (194), co2 uitstoten (78), verbranding (97), reduceren (115), opwarming beperken (58), transitie (309), klimaatwet (177), milieuorganisaties (148)

Waste and storage: afval* (929), opslag* (517), opslaan* (145), ondergronds* (78), covra* (73), restwarmte* (55), eindberging* (32), opbergen* (22), radioactief afval (228), berging* (13), kernafval (396), radioactieve afval (53), zoutkoepel* (13), halfwaardetijd* (5), opgeslagen (167), ondergrondse opslag (42), nadelen (227), opslag kernafval (36), ondergrondse (108), hoogradioactief (39), opgeborgen (44), hoogradioactief afval (34), kleilagen (27), afval opgeslagen (22), turkenburg (109), nucleaire afval (19), zoutkoepels (40), geaccepteerd (37), giftig (29), wim turkenburg (32), hoogradioactieve (12), overhoudt (11), opslagplaats (17), afval covra (11), toekomstige (297), afvalprobleem (45)

Geopolitics: kernwapens* (246), poetin* (251), iran* (303), conflict* (98), sancties* (133), geopolitiek* (71), militair* (46), militairen* (29), ontwapening* (18), president (791), arsenaal* (26), internationale (682), obama (307), verrijking* (20), staten (719), amerikanen (184), vladimir (55), vladimir poetin (49), russen (127), barack obama (58), barack (58), moskou (103), trump (666), koude (138), president poetin (37), donald (183), 2009 (264), donald trump (147), president vladimir (25), president obama (62), navo (93), syrië (104), congres (115), krim (39), oekraïne (222), iraanse (80), washington (97), 2012 (346), dreiging (106), trumps (79), teheran (35), clingendael (77), republikeinse (63), irak (62), conflicten (54), 2005 (153), wereldleiders (69), steve

(33), gevaar (294), clinton (40), spanningen (52), belangen (218), defensie (99), leiders (169), republikeinen (47), security (30), jinping (52), raketten (32), summit (27)

Safety: fukushima* (663), ramp* (519), tsjernobyl* (394), straling* (336), tsunami* (160), aardbeving* (83), radioactief* (375), kanker* (76), veiligheid* (481), terroristen* (60), beveiliging* (57), explosie* (48), terrorisme* (85), aanslag* (87), meltdown* (38), dosis* (26), dode* (61), catastrofe* (54), kernsmelting* (15), kernramp (252), evacueren* (12), ramp fukushima (107), ongeval* (36), 2011 (419), radioactieve (225), nucleaire (1081), 1986 (138), kernramp fukushima (76), millisievert* (13), radioactieve straling (51), geëvacueerd (34), nucleaire ramp (28), stoffen (153), dodelijke (50), gevaren (129), three mile (32), ontplofte (28), mile island (29), kernrampen (30), ontploffing (21), zeewater (49), nucleaire straling (15), schildklierkanker (16), veiligheidsmaatregelen (18), radioactieve wolk (12), voorzorgsmaatregelen (15), mensenlevens (25)

Cost: miljard* (1433), euro* (1806), miljoen* (1331), geld* (1476), kost* (552), prijs* (670), subsidie* (482), investeren* (636), investeringen* (593), duur* (540), goedkoop* (246), investering* (221), capaciteit* (208), kilowattuur* (129), economisch* (290), miljard euro (684), investeerders* (178), financiering* (112), cent* (129), financieel* (178), rendabel* (133), kwh* (111), ontmanteling* (80), investeert* (107), financieren* (72), miljoen euro (359), banken* (112), financiers* (30), uitgaven* (54), investeringsbank* (25), bank* (189), financiën* (113), kosten (1370), financieele* (64), business case* (14), levensduur* (120), gefinancierd* (55), miljarden (426), 100 (453), bedrag (134), investeerder* (35), overheid (1360), miljoenen (313), banen (308), goedkope (244), bedragen (95), 100 miljoen (57), tekort (170), lagere (200), subsidies (315), infrastructuur (315), kost miljard (46), fors (220), 100 miljard (41), denktank (109), daalt (124), gigawatt (157), kubieke (97), lager (269), stroomprijs (79), oplopen (86), dure (173), minstens (190), energierekening (111), groeiende (227), geïnvesteerd (193), stijgende (128), omlaag (207), bouwt (148), circa (127), euro subsidie (33), uitgeven (66), kapitaal (53), aandelen (90), aanleg (174), bouwkosten (40), commercieel (67)

Technology: uranium* (395), reactor* (414), reactoren* (276), thorium* (277), kernreactoren* (170), technologie* (664), ontwerpen* (45), thoriumreactor* (58), splijtstof* (48), reactorvat* (35), koeling* (29), kernsplitsing* (26), kernsplijting* (23), kettingreactie* (20), smr* (34), modulaire* (20), zoutreactor* (16), koelmiddel* (13), neutronen* (21), veiliger (189), fusiereactor* (16), plutonium (56), petten (119), finland (197), gesmolten (64), thoriumcentrales (83), nucleaire technologie (36), gesmolten zout (32), frankrijk (761), dodewaard (129), kostenoverschrijdingen (47), uranium thorium (23), isotopen (61), verrijkt (40), thorium brandstof (20), kloosterman (92), beton (62), veilige (160), leen kloosterman (34), flamanville (32), kernfusie (98), veiligste (70), urenco (73), reactorkern (18), iaea (41), ingenieur (78), winning uranium (13), centrifuges (13), modular (10), salt (10), verrijking uranium (9), gesmoltenzoutreactor (11), flux (10), flux reactor (10)

Ethics: kinderen* (722), mensheid* (303), wetenschap* (533), kleinkinderen* (90), filosoof* (126), waarden* (129), filosofie* (77), religie* (71), generaties* (271), moreel* (49), ethiek* (67), generatie* (365), principes* (39), filosofen* (28), moraal* (23), socioloog* (30), ethisch* (19), principe* (100), intellectueel* (25), planeet (348), wetenschappers (427), wetenschappelijke (265), morele (137), vooruitgang (253), genetische (80), hoogleraar (457), wetenschappelijk (177), universiteit (472), zeventig (218), tegelijk (265), culturele (56), welvaart (205), technologische (202), wetenschapper (97), kind (183), achteraf (112), term (118), tijdperk (111), revolutie (194), leeftijd (93), journalisten (116), begrip (144), armoede (172), kapitalisme (120), kranten (79), beschaving (68), ethische (54), durven (108), belangrijker (171), eeuwen (92), menselijk (114), ecologische (114), houding (143), gezondheidszorg (76), afweging* (56), genetisch (65), modificatie (41), genetische modificatie (41), denkers (50), migranten (49), planten (173), telkens (112), leerde (48), optimisme (105), werkkamer (33), christelijke (91), essay (60), veranderde (63), kunstmatige (63), ideologie (54), literatuur (39), culturen (35), optimisten (24), immigratie (65), pessimisten (24), perspectief (141), toekomstige generaties (99)

Politics: vvd* (1829), cda* (871), d66* (938), partij* (1520), pvda* (709), groenlinks* (902), verkiezingen* (531), kamer* (959), kabinet* (1107), rutte* (936), coalitie* (291), verkiezingsprogramma* (243), democratie* (487), links* (457), formatie* (138), rechts* (288), partijen (1460), pvv (378), vvd cda (253), christenunie (273), forum democratie (221), linkse (324), klaver (423), zetels (205), verkiezing* (42), fvd (197), politieke (1309), cda d66 (90), d66 groenlinks (100), lijsttrekker (133), kiezers (336), forum (396), kamerlid (320), jesse klaver (141), d66 christenunie (76), jesse (155), sgp (213), fractievoorzitter (135), motie (195), peilingen (146), kiezer (130), dijkhoff (334), rechtse (172), meerderheid (356), liberalen (84), groenlinks pvda (60), cda vvd (58), geert (92), politieke partijen (180), kamerverkiezingen (89), wilders (152), lijsttrekkers (72), pvda d66 (41), regeerakkoord (156), geert wilders (68), coalitiepartijen (57), klaas dijkhoff (164), 50plus (61), klaas (207), partijleider (71), partijen vvd (41), vvd pvda (42), pechtold (61), cda pvv (42), groenlinks d66 (65), vvd d66 (65), pvda groenlinks (46), baudet (203), kamerleden (107), achterban (98), links rechts (82), veestapel (80), diederik (182), buma (107), cda pvda (40), linkse partijen (69), samsom (375), alexander pechtold (35), kabinetsperiode (39), rob jetten (47), verkiezingscampagne (69), jetten (74), partij vvd (30), klaver groenlinks (31), vvd pvv (32), d66 pvda (39), diederik samsom (133), migratie (89), vvd partij (31), voorstander (207), pvv forum (31), vvd groenlinks (29), coalitie vvd (31), regeringspartij (51), statenverkiezingen (44), kaag (108), harbers (100), sigrid (47), sigrid kaag (47), kabinetten (57)

Choice of site: provincie* (759), zeeland* (571), eemshaven* (370), borssele* (960), maasvlakte* (151), groningen* (699), brabant* (234), locaties* (145), gemeente* (451), zeeuwse (274), gedeputeerde (199), locatiekeuze* (15), provinciale (242), bouwlocaties* (8), provinciebestuur (95), 2033 (128), provinciale staten (102), epz (196), delta (377), pzem (192), provincies (187), rwe (482), draagvlak (458), provincie zeeland (42), energiebedrijf (326), regionale (281), provincie groningen (42), eemshaven maasvlakte (29), aandeelhouders (189), limburgse (95), middelburg (78), drenthe (66), brabantse (102), groningse (143), landschap (319), energiebedrijf delta (35), rijksoverheid (64), vlissingen (64), landelijk (129), zeeuwen (46), eric bie (25), leveringszekerheid (102), tno (92), maasvlakte eemshaven (18), ruimtelijke (83), gelderland (63), essent (153), carlo wolters (24), westerschelde (33), friesland (80), aandeelhouders pzem (19), groningen maasvlakte (16), maasvlakte borssele (17), haven (138), vergunning (102), overname (66), provinciaal (30), overijssel (39)