



**University of
Zurich**^{UZH}

Geospatial Sentiment Analysis of different Energy Sources using historical Twitter Data

GEO 511 Master's Thesis

Author

Marco Stutz
17-704-818

Supervised by

Prof. Dr. Ross Purves


Faculty representative

Prof. Dr. Ross Purves

31.01.2024

Department of Geography, University of Zurich

Geospatial Sentiment Analysis of different Energy Sources using historical Twitter Data



Marco Stutz

Supervisor: Prof. Dr. Ross Purves

Faculty representative: Prof. Dr. Ross Purves

University of Zurich

Department of Geography

Geocomputation

January 31, 2024

Acknowledgements

After many interesting years that taught me so much more than just geographic knowledge, after thousands of funny moments and dozens of kind, helpful and respectful people I met here, this Master's thesis is my last step at the Department of Geography. And it's definitely one of the toughest but also one of the most fascinating ones. All in all, I'm extremely glad to say that I would go the exact same route again if time machines existed.

Among lots of other people, my supervisor Prof. Dr. Ross Purves has been a part of this route as well. I want to express my biggest thanks for his excellent support, his sophisticated suggestions and his tolerance regarding my mistakes and own ideas. This doesn't only apply to this thesis but also to all the courses he lectured during my studies.

Moreover, I want to thank my GIS colleagues for being a part of this experience as well. I'm super happy our lifelines crossed at the GIUZ and will reminisce about all the awesome moments we had during our time at the Irchel campus. Thank you!

Lastly, I would like to express my sincere thanks to my Swiss army fellow Gfr Sager who actively advised me during the final layout process.

Abstract

In this thesis, the sentiment of Twitter users from Germany, Switzerland and Austria about various energy sources between 2007 and 2023 is monitored, while underlying causes of sentiment changes and country-specific developments are further examined. To do so, an aspect-based sentiment model on the basis of a German BERT is fine-tuned for this specific sentiment analysis task and later applied on more than 2.6 million tweets. Furthermore, term- and n-gram frequencies, topic modelling and microreading is performed to uncover latent emphases of discourses. Lastly, sentiments are geo-spatially separated via a home location prediction of Twitter users. It was found that coal and nuclear energy were mostly negatively perceived while Twitter had more or less neutral sentiments about gas and wind energy and positive sentiments about water and especially solar power. Sentiment variations were caused by supranational as well as local events, disasters, political decisions, media articles, social movements and geopolitical incidents. Most of these led to a shift in priorities, either emphasising advantages or drawbacks of the energy sources. The latest energy crisis intensified by the Russian invasion shifted priorities away from environmental and climate arguments towards a strong focus on affordable and independent energy supply, strengthening coal and nuclear power.

Keywords: Climate change, energy crisis, energy strategy, energy transition, public support, sentiment analysis, Twitter

Contents

1	Introduction	1
1.1	Motivation and relevance	1
1.2	Research aims	2
1.3	Outline	4
2	Background and literature	5
2.1	Energy sources: A pervasive dilemma	5
2.1.1	Characteristics of energy classes	5
2.1.2	Climate action and public support	6
2.2	Twitter as a data source	9
2.2.1	Basic information	9
2.2.2	Advantages of using Twitter for research	9
2.2.3	Cautions and limitations	10
2.3	Sentiment Analysis	12
2.3.1	History and evolution	12
2.3.2	Transformer-based BERT models	13
2.3.3	Aspect-based sentiment analysis	14
2.3.4	German sentiment analysis regarding energy perception	15
2.4	Content Analysis	17
2.4.1	Term frequency methods	17
2.4.2	Topic modelling	18
2.4.3	Phrase extraction	18
2.5	Geospatial separation	20
2.5.1	Geoparsing	20
2.5.2	Toponym recognition	21
2.5.3	Toponym resolution	23
3	Methods	25
3.1	Thesis extent	25
3.2	Data download and software	26
3.3	Preprocessing	29
3.3.1	Duplicates removal	29
3.3.2	Bot detection	30
3.3.3	Aspect separation	35
3.3.4	Thematical relevance	36
3.4	Sentiment Analysis	38
3.4.1	Unsuitable document-level sentiment analysis models	38
3.4.2	A custom aspect-based sentiment analysis model	38

3.4.3	Applying the fine-tuned ABSA model	41
3.5	Content Analysis	42
3.5.1	Topic modelling	42
3.5.2	Term frequency methods	43
3.5.3	Micro-reading	43
3.6	Geospatial separation	45
3.6.1	Assumption to infer the home location from the tweet content	45
3.6.2	Geoparsing, disambiguation and inclusion of GPS tags	45
4	Model performance	47
4.1	Sentiment analysis model	47
4.2	Geospatial separation approach	48
5	Results	49
5.1	Sentiment of all German-speaking Twitter users	49
5.2	Country-specific sentiment differences	50
5.3	Nuclear energy	52
5.4	Coal energy	58
5.5	Solar energy	65
5.6	Wind energy	71
5.7	Water energy	77
5.8	Gas/oil energy	83
6	Discussion	89
6.1	Main findings and evaluation of hypotheses	89
6.2	The clash of interests	90
6.2.1	Fossil fuels: Constant energy supply as the carrying argument	91
6.2.2	Renewable energy sources: Positively perceived with one limiting factor	94
6.2.3	Nuclear power: A special case	97
6.2.4	Implications for public support and the renewable energy transition	99
6.3	Methodological limitations and validity	103
6.3.1	Data retrieval and preprocessing	103
6.3.2	Sentiment analysis	103
6.3.3	Content Analysis	104
6.3.4	Geospatial separation	104
6.3.5	Validity	105
7	Conclusion and further work	107
8	Bibliography	108

9 Appendix

132

List of Figures

1	The process of Geoparsing.	21
2	The methodic steps conducted in this thesis.	25
3	The temporal distribution of tweets per energy source.	26
4	The final sentiment timeline showing daily sentiments about the chosen energy sources and respective uncertainties visualised via buffers. Graphs are smoothed over 365 days, hence not starting in 2007.	49
5	The country-specific monthly average sentiments between 2007 and 2023 (slightly smoothed by a Gaussian filter).	51
6	The sentiment timeline for nuclear energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	52
7	The monthly average sentiment about nuclear energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	53
8	The temporal evolution of the ten different topics for nuclear energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.	54
9	The word clouds for positive and negative tweets about nuclear energy.	55
10	The average sentiment per nuclear topic of the whole study period.	56
11	The country-specific sentiment timeline for nuclear energy per country.	57
12	The sentiment timeline for coal energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	58
13	The monthly average sentiment about coal energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	59
14	The temporal evolution of the ten different topics for coal energy and their characteristic words derived from the LDA model, slightly smoothed by a Gaussian filter.	60
15	The word clouds for positive and negative tweets about coal energy. Note that 'braunkohl' appeared twice as it was part of connected words like 'braunkohl-abbau' which couldn't be handled by the word cloud library.	62
16	The average sentiment per coal topic of the whole study period.	62
17	The country-specific sentiment timelines for coal energy.	64
18	The sentiment timeline for solar energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	65
19	The monthly average sentiment about solar energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	66

20	The temporal evolution of the ten different topics for solar energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.	66
21	The word clouds for positive and negative tweets about solar energy.	67
22	The average sentiment per solar topic of the whole study period.	68
23	The country-specific sentiment timelines for solar energy.	70
24	The sentiment timeline for wind energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	71
25	The monthly average sentiment about wind energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	72
26	The temporal evolution of the ten different topics for wind energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.	72
27	The word clouds for positive and negative tweets about wind energy.	74
28	The average sentiment per wind topic of the whole study period.	75
29	The country-specific sentiment timelines for wind energy.	76
30	The sentiment timeline for water energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	77
31	The monthly average sentiment about water energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	78
32	The temporal evolution of the ten different topics for water energy derived from the LDA model, slightly smoothed by a Gaussian filter.	79
33	The average sentiment per water topic of the whole study period.	80
34	The country-specific sentiment timelines for water energy.	82
35	The sentiment timeline for gas/oil energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.	83
36	The monthly average sentiment about gas/oil energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).	84
37	The temporal evolution of the ten different topics for gas/oil energy derived from the LDA model, slightly smoothed by a Gaussian filter.	85
38	The word clouds for positive and negative tweets about gas/oil energy.	86
39	The average sentiment per gas/oil topic of the whole study period.	86
40	The country-specific sentiment timelines for gas/oil energy.	88

41	The four typical dimensions of arguments made by people to express their sentiment about energy types (with example keywords). Since no energy source can fully cover all four dimensions, a trade-off emerges.	91
42	The final sentiment timeline accompanied by the underlying dimensions and their temporally variable importance observed from the data. Solid lines represent a greater importance than dotted lines.	99
43	Shift of prioritized dimensions observed in data for nuclear power. Starting in 2018, the climate and energy supply topics replaced previously dominant environmental and health arguments.	133

The photos on the title page: Photo (left) by Brina Blum, 2016, <https://unsplash.com/photos/gray-building-with-smoke-under-white-clouds-rsrhoWJe8j0>; Photo (middle) by Karsten Würth, 2016, <https://unsplash.com/photos/windmill-on-grass-field-during-golden-hour-0w-uTa0Xz7w>; Photo (right) by Andreas Gücklhorn, 2017, <https://unsplash.com/photos/aerial-photography-of-grass-field-with-blue-solar-panels-Ilpf2eUPpUE> (Unsplash license).

List of Tables

1	The attributes downloaded by the Twitter API.	26
2	Search queries defined to retrieve tweets via the API.	28
3	The indicators used to train the bot detection model.	33
4	Performance of the used bot detection model on the Botometer labelled testing set. .	34
5	The total amount of remaining tweets after the pre-processing tasks.	37
6	Final hyperparameters chosen to train the aspect-based sentiment analysis model. .	40
7	Sentiment model performance on the testing data.	47
8	Performance of the home location predictor.	48

List of abbreviations

AI	Artificial intelligence
API	Application Programming Interface
BOW	Bag-of-Words
BERT	Bidirectional Encoder Representations from Transformers
CNN	Convolutional neural network
EU	European Union
GIR	Geographic Information Retrieval
GPS	Global Positioning System
IAT	Inter-arrival time
IPCC	Intergovernmental Panel on Climate Change
LDA	Latent Dirichlet Allocation
ML	Machine Learning
NER	Named Entity Recognition
NLP	Natural Language Processing
POS	Part-of-speech
PV	Photovoltaic
RF	Random Forest
TF-IDF	Term frequency - inverse document frequency
USA	United States of America

List of condensed terms

Energy words	Keywords referring to a predefined energy source
Energy tweets	Tweets that mention at least one of the predefined energy sources
Water energy	Energy generated from the power of water (analogous for other sources)
Water sentiment	The sentiment about water energy (analogous for other sources)

1 Introduction

1.1 Motivation and relevance

“Energy is everywhere and drives everything. Our modern lives, both individual and societal, have come to depend on its abundance, convenience, and potential. It is the motive force within our bodies, propelling our vehicles, lighting our world. Consider a power outage, or a dead cell phone battery; living without energy, for even ten minutes, demonstrates how indelible its imprint is on daily activities.” (Coyle & Simmons, 2014, p. 1)

This dramatically epic paragraph serves as the introduction of Eugene Coyle’s and Richard Simmons’ book *Understanding the Global Energy Crisis*. It impressively shows the societal value of energy and how our civilization is strongly dependent on a reliable energy supply, whether it is for transportation, mobility, food preparation, heating, water purification or many more needs (Kalt et al., 2019). The crisis that the authors refer to was already present back in 2014 when their book was published, evoked by the contradiction between rising energy demand due to population growth and economic expansion (Farghali et al., 2023) and the urgent need for sustainability to mitigate the severe impacts of climate change (Owusu & Asumadu-Sarkodie, 2016). Since the latter is not possible by continuing current activities where more than 80% of global energy stems from fossil fuels which do also contribute to 90% of all carbon dioxide emissions (United Nations, 2022), a sustainable energy supply became one of the world’s central challenges. And this challenge became even more urgent with the Russian invasion of Ukraine (Farghali et al., 2023). On a global scale, most countries are dependent on fossil fuel imports, leading to over 6 billion people that are potentially vulnerable to such geopolitical conflicts (UNECE, 2022). Central Europe, especially Germany is heavily dependent on Russian gas (Vrana et al., 2023) while most other EU countries do also import natural gas (Zachmann et al., 2022). Due to the cut relationship with Russia and introduced sanctions, Russia reduced or even cut off gas exports to certain countries in the EU (Chyong et al., 2023). Consequently, gas prices skyrocketed, resulting in winter 2022’s energy price being ten times higher than the average of the past five years (Jayanti, 2022). These increases in the market price forced companies to cut production and lay off thousands of employees while citizens of most countries had to spend a record proportion of their income on energy (Vrana et al., 2023). Not surprisingly, a survey executed in Germany in 2022 found that 83% of the participants started to save energy, especially targeting the heating system (Siebel, 2022).

In parallel, the same survey documented the priorities of the participants. Most of them assigned the biggest importance to energy security, closely followed by low energy prices. Only 12% of all participants ranked climate and environmental protection higher than the aforementioned issues (Siebel, 2022). This corresponds to results from research about renewable energy support mechanisms which found a positive correlation between financial resources and the willingness to support renewable energy solutions (Sokołowski et al., 2023), also known as the *End of the world vs. end of the month* dichotomy (Martin & Islar, 2021). In the context of the energy crisis, long

ongoing debates and newly emerged ones about particular energy sources, energy usage and energy security reignited in newspaper articles, on social media platforms and in private settings. Due to the necessity of sustainable energy production, geopolitical gas and oil boycotts and risen price levels for traditional energy, the potential of renewable energy sources, fossil and nuclear fuels is intensively discussed, leading to a change in support for various energy policies, as a recent survey in Switzerland showed (Steffen & Patt, 2022).

Such opinion monitoring practices are important to track the public mood regarding different topics, in this case, energy sources and their use strategies. Especially during crisis situations, being aware of public views helps government officials to communicate accordingly, helping to avoid panic (Siebel, 2022). Furthermore, and more importantly, public support for particular energy sources is key to legitimising their usage, building a stepping stone for its successful implementation and financial maintenance of respective energy projects (Dehler-Holland et al., 2022). On the contrary, disapproval and negative sentiments towards certain energy types and corresponding decisions taken can be significant barriers when it comes to achieving energy targets, especially in stable democracies (Segreto et al., 2020). A prominent example of the consequences of lacking public support regarding climate policies were the *Yellow Vests* in France (Tatham & Peters, 2023). As a bottom-up movement, their protests and riots eventually led to the withdrawal of a planned diesel tax (Sokołowski et al., 2023). Hence, monitoring the sentiment of citizens and its change following policy measures or events can give crucial information about the choice of suitable future energy policies and is, thus, a necessity that should not be neglected to achieve the sustainable energy transition.

1.2 Research aims

Traditionally, opinion polling about public support of energy sources is mostly conducted through surveys (Kallbekken, 2023; S. Y. Kim et al., 2021; T. M. Lee et al., 2015; Y. Liu et al., 2019; Tatham & Peters, 2023; Vringer & Carabain, 2020). Although this is a widely used data collection method, it has several crucial limitations like a low temporal resolution (Loureiro & Alló, 2020) and small response rates (Vaske et al., 2023). For instance, a recent online survey about climate change awareness and preferred energy sources carried out by the famous Swiss polling institute *Sotomo* only reached less than 4000 participants (Lanz & Morgenthaler, 2022). Hence, Kallbekken (2023) pleads for more diverse research methods as it is so important to monitor public support and concerns of people regarding energy and climate policies. As alternatives to surveys, the author names conjoint analyses, laboratory experiments, interviews, nature field experiments or machine learning techniques. The application of the latter has risen sharply in recent years (Khairnar & Kinikar, 2013) as several researchers started using social media posts to gather public opinions on various topics. Especially Twitter¹, a famous microblogging platform, was found to have great potential as an information source (Ahmed et al., 2017) as users – voluntarily – share their opin-

¹I will refer to the social media platform as *Twitter* throughout the thesis although it has been rebranded to *X* by the end of July 2023 (Mac & Hsu, 2023).

ions and feelings on any possible aspects of life (Pak, Paroubek, et al., 2010). Hence, tweets may contain all sorts of basic emotions and sentiments about various topics (Garske et al., 2021). With hundreds of millions active daily users (Ahlgren, 2023; Dean, 2022), Twitter is a rich source of such information. Since it is not primarily a direct messaging service, most users have publicly open profiles and participate in discussions of any kind, as they can not only post their own tweets but also write responses, retweet posts or quote other tweets (Zheng et al., 2018). Recently, due to the tensions arisen from the Russia-Ukraine conflict, discussions about energy supply and policy have also intensified on Twitter. This creates the possibility to examine Twitter users' current but also historic sentiments about different energy sources and, thus, gives a hint at the degree of public support and its temporal development.

In accordance to the aforementioned necessity of public support detection, this thesis uses historical Twitter data to perform a geospatial *sentiment analysis*. The sentiments of German-speaking Twitter users about different energy sources are monitored between January 1, 2007 and December 31, 2022. The goal is to detect periods of positive and negative perceptions about the aforementioned energy types and to uncover possible drivers of temporal sentiment variations. Furthermore, the thesis aims to assess these information in a geospatial manner, uncovering country-specific sentiment changes (Austria, Germany and Switzerland). Derived from the observed sentiments, this should allow insights into the level of support of the energy sources, its development over time and its spatial differences, and finally to assess the chances of success of possible energy policy measures per country on the way to a sustainable energy future. Accordingly, the research questions and the respective hypotheses are as follows:

RQ1: How did the sentiments of German-speaking Twitter users towards different energy sources change between 2007 and 2023?

H1: Due to the rising climate change awareness, renewable energy sources and nuclear power gained popularity over the years while the sentiment about fossil fuels became more negative.

RQ2: Which events or circumstances were responsible for these sentiment variations?

H2: Energy availability, prices and projects, political decisions and popular movements, disasters or technological progression are driving forces behind sentiment variations.

RQ3: How did sentiments differ between Twitter users from Germany, Austria and Switzerland?

H3: Country-specific differences of sentiments mainly base on the respective energy mix of the country and national or regional events. Hence, more positive sentiments about hydropower are expected for Swiss and Austrian users while more positive perception of coal power is expected for German users.

1.3 Outline

The following chapter builds the theoretical frame of the thesis and summarises existing methods and concepts found in the literature regarding the main tasks. Chapter 3 then describes the data and the respective methods chosen while chapter 4 quickly documents the performance of the sentiment model and the geospatial separation approach. Chapter 5 then presents the results regarding the three research question which are then, in chapter 6, compared with existing findings from the literature and discussed under a broader perspective, also touching on public support and implications for the renewable energy transition. Moreover, the main limitations and sources of uncertainties will be discussed. Lastly, the findings of the thesis are concluded in chapter 7.

2 Background and literature

2.1 Energy sources: A pervasive dilemma

The importance of energy for human civilisation can barely be put into words as it is the prerequisite for our every-day lives (Coyle & Simmons, 2014). Generally, sources of energy can be split into three main classes: fossil fuels, nuclear sources and renewable energy (Qazi et al., 2019; Rahimnejad et al., 2015). Each of them have great advantages but also major disadvantages, forming the basis of ongoing energy debates.

2.1.1 Characteristics of energy classes

Fossil fuels

Coal, gas and oil are the most common fossil fuels (National Geographic Education, 2023). Originated from decomposed plants and putrid animals, they contain carbon and hydrogen which releases massive amounts of energy during combustion (National Geographic Education, 2023). Since energy production is independent of weather conditions and energy can easily be stored, fossil fuels are considered reliable regarding energy security (Roggenkamp et al., 2021). In combination with this efficiency and controllability, they're rather cheap energy sources (Barreto, 2018) and therefore responsible for 80% up to 90% of the global energy supply (Celik, 2021; Nicoletti et al., 2015; Okedu, 2018; United Nations, 2022). However, combustion products – especially carbon dioxide (Rahimnejad et al., 2012) – do have major negative impacts on the greenhouse effect, the ozone layer, acid rains and environmental pollution (Nicoletti et al., 2015). In fact, they're the largest contributor to greenhouse gas emission (IPCC, 2023) and major drivers of climate change (Ang et al., 2022). Additionally, they can have severe negative health consequences (Gasparotto & Martinello, 2021).

Nuclear sources

Another popular energy source is nuclear energy. Just like fossil sources, the utilised natural uranium is also non-renewable (Rahimnejad et al., 2015). According to the IPCC, uranium resources are expected to last another 120 years (Bruckner et al., 2014). Compared to the aforementioned fossil fuels, the big advantage of nuclear energy is its climate- and ecological friendliness. As it's driven by nuclear fission, barely any carbon dioxide (Amponsah et al., 2014; Dehler-Holland et al., 2022; Vaillancourt et al., 2008) or pollutive gases are emitted (Turconi et al., 2013). Nevertheless, health can still be at a massive risk due to the radioactive material that is released (Dai et al., 2019; Dehler-Holland et al., 2022). Cardis et al. (2007) found that cancer risk increased when people are exposed to radioactive radiation while this risk is dependent on the duration of exposure. These problems continue once the radioactive fuel rods have to be disposed (Bruckner et al., 2014). While accidents are quite rare, implications are exorbitant and long-lasting if something happens as seen in Chernobyl 1986 or Fukushima 2011 (Aub et al., 1952; Morino et al., 2011; Wheeler, 1988).

Renewable energy

Specific examples of renewable energy sources are solar energy, hydro energy, wind energy, bioenergy, geothermal energy or hydrogen energy (Mohtasham, 2015). In opposition to fossil fuels and nuclear energy, they're sustainable since they use regenerative natural sources – predominantly driven by solar insolation and geothermal heat – for energy production (Turkenburg, Faaij, et al., 2000). They provide power with almost zero emission of pollutants or greenhouse gases which makes them by far the most climate- and eco-friendly energy sources (Ang et al., 2022; Mardani et al., 2015; Mohtasham, 2015). Hence, they are indispensable if climate change is to be combated (IPCC, 2023). However, their dependence on natural forces is also a major drawback since it leads to an intermittent and unreliable energy supply (Ang et al., 2022; Azarpour et al., 2013; Behabtu et al., 2020; Leonard et al., 2020). Another prevalent issue is the efficiency as the naturally originated energy has to be converted into other forms of energy (e.g. electrical, thermal or chemical) which inevitably leads to a reduction of energy potential (Azarpour et al., 2013; Moriarty & Honnery, 2012). This makes storage a lot more complex compared to oil barrels, coal or gas (Moriarty & Honnery, 2012). According to thorough calculations by Holechek et al. (2022), renewable energy sources would theoretically only be able to fully replace fossil fuels if nuclear power is expanded by at least 30% globally.

All in all, none of the available energy sources can meet the needs of efficiency, affordability, storage, reliability, environmental protection, climate compatibility and sustainability at the same time. It's a pervasive dilemma that requires compromise and prioritization.

2.1.2 Climate action and public support

While the priority in the 20th century was clearly on a reliable energy security, this changed in the past decades, driven by the urge to fight climate change. Since the latter is a global challenge, it requires global actions (Vidadili et al., 2017). With the Paris Agreement in 2015, 197 countries have committed to ambitious efforts to limit global warming to +2°C compared to pre-industrial levels (United Nations, 2015) and reach a state of carbon neutrality by the middle of this century (Farghali et al., 2023). To achieve this, several transformative measures have to be taken evoking new economic, social, ecological and political challenges (Fuso Nerini et al., 2019). Among others, a transition towards renewable, sustainable energy is necessary as the fossil fuel dominated energy sector has a large impact on climate change (IPCC, 2023). Thus, unsustainable energy sources have to decline and eventually vanish (Markard et al., 2023). Accordingly, the term *energy crisis* has gained more attention by various researchers in the past years. While Chevalier (2009) and Coyle and Simmons (2014) named increasing energy demand and continued dependence on fossil fuels (in consideration of climate change) as reasons for the energy crisis, Sokołowski et al. (2023) and Farghali et al. (2023) speak of 2022 as the beginning of the energy crisis following the Russian invasion in Ukraine. In this thesis, I will refer to the dependence on fossil fuels and the urge for

renewable energy sources as *climate crisis* while *energy crisis* will refer to the challenged energy supply security caused by the invasion. As Europe is heavily dependent on Russian gas and oil, the latter led to a drastic increase in energy prices, which made financing heating and transportation costs difficult for many households, especially in Winter 2022 (Vrana et al., 2023). Although it has long been clear that fossil fuels were not compatible with the climate goals, these events even increased the pressure of finding capable energy solutions at a faster rate. In autumn 2022, Vrana et al. (2023) found German Twitter users to feel fearful regarding this uncertain and challenging situation.

The relevance of public support

Unsurprisingly, energy policies proposed to tackle the urgent challenges can have large impacts on citizens (Vringer & Carabain, 2020) and can cause social tension and division (Sokolowski et al., 2023). In extreme cases, climate action programs even resulted in a forceful displacement of vulnerable social groups (Fuso Nerini et al., 2019). If such interventions are too drastic, social acceptance diminishes and people start to resist, which can lead to policy goals not being effective or not being achieved at all (Vringer & Carabain, 2020; Wüstenhagen et al., 2007), see *Yellow Vests* in France (Kallbekken, 2023; Sokolowski et al., 2023). Hence, while the focus was primarily on market acceptance for a long time, *social acceptance* or *public support* started to gain more and more attention in recent years (Wüstenhagen et al., 2007). However, according to Scharpf (1999)'s very famous book called *Governing in Europe: Effective and Democratic?*, a further distinction between public support for the goal (*input legitimacy*) and public support for the associated interventions (*output legitimacy*) has to be made. So, public support for a sustainable energy transition doesn't automatically imply public support of interventions or policies to achieve this goal (Tatham & Peters, 2023; Vringer & Carabain, 2020).

Aspects affecting public support for energy projects

Based on the previous chapter, it's crucial that politicians, customers or generally citizens can see the necessity and effectiveness of energy policies and interventions (Geels & Verhees, 2011; S. Y. Kim et al., 2021). Research has identified various factors that influence public support of energy projects. As per Vringer and Carabain (2020), interventions are easier to achieve if the outcome of the measures have a local or regional impact and lie in the near future. Moreover, contributing to the regional economy by involving local companies helps to strengthen their support of renewable energy projects (Olson-Hazboun et al., 2016; Vringer & Carabain, 2020).

As Holt (1999) already claimed in 1999, energy has to be affordable to people. Accordingly, energy costs were often found to have a decisive impact on the public support (Bunting, 2004; Jenkins et al., 2016; Teisl et al., 2015).

Another factor influencing public support is the risk an energy project poses to the citizens. As per Boudet (2019), it's not primarily the *objective risk* quantified by scientists and experts that

shapes the policy preferences of people. Instead, it's rather the perception of social, economic and environmental risks and benefits that influence people's opinions about specific energy sources and projects (Slovic, 1987). Additionally, aesthetics is another non-neglectable issue (Rand & Hoen, 2017). This is often attached to the specific place an energy project should be launched and is dependent on the current landscape usage (Cotton & Devine-Wright, 2013), the proximity to protected or population-dense areas (McAdam & Boudet, 2012) or past experiences with similar energy projects (Bugden et al., 2017).

It was further observed that character traits like age also played a significant role as younger people were willing to spend a larger proportion of their income on climate mitigation measures (Abdar et al., 2020; Sokołowski et al., 2023). Moreover, the individual beliefs, perception and the awareness of climate change and environmental risks (Dreyer et al., 2017; T. M. Lee et al., 2015; Noblet et al., 2015; Qazi et al., 2019), political ideology (Boudet, 2019), personal belief in contribution to climate goals (Olson-Hazboun et al., 2016; Vringer & Carabain, 2020) and trust in local authorities and policymakers (Segreto et al., 2020) impact support of renewable energy sources.

Due to the complexity of the energy crisis, the trade-offs between energy sources, the counter-productiveness of policies lacking public support and the versatility of factors driving the latter, opinions regarding the future of energy supply may differ between stakeholders and individuals (Komendantova & Neumueller, 2020). Hence, it's crucial to monitor existing public acceptance of energy sources to establish suitable and effective energy policies (Komendantova & Neumueller, 2020; Qazi et al., 2019). Therefore, knowledge of past sentiment changes following political decisions, projects and events is helpful as well. While the support for projects can't be directly inferred from public sentiments (Scharpf, 1999), positive sentiments are considered a necessary prerequisite for the support of projects. Thus, this thesis aims to monitor the sentiments and underlying drivers of sentiment variations.

2.2 Twitter as a data source

Most of the existing opinion mining studies were conducted using surveys or other non-automated methods. However, as Kallbekken (2023) suggests, more insights into public support dynamics can be provided using a variety of methods. One of them is natural language processing (NLP) which will be the focus of this chapter.

Microblogging platforms pose a great possibility to be used for opinion mining tasks via NLP. Herefore, especially Twitter was used very frequently for data collection purposes of various academic research fields (Karami et al., 2020). According to Karami et al. (2020), the growth of Twitter used in research was massive in the last decade. During their systematic review of English research papers, the authors found that Twitter was prominently used for sentiment analysis, social network analysis, big data mining, topic modeling and content analysis (Karami et al., 2020).

In the following paragraphs, some basic information about the platform, motivation for its choice for this thesis and some limitations will be explained.

2.2.1 Basic information

The first ever tweet was posted by the former Twitter CEO Jack Dorsey on March 21, 2006 (Dean, 2022). Three years and two months later, one billion total so-called *tweets* have been posted (Dean, 2022). While 1.3 billion Twitter accounts exist nowadays, 368.4 million users are active, posting approximately 500 million tweets per day (Ahlgren, 2023). While Twitter has become a worldwide phenomenon, USA, Japan, India, Indonesia and Brazil account for the biggest number of users (Ahlgren, 2023). Users can not only post tweets, but also respond to other tweets, engage in discussions, repost a tweet (*retweet*), like a tweet or tag other users (Arigo et al., 2018). Until October 2018, the character limit was restricted to 140 and then expanded to 280 (Karami et al., 2020).

2.2.2 Advantages of using Twitter for research

Based on these numbers, it seems logical that Twitter’s biggest advantage compared to the traditional information sourcing alternatives is the extensive amount of free data it provides across different demographic groups and geographic locations (S. Y. Kim et al., 2021). Especially compared to laborious surveys, experiments and interviews (Li et al., 2019), this is a major advantage as it boosts (voluntary) participation rate a thousandfold. Apart from the data richness, Twitter also offers the possibility to go back in time as all Tweets from publicly open accounts ² are still available today (Müller-Hansen et al., 2023).³ Although Twitter’s search engine is quite advanced, collecting tweets is greatly simplified and – for the retrieval of extensive data masses – only possible

²Users can also set their profile as private so that only followers can view them (Arigo et al., 2018)

³Presumed they haven’t been deleted by Twitter for violating rules or by the user himself/herself

via the access to the Twitter Application Programming Interface (API) (Karami et al., 2020). Not only does it allow users and researchers to gather massive data volumes, it also lets web builders embed tweets into their websites (Makice, 2009), create Twitter bots (Alothali et al., 2018) or build their own third-party Twitter applications (Reinhardt, 2009). Due to Twitter's popularity within the scientific community, even an *Academic Research API* version with a monthly tweet cap of 10 million tweets was launched in January 2021 (Tornes, 2021). All levels of API accesses were freely available until 2023 (Twitter (X) Developers, 2023).

Not only does the API allow the retrieval of large data loads, it also enables data collection at a faster rate than surveys or interviews (Ahmed et al., 2017). This is predominantly important to investigate live events (Jackoway et al., 2011) and is used for disaster management (Chen et al., 2016) and even for advanced, multilingual disaster location mapping (Sufi & Khalil, 2022). In combination with the constant supply of tweets over time, this gives the valuable possibility to analyse a phenomenon over a long timespan and to build time-series accordingly (E. L. Lai et al., 2016).

Lastly, Li et al. (2019) point out the nature of Twitter data to be composed of "crowd-sourced public opinions instead of those from a directed and targeted audience" (Li et al., 2019, p. 2). In the sense of the energy crisis, this is important as it enables policymakers to get an overview of ongoing discussions and public sentiments towards the energy crisis (Vrana et al., 2023).

2.2.3 Cautions and limitations

Despite the advantages, there are some key limitations to keep in mind when working with Twitter data. Li et al. (2019) emphasize statistical and spatial representativeness as a problem. However, this is a critical issue in any type of study, including surveys and interviews. For the purpose of this thesis, spatial representativeness shouldn't raise too many problems as German tweets are more or less only found in the examined countries. Still, people on Twitter can't fully represent *the public* because, for instance, only 10% of Germans were found to use Twitter regularly (Schmidt et al., 2022). Moreover, age and gender distribution has a clear bias towards male users between the age of 25 and 34 years (Ahlgren, 2023). Plus, data quality is impaired on all social media platforms as data is quite noisy (Jackoway et al., 2011; Li et al., 2019), containing lots of abbreviations, typos, strong language or emojis (Ahlgren, 2023). Furthermore, the opportunity to post tweets via an automated approach using the API leads to many bots influencing Twitter discussions (Lei et al., 2022). It is estimated that the amount of bot accounts is between 9 and 15% (Edry et al., 2021). Here, Novotny distinguishes between social bots, traditional spammers and fake followers (Novotny, 2019). While most (social) bots are helpful implementations (e.g. for downloading video content or making threads more easily readable), some pursue malicious goals (scamming, political agenda, fake news, propaganda, cyberbullying, harassment, etc.) (Wu et al., 2022). Bots can purposefully divert attention from controversial political decisions as found during the Syrian civil war (Alothali et al., 2018). Identifying bots is usually done via graph-based detection considering account interactions, crowdsourcing via human annotations or machine learning approaches (Alothali et al.,

2018).

Unfortunately, data downloading and automated bot identification was made massively more complicated since Elon Musk bought Twitter in 2022. He decided to introduce premium features only for payments and finally shut down the free API access for all versions in February 2023 (Twitter (X) Developers, 2023). Although there is still a free write-only access to the API, retrieving tweets costs at least 5'000\$ per month and is limited to 1'000'000 tweets (Twitter Developer Platform, 2023b). For developers and researchers, there's an enterprise package which allows to retrieve up to 200 million tweets per month at a fee of 210'000\$ (Stokel-Walker, 2023). This massively impedes research using Twitter data. Before these restrictions, though, collecting public opinion data via Twitter was less expensive, less labour-intensive and way faster than the alternative survey method (Abdar et al., 2020).

2.3 Sentiment Analysis

To make use of extensive Twitter data to examine public perceptions of energy sources, an automated process is required. This process is called *sentiment analysis*, sometimes also referred to as *opinion mining* or *subjectivity analysis* (Pang, Lee, et al., 2008). B. Liu (2012, p. 7) describes it as a "field of study that analyzes people's opinions, sentiments, evaluations, appraisals, attitudes, and emotions towards entities such as products, services, organizations, individuals, issues, events, topics, and their attributes." In the following, I will briefly give an overview of the evolution of sentiment analysis, the known methods and related work that is relevant in the scope of this thesis.

2.3.1 History and evolution

First computer-based sentiment analyses were already performed in the 1990s. Wiebe et al. (1999) collected texts from different sources and let three human annotators classify each sentence into *subjective* and *objective*. This training data was then used to build multiple machine-learning algorithms for subjectivity classification. The paper demonstrated how important a corpus of accurately annotated training data is for reliable machine-learning models. In the early years of the 2000s, sentiment analysis tasks predominantly focused on subjective text data and tried to assess the polarity (positive or negative). This was mainly done to analyse product reviews (Mäntylä et al., 2018). A pioneer paper on sentiment classification for online reviews was published by Dave et al. (2003).

With the emergence of social media, a shift of sentiment analysis towards microblogging platforms was observed and the research field grew fast (Mäntylä et al., 2018). In 2008, Pang, Lee, et al. (2008) published an almost 100 pages long paper about the incredible potential of sentiment analysis algorithms for everyday applications (marketing, customer services, etc.). Moreover, the authors provided an overview of the different techniques for sentiment analysis as well as their strengths and weaknesses. It is still the most cited paper in the research field. Four years later, another widely respected piece of work entered the research area when B. Liu (2012) released his extensive book called *Sentiment Analysis and Opinion Mining*, which is some kind of a *bible* for sentiment analysis as it carefully describes everything from the definition of a *sentiment* to all available classification approaches that were known in 2012. However, it was already in 2009 when Go et al. (2009) first used Twitter to build several classifiers which were able to differentiate between positive and negative sentiments of tweets. The Naïve Bayes algorithm was found to provide the most accurate results. One year later, Pak, Paroubek, et al. (2010) built upon that and were able to integrate a *neutral* sentiment class as well. Just like Go et al. (2009), they took tweets containing happy emoticons and such containing sad emoticons as training data for the classifier. In addition, they queried tweets from objective newspaper Twitter accounts to train the neutral sentiment class. Via POS tags and n-grams, a Naïve Bayes classifier was created to successfully differentiate between positive, negative and neutral sentiments in tweets. Until about 2015, most accurate sentiment analysis tasks were performed using traditional machine learning algorithms. Then, Tang et al.

(2015) released one of the first articles about deep learning for sentiment analysis, using different types of neural networks to classify texts into sentiments. In 2018, L. Zhang et al. (2018) published a famous survey about deep learning for sentiment analysis, explaining what deep learning is, how it works, which network types exist and which tasks it can be used for, one of them being sentiment analysis. While early sentiment analysis techniques have predominantly been used for English texts, the research field started to swap over to German scientists as well. Already in 2010, Remus et al. (2010) built a sentiment lexicon (called *SentiWS*) containing more than 32’000 German word forms which were assigned a POS tag and a respective polarity weight between -1 and 1. SentiWS has often been used for German sentiment analyses. It was also part of the research of Schmidt et al. (2022) who did a sentiment analysis on tweets of German politicians and political parties during the federal election year in 2021. Apart from the lexicon-based approach via SentiWS, the authors also tested the performance of traditional machine learning techniques like Naïve Bayes and Support Vector Machine and most recent deep learning algorithms based on transformers. The transformer-based, so-called *BERT models* outperformed traditional ML and lexicon-based approaches, especially when using a combination of large Twitter data as training datasets.

2.3.2 Transformer-based BERT models

The beginning of this new cutting-edge NLP technology, which is also the base of the currently hyped *ChatGPT*, was made by Google’s Deep Learning department *Google Brain* in 2017 via their infamous paper called *Attention Is All You Need* where the authors presented their transformer architecture (Vaswani et al., 2017). Based on this, Devlin et al. (2018) then introduced the language representation model *BERT* (Bidirectional Encoder Representations from Transformers) which became the new state-of-the-art model for tasks like question answering, translation and also sentiment analysis. Pre-trained BERT models statistically learnt the co-occurrences of words and their relationship via *Masked Language Modelling* (Salazar et al., 2019) and *Next Sentence Prediction* (Shi & Demberg, 2019) and became thus aware of the main structures of a language (Chalamkate et al., 2023; Hoang et al., 2019). Therefore, input text is converted into word embeddings that encode semantic and syntactic features (Al-Rfou et al., 2015; Mikolov et al., 2013). Thanks to the powerful self attention mechanism (Vaswani et al., 2017), pre-trained models can then be fine-tuned to be able to solve a specific task (De Greve et al., 2022; Hoang et al., 2019). One of the most popular tasks for Twitter data is sentiment analysis (Karami et al., 2020). Researchers from *Meta AI* (previously *Facebook AI*) did a replication study of Devlin et al. (2018) by evaluating the effects of different hyperparameters and the training set size (Y. Liu et al., 2019). They found Devlin et al. (2018)’s BERT model to be undertrained and optimized the model by including more training data, longer training sessions, longer sequences and a dynamically changing masking pattern which led to a better overall performance. Their model goes by the name of *RoBERTa* (Robust optimized BERT approach). In 2020, the Meta AI team took this RoBERTa model and pre-trained it on text

documents in 100 languages to create a powerful cross-lingual language model (*XLM-R*) (Conneau et al., 2019). Similarly, Barbieri et al. (2022) used 198 million Tweets to train XLM-R in more than 30 languages to make it specifically capable for NLP on microblog texts. Accordingly, they named the model *XLM-T* (for Twitter). Furthermore, they explicitly fine-tuned the model for sentiment analysis in various languages as it is the most studied task on Twitter data. Therefore, the SB-10K corpus, containing 10’000 annotated German tweets, was used which led to a slight performance improvement compared to XLM-R (Cieliebak et al., 2017). Apart from multilingual approaches, there are also a few German-only models, for example by Guhr et al. (2020) who chose the BERT architecture for a German sentiment analysis. It was trained on a large German corpus including Twitter, newspaper and online reviews data and reached an F1 score of 0.94, making it quite powerful for sentiment classification tasks.

2.3.3 Aspect-based sentiment analysis

Standard sentiment analysis traditionally examines the sentiment of the author on a document-level (e.g. a full tweet) or sentence-level. However, this is often insufficient as document or sentences may contain different talking points (called *aspects*) and different sentiments (B. Liu, 2012). To illustrate this issue, B. Liu (2012, p. 11) used the example ”The iPhone’s phone quality is good, but battery life is short”. This short phrase consists of two aspects (‘phone quality’; ‘battery life’) and two corresponding sentiments (‘good’: *positive*; ‘short’: *negative*). Standard sentiment analysis algorithms would probably label this phrase as *neutral* which is not accurate enough when the user is specifically interested in the phone quality or the battery life of the phone. To tackle this issue, a special case of sentiment analysis called *ABSA* (Aspect-based Sentiment Analysis) was brought up (B. Liu, 2012). ABSA is usually divided into three subtasks: Aspect term extraction, aspect category classification and aspect polarity classification (De Greve et al., 2022; Kersting & Geierhos, 2020). These terms are explained by the survey of W. Zhang et al. (2022): *Aspect terms* are the words that represent an aspect (in the previous examples: ‘phone quality’, ‘battery life’). These aspect terms can be assigned to an *aspect category* (e.g. ‘phone quality’, ‘battery’). Further, there are *opinion terms* signalling corresponding sentiments (‘good’, ‘short’) which further define the *polarity classification* (‘negative’, ‘neutral’ or ‘positive’).

While this specific research field was dominated by English language for a long time, the *GermanEval17* dataset served as one of the first (and still quite few) tasks that focused on aspect-based sentiment analysis on German texts (Wojatzki et al., 2017). The dataset consists of more than 26’000 reviews about the German railway transport system (*Deutsche Bahn*) gathered from Twitter, Facebook and Q&A websites in 2015 and 2016 (Wojatzki et al., 2017). All documents were elaborately annotated for four subtasks: Relevance classification, document-level sentiment analysis (Is the overall sentiment negative, neutral or positive?), aspect-based sentiment analysis (assign a sentiment polarity to all aspects of predefined categories, e.g. punctuality of trains) and opinion target extraction (finding linguistic features defining the sentiment) (Aßenmacher et al., 2021).

While most scholars focused on document-level sentiment analysis using traditional ML techniques or neural networks (Attia et al., 2018; Biesialska et al., 2020; Cieliebak et al., 2017; Guhr et al., 2020), only M. Schmitt et al. (2018) dedicated himself to the aspect-based subtask using CNNs that achieved a micro-averaged F1 score of 0.55 on the test set. It was four years later - after the introduction of BERT models by Devlin et al. (2018) - when Aßenmacher et al. (2021) revisited the GermEval17 dataset making use of pre-trained German BERT models to test their performance on this challenging subtask. They managed to achieve an F1 score of 0.79 that demonstrated the power of transformer-based models. One year later, De Greve et al. (2022) fine-tuned a pre-trained BERT model for ABSA on German tweets about books presented at a literature prize and were able to extract the aspect category (main motif being discussed in the text) and the corresponding polarity classification. Achieved micro average F1 scores for these two subtasks were at 0.8 and 0.72 respectively. Another German ABSA was carried out by Kersting and Geierhos (2020) who worked with physician customer ratings from Austria, Germany and Switzerland. However, from the three main subtasks of ABSA, only an aspect term extraction and the corresponding aspect category classification were conducted.

2.3.4 German sentiment analysis regarding energy perception

There are various non-German sentiment analysis studies about specific energy sources like nuclear energy (Jeong et al., 2021; Z. Liu & Na, 2018), solar energy (S. Y. Kim et al., 2021; Nuortimo et al., 2018), wind power (Vågerö et al., 2023), hydropower (Jiang et al., 2016; Yin & Fan, 2023) or renewable energy sources overall (Abdar et al., 2020; Ibar-Alonso et al., 2022). However, except for Vågerö et al. (2023) (2006 until 2022) and Yin and Fan (2023) (1971 until 2020), sentiment was analysed over a short period of time which could not provide useful insights into how the public perception changed.

Albeit being quite rare, there are also some German scholars who recently performed sentiment analyses for energy perception monitoring. For example, there is a recent study covering the Twitter discussion of EU citizens about the energy crisis in 2022 by Vrana et al. (2023). Simple sentiment analysis was performed using a multilingual sentiment lexicon in September 2022. Moreover, the authors conducted a content analysis via frequent word pairs to compare sentiments with central aspects of the discussions. Apart from the narrow time frame, I'd argue that this is more of an *emotion analysis* as emotions like *fear* or *trust* were monitored instead of classic polarities proposed by B. Liu (2012). The author managed to monitor how people felt about the energy crisis in autumn 2022 while no information was gained about the sentiments towards specific energy sources (Vrana et al., 2023). The latter was done by Dehler-Holland et al. (2022) who combined a structural topic model with a lexicon-based sentiment approach to monitor sentiment variations towards wind power using German newspaper articles that were published between 2009 and 2018. Via this technique, they could classify articles into topics related to wind power and weighted the sentiments of these

topics on a monthly basis. The authors found that the support of wind power got increasingly challenging over the years (Dehler-Holland et al., 2022). Another insightful study was conducted by Müller-Hansen et al. (2022) between 2017 and 2020 who examined the perception of coal energy in Germany. They also used a simple lexicon-based approach (SentiWS by Remus et al. (2010)) which assigned a polarity score to each word. For the sentiment of the tweet, the average word polarity was chosen. Eventually, the authors found a decreasing public perception trend regarding coal energy, accompanied by an increasing polarization (Müller-Hansen et al., 2022). There's also a recent sentiment analysis study by Xu et al. (2022) examining public sentiment of nuclear energy in German-speaking countries. However, apart from the abstract, I could not access the full text of this paper.

A research gap

Conclusively, there is no German sentiment analysis research about the public perception towards different energy sources over the timespan of several years which would enable a *bigger picture* of the topic. Moreover, the existing German studies are limited to lexicon-based sentiment approaches which are methodologically limited due to the complexity of the German language and the noisy social media data. Hence, broadening the perspective methodologically by using state-of-the-art models in combination with an aspect-based sentiment analysis model, thematically by covering different prominent energy sources, and temporally by expanding sentiment analysis over several years is necessary to get more insight into the public energy perception.

2.4 Content Analysis

Once an insight into the temporally shifting sentiments about different energy sources is guaranteed, it will further be necessary to inspect possible underlying causes and prominent talking points. Due to the interaction setting, social media data is often clustered around ongoing events or themes which are of interest to be revealed for academic research fields (E. L. Lai et al., 2016). Such a *content analysis* could generally be summarised as a technique of transforming unstructured text data into a format that allows the systematic and quantitative description of its substance (Marine-Roig, 2022). In its early days of research, content of Twitter data has been analysed using manual coding regarding Twitter accounts of television channels (Greer & Ferguson, 2011), sport athletes (Hambrick et al., 2010) or universities (Linville et al., 2012) before moving on to more sophisticated methods which will be mentioned as follows.

2.4.1 Term frequency methods

Already in 2013, Benhardus and Kalita (2013) were interested in uncovering trending topics on Twitter. The authors did this by analysing the frequency of terms appearing in tweets over a certain time period and further used unigrams and bigrams as well as normalization methods to monitor the topics of online discussions. Such content analyses were also carried out to uncover events that could be responsible for observed sentiment variations. For Twitter data, this was done by Schmidt et al. (2022), who investigated the sentiment of political parties during the federal election 2021 and performed a term frequency analysis which is the simplest kind of BOW analysis according to Purves et al. (2022). By grouping the Twitter data, Schmidt et al. (2022) monitored words that usually appeared in negative and positive tweets. As a visualisation, they chose word clouds. A derivative of term frequency is the TF-IDF (term frequency - inverse document frequency) method which compares the abundance of a certain term in a document with its frequency in the whole corpus (Purves et al., 2022). By doing so, terms frequently occurring in just a few documents are enhanced, while 'standard' terms that are part of many documents within the corpus but appear less frequently in these documents get assigned a smaller weight (Hiemstra, 2000). TF-IDF can further be used to transform text data into vector representations that serve as inputs for machine learning algorithms (Ahmed et al., 2017).

Instead of single words, Vrana et al. (2023) monitored frequently used word pairs to get a broader insight into the ongoing debates regarding the European energy crisis. To focus on important terms, stop words were removed beforehand and a threshold for the occurrence of word pairs was defined, removing pairs that occurred less than ten times. Each word was then represented as a node and edges between these nodes were added if the words co-occurred. This resulted in a semantic network which was further cleaned based on betweenness centrality (Vrana et al., 2023). Building on the word pair approach, the *distance* between words can further be taken into account to make typical patterns visible (Purves et al., 2022). This is predominantly a good option for longer documents. This method can further be expanded to "meaningful combination of words,

termed collocates” (Purves et al., 2022, p. 60). Meaningful are such combination if they statistically co-occur more frequently than expected by random choice (El-Kishky et al., 2014).

2.4.2 Topic modelling

An advanced method to analyse large text documents are topic models which also attempt to make latent themes of a given text corpus visible (E. L. Lai et al., 2016). One widely used approach of topic models was introduced by Blei et al. (2003) who came up with *Latent Dirichlet Allocation (LDA)* that uses the statistical distribution of words in a corpus to build groups (Jelodar et al., 2019). Among other scholars who made use of LDA for Twitter content analysis (e.g. (Xue et al., 2020)), Karami et al. (2020) used this technique for an extensive systematic review about research with Twitter data to monitor common topics discussed. Once each document (in the case of Karami et al. (2020) abstracts of research papers) was converted into a BOW representation, LDA examines the occurrences of words in the documents of the whole corpus. Based on this word distributions, topics are identified and are represented by word probabilities (Jelodar et al., 2019). As it can efficiently group unstructured text data, LDA is commonly used to gain a better understanding of tweets (Negara et al., 2019; Xue et al., 2020; S. Yang & Zhang, 2018). In their review, Karami et al. (2020) let the LDA model identify 40 different topics. For instance, it was found that words like 'tourism', 'species' and 'human' frequently co-occurred which allowed to form a topic 'Nature/Tourism' (Karami et al., 2020). Moreover, LDA is a good option to explore a corpus of documents and their dominant themes (Purves et al., 2022).

From LDAs, more advanced statistical models called *Structural Topic Models (STM)* were derived to "make inference about social and political processes that drive discourse and content" as Roberts et al. (2016, p. 1) put it. Dehler-Holland et al. (2022) decided to use a structural topic model (STM) as an unsupervised method to split their newspaper articles about wind power into meaningful categories. The functionality of an STM is very similar to the one of LDAs as it also assumes that a topic consists of the frequent usage of similar words (Dehler-Holland et al., 2022). The main innovation of STMs compared to LDAs is the implementation of covariates (other features apart from pure BOW-represented texts, e.g. time or publisher) to improve the classification (Roberts et al., 2016). As preprocessing is more elaborate, Dehler-Holland et al. (2022) applied POS-tagging, lemmatization and stop word removal on the data and finally transformed the corpus into a document-term matrix. For their investigation about legitimacy towards wind power, the authors relied on the covariate 'time' since they view legitimacy as time-dependent, and let the STM create 44 different topics which were then undertaken a sentiment analysis (Dehler-Holland et al., 2022).

2.4.3 Phrase extraction

Although all these models are based on a BOW representation of text documents and, thus, ignore the grammatical order of words, they "can be surprisingly effective when we are dealing with large

corpora.” (Purves et al., 2022, p. 57). Leaving the BOW-approach, M. Yang et al. (2018) went a step further believing that words do not contain enough information to gain sufficient insights into social media discussions and thus used a phrase-mining algorithm to extract frequent word sequences. This data-driven algorithm created by El-Kishky et al. (2014) browses the corpus for frequent contiguous patterns and then statistically merges them into meaningful sentences. Several advanced phrase mining methods have been proposed in the recent years. One of them was developed by (Tripathy et al., 2021) in 2021 and implements a POS tagger which assigns a grammatical feature to each word (Purves et al., 2022). Here, especially nouns are tagged due to their significance for phrases (Tripathy et al., 2021). Based on the POS tags, extractive summarization is applied to retrieve important information in form of frequent bigrams, followed by the comparison of these bigrams with the original phrases (similarity measure) (Tripathy et al., 2021). Schopf et al. (2022) combined this POS-tag approach with the power of BERTs to extract sequences that best describe the full document. Here, the similarity of the extracted phrases to the full document is based on the vector embeddings coming from BERT models. As these BERT models also exist for German language, a phrase extraction for German documents is possible as well (Schopf et al., 2022).

2.5 Geospatial separation

According to the aims and research questions of this thesis, a distinction between the three German-speaking countries (Germany, Austria and Switzerland) is made to compare the sentiments on a country scale. Hence, a geographic component comes into play and is, thus, crucial to be extracted from the dataset. In other studies, this valuable location information was retrieved for traffic, disaster or crime management and disease surveillance (X. Hu et al., 2022).

On Twitter, geographic information can be shared via three different methods. First of all, users can enable Twitter to access their GPS data to add the precise location to a tweet (Twitter, 2023). This location is then shown in the tweet itself, hence, such tweets are typically called *georeferenced tweets* (Hahmann et al., 2014, p. 4). Since this feature is turned off by default, only about between 0.4% and 1% of Twitter users add the location to their tweets (Carley et al., 2016; Hecht et al., 2011; Morstatter et al., 2013; Ryoo & Moon, 2014).

The second option to add location information is the Twitter profile itself (Zheng et al., 2018). However, this is a unstructured field meaning that neither GPS data has to be provided, nor does the user have to select a location from a database. Instead, any possible combination of characters can be manually typed in. Not surprisingly, Hecht et al. (2011) found that 34% of Twitter users provided a fictional or sarcastic place names instead of a real locations. Those users who typed in a real location were mostly referring to a city (Hecht et al., 2011). While comparing the profile location with the GPS coordinates attached to tweets, B.-S. Lee (2012) found that close to 50% of all users post most of their tweets from their profile location. An even lower amount (34%) was found by Leetaru et al. (2013), indicating the unreliability of this place declaration for fine-grained geolocation purposes (Mahmud et al., 2012).

The third and most frequently used option to incorporate geospatial information is via the tweet text itself (Zheng et al., 2018). As part of a flood detection study, Arthur et al. (2018) found 33% of their retrieved tweets to mention a place name (so-called *toponym* (Gritta et al., 2020)). Especially for disaster management purposes (Grace, 2021), a tweet gets instantly more valuable if a location is provided (Munro, 2011). When reporting local events, Twitter users prefer to include place names into the tweet's text instead of adding GPS tags (Grace, 2021). This makes sense as not all people who tweet about a place *are* at that place by the time of tweeting. This issue was illuminated by a study of Hahmann et al. (2014) who found a low correlation between the exact location tweets were sent from and the toponyms mentioned in those tweets. Since it's the goal of the thesis to compare energy source related sentiments of users from Germany, Austria and Switzerland, the home location is of interest. Thus, compared to most other studies, no fine-grained location extraction is necessary.

2.5.1 Geoparsing

The process of identifying toponyms and translating them into geographic objects is usually called *geoparsing* and is divided into two steps: *Geotagging* and *Geocoding* (Gritta et al., 2020). While

geotagging is also known as *toponym recognition*, geocoding is often referred to as *toponym resolution*. The first step identifies a toponym in an unstructured text document (such as a tweet), before the second one disambiguates the toponym and links geographical coordinates to it (De Bruijn et al., 2017; Gritta et al., 2018, 2020; Lieberman et al., 2010). In this sense, geotagging is a special case of *Named Entity Recognition (NER)* (Gritta et al., 2020) which generally identifies proper nouns as people, organisations or – like in this case – places (Purves et al., 2022). Although terminologies can be very confusing as they vary from author to author⁴, I will refer to geoparsing as a combination of toponym recognition (= geotagging) and toponym resolution (= geocoding) in this thesis (Figure 1) (as proposed by Gritta et al. (2018), X. Hu et al. (2022), Y. Hu and Adams (2020), Wang and Hu (2019), and Wang et al. (2020)). Due to the noisy Twitter data quality, language diversity, ambiguous place names, metonyms and limited context information, geoparsing is a challenge (Gritta et al., 2018). Different toponym recognition and resolution approaches exist to tackle this challenge, varying in complexity and accuracy.




Figure 1: The process of Geoparsing.

2.5.2 Toponym recognition

Usually, four techniques are distinguished (rule-based, gazetteer-based, statistical learning-based and hybrid approaches of aforementioned techniques) when it comes to extracting place names from unstructured texts (X. Hu et al., 2022; Zhou et al., 2023). While simple rule-based approaches focus on typical suffixes (in German e.g. ‘-berg’), prepositions or typical POS-combinations to identify locations (X. Hu et al., 2022), gazetteer-based approaches compare words of the text document with dictionaries that store location names and corresponding geospatial information (*gazetteers*) (Sagcan & Karagoz, 2015). Some gazetteers also contain additional facts like population size, alternative names or administrative levels (Y. Hu & Adams, 2020). Still, these approaches are quite limited, especially due to place name ambiguity, and hugely depend on the defined rules or the volume of the gazetteer (De Bruijn et al., 2017).

Gazetteer-upgrades

Hence, gazetteer approaches were upgraded by including other techniques and heuristics to improve their performance. Particularly optimized for tweets, Paradesi (2011) inspected the words before

⁴There are authors who use geoparsing as a synonym for geotagging, leaving out the geocoding part (Gelernter & Balaji, 2013; McCurley, 2001), such that equate geotagging with geocoding (Ghahremanlou et al., 2015) and such treating geotagging as consisting of toponym recognition and toponym resolution (De Bruijn et al., 2017)

a term found in gazetteers to evaluate whether it is actually a toponym or not as prepositions like ‘near’ or ‘in’ are often followed by place names. Mahmud et al. (2012) used the *USGS* gazetteer with a combination of classifiers that further took non-toponym words and hashtags of the tweet into account. They additionally incorporated heuristic classifiers such as number of times certain toponyms are mentioned and tweeting behaviour (tweet volume per time unit). A Naïve Bayes Multinomial method was used to train the classifiers which led to a performance improvement when predicting the home location of users (Mahmud et al., 2012).

Statistical learning approaches

In comparison to gazetteers, statistical learning-based methods like NER use NLP to analyse grammatical patterns and structures of a text to classify words into categories of which one usually is location (Al-Rfou et al., 2015). This especially enables an appropriate distinction between place names and people (Amitay et al., 2004). Most of these algorithms were trained using traditional machine learning approaches like Random Forest (RF) or Long Short-term Memory (LSTM) (X. Hu et al., 2022). A famous NER is the *Stanford NER* which was pre-trained on a large English corpus (CoNLL) in 2003 and uses conditional random fields (CRF) to recognise words that represent a location (Finkel et al., 2005). Stanford NER served as the geotagging algorithm in multiple studies, e.g. by DeLozier et al. (2015) or Karimzadeh et al. (2013). The latter came up with their *GeoTxt* geoparser which is specifically optimized for unstructured micro-text such as tweets. However, apart from the linguistic limitation, the Stanford NER was found to perform badly when it comes to informal, user-generated content (P. Liu et al., 2022). In 2015, Al-Rfou et al. (2015) built their own NER system using word embeddings and *Wikipedia* articles which could outperform existing algorithms like OpenNLP and NLTK and reach a good performance especially for tagging persons (Al-Rfou et al., 2015). In 2020, the Stanford NLP group published *Stanza*, an open-source python library for toponym recognition (and various other NLP tasks) which supports 66 languages including German (Qi et al., 2020). It was pre-trained on 112 datasets and operates via BiLSTM and a CRF. A direct competitor is *SpaCy* which is also a free open-source NLP library for python (spaCy.io, n.d.). It can be used for various NLP tasks and supports multiple languages as well, one of them being German. It focuses on a simple implementation as aims to be as user-friendly as possible. While SpaCy wasn’t that good at recognizing persons and organisations, it reached an F1-score of 0.74 regarding toponym recognition (X. Schmitt et al., 2019).

Deep learning approaches

More complicated but also more powerful deep learning algorithms have been built by Wang et al., 2020 who came up with the so-called *NeuroTPR* and by Zhou et al., 2023 who recently introduced their cutting-edge *TopoBERT*. NeuroTPR is based on a bidirectional recurrent neural network specifically trained for noisy microblog texts. Its limitations were addressed by TopoBERT which uses a onedimensional CNN in combination with the capabilities of a pre-trained BERT language

model. It was fine-tuned using three large datasets and reaches a F1-score of 0.865 for toponym recognition, which is better than the performance of NeuroTPR (0.728) tested on the same dataset (Zhou et al., 2023). Such BERT-based systems are particularly powerful when it comes to larger context observations (X. Hu et al., 2022) and social media texts that often consist of informal language, abbreviations, inconsistent capitalization or misspellings (Wang et al., 2020). The large majority of these models are limited for this thesis as they are only available for English, French or Chinese text data (Chalamkate et al., 2023). Apart from multilingual SpaCy (spaCy.io, n.d.) and Stanza (Qi et al., 2020), there are barely any toponym recognition models specifically designed for the German language.

2.5.3 Toponym resolution

Depending on the method, toponym recognition and resolution are closely related. While gazetteers already contain stored coordinates and network-based methods like the oft-cited approach by Compton et al. (2014) directly infer user locations from GPS locations of Twitter friends, NLP toponym recognition methods, that solely focus on predicting whether a word is a place name or not, demand an additional toponym resolution part to retrieve corresponding coordinates. Most statistical-learning based NLP approaches still used common gazetteers as a basis. As for the gazetteers, Acheson and Purves (2021) found that the *Google Geocoding API* performs a lot better at toponym resolution than the famous *GeoNames*⁵, at least for locations mentioned in scientific articles. For toponyms not present in the chosen gazetteers or such with multiple entries, scholars additionally developed geocoding and disambiguation methods. As Weissenbacher et al. (2015) summarised, there are mainly two methods to do so: One inspecting the context words in the document and one using heuristics like distance or population count. Based on the literature, I’d suggest that there is a third method making use of tweet metadata.

Context, heuristics and metadata

For GeoTxt, Karimzadeh et al. (2013) applied ranking rules using *Levenshtein Distance* – a quite old similarity measure introduced by Levenshtein et al. (1966) – to compare the toponyms found in the text with the entries in GeoNames and population count majority to select the most suitable entry. Moreover, other toponyms in context are taken into account to enhance disambiguation. A similar approach was chosen by Weissenbacher et al. (2015) who also considered a population heuristic to disambiguate toponyms but further included the distance to other found (and matched) toponyms in the same document. Paradesi (2011) compared the locations of the found toponyms with the user location found in the profile to disambiguate frequently occurring place names. In a study about geocoding Thai tweets, Chalamkate et al. (2023) used the Google Geocoding API as it provides comprehensive and accurate information about Thai names. The authors developed an approach to geocode toponyms from Twitter data that could not be matched with the Geocoding

⁵<https://www.geonames.org/>

API. They tested different clustering methods to estimate the location of unmatched toponyms and found topology words (expressions found in the tweet containing an unmatched toponym) to be the most accurate predictors of the geolocation (Chalamkate et al., 2023). Not solely focussing on topology words, DeLozier et al. (2015) came up with *TopoCluster* which is a geoparsing approach based on the assumption that a set of words can predict locations. Then, they trained a model based on a Wikipedia dataset to associate a cluster of words with corresponding coordinates. Similarly, Hahmann et al. (2014) looked at toponyms and other words that are characteristic for a certain region (e.g. 'beach' for a coastal place).

De Bruijn et al. (2017) tackled the problem of ambiguity and multiple toponym mentions (within the same tweet) by taking metadata of tweet groups referring to the same toponym into account. Particularly, the user's time zone, the residence location from the profile and tweets with GPS-enabled location were considered for the disambiguation process. The authors inspired Vågerö et al. (2023) who focussed the user defined location field specified in the Twitter profile and compared this with Wikipedia data to skip fictional place names. Due to De Bruijn et al. (2017)'s grouping approach, tweets without GPS locations could be geotagged on multiple geospatial scales. The performance could be improved compared to Grover et al. (2010) who worked with two gazetteers and also considered the context words but didn't group the tweets. Time zone metadata was also used by Arthur et al. (2018) for location inference. The authors closely followed the approach proposed by Schulz et al. (2013) who also worked with GeoNames and a combination of metadata that hints location (they call these information *spatial indicators*). Apart from time zone, the tweet text itself, the location field in the profile, other user account information or shared links can indicate a location or at least narrow the selection of ambiguous toponyms or possible regions (Schulz et al., 2013). Compared to many other studies, Arthur et al. (2018) did not locate tweets or users but focussed on retrieving the location of flood events mentioned in tweets. Similarly, Sufi and Khalil (2022) built an algorithm that scans social media posts for disaster keywords and extracts the mentioned location via a NER system to retrieve the location of the disaster. Hence, the aims of Arthur et al. (2018) and Sufi and Khalil (2022) are similar to the one of this thesis. However, while flood event locations should be geocoded as accurate as possible, retrieved tweets document should only be assigned to one of the three German-speaking countries here.

3 Methods

As described in the previous chapters, there exist many different approaches to perform sentiment, content and geospatial analyses using Twitter data. With regard to the reproducibility, critical reflection and assessment of the results, knowledge of the exact methodology will be provided in this chapter. Figure 2 gives a broad overview of the overall methodological procedure that has been chosen to address the research questions. Following the thesis extent justifying the thematical, spatial and temporal dimension of this work, the individual methods will be described in detail in the following subsections.

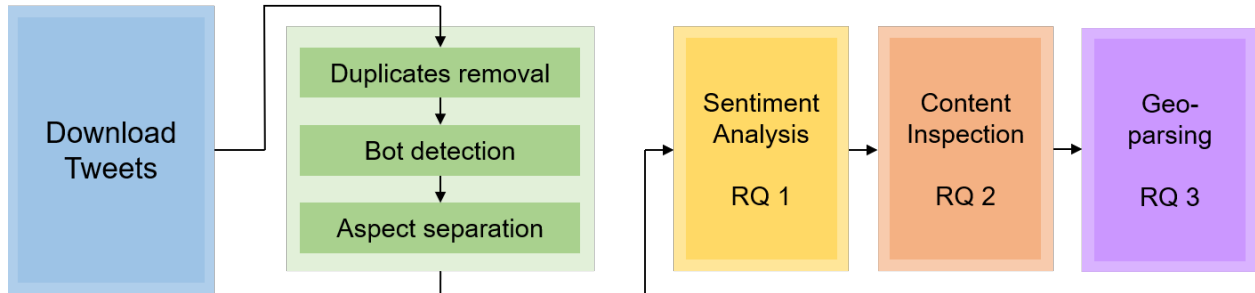


Figure 2: The methodic steps conducted in this thesis.

3.1 Thesis extent

In context of the *energy crisis* (see 2.1.2), Swiss newspapers, political debates and private discussions are frequently centered around energy supply and possible solutions regarding energy sources. Due to the lack of existing work about sentiment towards energy sources in German-speaking countries, the necessity of public support and the different energy policy approaches of the surrounding countries, this thesis doesn't only focus on the case of Switzerland but further considers its neighbour countries Austria and Germany.

In accordance to the selected countries, the dominant energy sources used in Germany, Austria and Switzerland have been chosen (IEA, 2020). Although biofuels and waste also play a certain role for energy supply in all three countries, this category has been left out as I perceived it as less discussed and rather uncontroversial compared to the other energy types. Moreover, I decided to summarise gas and oil into the same energy source class as they are closely related and were found to be mentioned together quite often during the visual inspection of tweets. Conclusively, six energy source categories have been set up: Nuclear, coal, solar, wind, water, gas/oil.

An advantage of using Twitter is the possibility to retrieve historic data. Although the first tweet was posted in 2006 already (Dean, 2022), I decided to choose *2007-01-01 00:00:00* as a starting date due to the small tweet volume in the year before. Since my work for this thesis started in January 2023, the end date was set to *2022-12-31 23:59:59*. Due to the sparse tweet availability in the platform's early years, the first tweet that was retrieved was posted on *2007-04-26 10:08:52*, while the last one was from *2022-12-31 23:56:06*.




Figure 3: The temporal distribution of tweets per energy source.

3.2 Data download and software

Access to the Twitter API made data retrieval straightforward, at least before the free API services were closed in February 2023.

Fortunately, since I already decided to use Twitter for my Master's thesis in December 2022, I could successfully apply for the *Academic Research* access and gather the data I needed for my work before the payment era started. To be granted access, a link to my name found on the GIUZ' departement website and a description of how the data will be used had to be provided. Between February 4th and February 8th, a total of 5'136'480 tweets were downloaded (see Figure 3). Apart from the temporal time span, multiple parameters could be specified (see Twitter Developer Platform (2023a)). For the tweets, all available optional fields such as geo-information, public metrics and entities have been selected. However, no optional user fields were included as I missed this option although they would have been useful for preprocessing. Table 1 lists all retrieved fields.

tweet fields
'created_at', 'id', 'entities', 'possibly_sensitive', 'lang', 'conversation_id', 'text', 'edit_controls', 'reply_settings', 'in_reply_to_user_id', 'public_metrics', 'edit_history_tweet_ids', 'referenced_tweets', 'author_id', 'geo', 'withheld'

Table 1: The attributes downloaded by the Twitter API.

The actual tweet retrieval was then conducted using suitable keyword queries. For each of the energy sources, a respective query was built as shown by [Table 2](#). Every tweet posted between January 2007 and December 2022 which contained at least one of those words was downloaded. Compared to Müller-Hansen et al. (2022), keywords like 'Kohleausstieg' or 'kohlefrei' were not included as they already contain a sentiment bias. Since the Twitter API is not case-sensitive, also tweets containing the keywords as lower-cased were downloaded. However, only words exactly matching the keywords provided were retrieved. Hence, the plural forms were specifically included as well. Moreover, tweets containing the energy source (e.g. 'Kohle') and one of the words 'Energie' or 'Strom' were also selected and downloaded. Accordingly, the sole occurrence of an energy source word ('Atom', 'Kohle', 'Sonne', 'Wind', 'Wasser', 'Gas', 'Öl') was not meaningful enough for a tweet to be downloaded. This decision was inspired by Müller-Hansen et al. (2023) and taken in favor of a better precision since these words can have multiple meanings in the German language and don't necessarily refer to an energy source. However, this reduced the recall as it's assumed that some users tweeted about energy without making use of these two additional words. Similar issues were evoked by the decision to implement 'PV-Anlage' instead of only 'PV' due to its shortness. Moreover, microreading later revealed that some people also used the abbreviations 'WEA' for wind power plants and 'WKA' for hydropower plants. Both terms were not included in the search queries, hence reducing the recall. On the other hand, the precision didn't reach a maximum either as some keywords referred to non-energy topics (e.g. people tweeting about a famous song called *Kernkraft 400*). Moreover, especially the usual abbreviations for nuclear energy ('AKW' and 'KKW') could easily refer to other non-German words. However, due to the character limit of 280 characters and the popularity of these abbreviations, I decided to include them in order to expand the recall.

As described before, the spatial focus of the thesis lies on Germany, Austria and Switzerland. However, the only download criteria were the defined keywords. Hence, tweets posted by users from the Principality of Liechtenstein or German-speaking users abroad were also downloaded. Even though a language feature parameter ('lang: de') was involved into the search query, this didn't seem to perform perfectly. Thus, also tweets in other languages were retrieved. Since Twitter automatically assigns a language label to the metadata of each tweet, I performed a simple language sensitivity analysis and found 99.86% of all tweets to be in German, followed by 0.09% in Dutch and 0.02% in English.

The API also allowed to exclude retweets if wanted so. Due to the retrieval limit of 10'000'000 tweets per month and the time pressure due to the announced API shutdown, retweets were excluded. According to Sharma and Gupta (2022) and Hutchinson (2021), about 50% of all tweets are retweeted. Since especially tweets of famous people can have thousands of retweets, tweet load would have been too big. Moreover, the 'public metric' attribute was included during the downloading process. Thus, it's still possible to view how often a certain tweet was retweeted. Consequently, not all existing tweets containing a sentiment about a chosen energy source were derived while some undesired posts were included, limiting the validity of the results.

For the downloading process, Python 3.10 in combination with the *tweepy* package (Roesslein, 2018) was used. Due to the rate limit of maximum 500 tweets per request, the data was downloaded in batches and then merged afterwards. For more powerful operations, an *Ubuntu* instance provided by the University of Zurich was employed.

Energy source	Keyword query
nuclear	'(Atomkraft OR Atomkraftwerk OR Kernkraft OR Kernkraftwerk OR Atomenergie OR Kernenergie OR Nuklearenergie OR Brennstäbe OR AKW OR KKW OR AKWs OR KKWs) lang:de -is:retweet'
coal	'(Kohlekraft OR Kohlekraftwerk OR Kohlekraftwerke OR Kohlestrom OR Kohleenergie OR Kohlenstrom OR Kohlenenergie OR Braunkohle OR Steinkohle OR (Kohle (Energie OR Strom))) lang:de -is:retweet'
wind	'(Windenergie OR Windkraft OR Windpark OR Windkraftwerk OR Windkraftwerke OR Windrad OR Windräder OR Windkraftanlage OR Windenergieanlage OR Windenergieanlagen OR Windkraftanlagen OR (Wind (Strom OR Energie))) lang:de -is:retweet'
solar	'(Solarenergie OR Sonnenenergie OR Solarkraft OR Sonnenkraft OR Photovoltaik OR PV-Anlage OR Photovoltaikanlage OR Photovoltaikanlagen OR Sonnenkollektoren OR Solarzelle OR Solarzellen OR Solarmodul OR Solarstrom OR (Sonne (Energie OR Strom))) lang: de -is:retweet'
water	'(Wasserkraft OR Wasserkraftwerk OR Wasserkraftwerke OR Laufwasserkraftwerk OR Flusskraftwerk OR Laufkraftwerk OR Laufwasserkraftwerke OR Flusskraftwerke OR Laufkraftwerke OR Pumpspeicherkraftwerk OR Pumpspeicherkraftwerke OR Speicherkraftwerk OR (Stausee OR Staumauer) (Strom OR Energie) OR Speicherkraftwerke) lang:de -is:retweet'
gas	'(Erdgas OR Erdöl OR Erdoel OR Heizöl OR Mineralöl OR (Gas OR Öl) (Strom OR Energie OR heizen) OR Erdölförderung) lang:de -is:retweet'

Table 2: Search queries defined to retrieve tweets via the API.

3.3 Preprocessing

In order to use the tweets for the main tasks corresponding to the research questions, namely sentiment analysis, content inspection and geoparsing (Figure 2), they have to be cleaned to prevent unwanted noise, bias or distortion. Especially due to the API which allows the automated posting of tweets, a lot of spams and politically motivated content is generated (Edry et al., 2021). This needs to be tidied up as efficiently as possible since it’s the goal of the thesis to monitor the sentiment of human users. Moreover, even if not sent via an automated approach, Twitter users can distort a discussion by copy pasting their own or other user’s tweets. As the sentiment towards an energy source is averaged over all tweets, spams can misrepresent the discussion (Balet et al., 2023; Dehler-Holland et al., 2022) and thus massively warp the sentiment. To avoid this unequal weighting of opinions, duplicate tweets and bot accounts are excluded for the further analysis. In addition, an aspect separation classifier was set up to reduce errors arising from the sentiment analysis task. These steps and their limitations will be described in detail in the following subsections.

3.3.1 Duplicates removal

To remove duplicate tweets (spams), different approaches have been adopted in the literature. While Feizollah et al. (2019) used *MD5 hashes* by Rivest (1992) to detect *real duplicates* (called *strict duplicates* by Nauman and Herschel (2022)), Tao et al. (2013) came up with some *near duplicate* (called *fuzzy duplicates* by Nauman and Herschel (2022)) detection methods. Here, the contents of two or more documents don’t have to be congruent to be considered spams. Instead, the authors labelled posts as near duplicates based on a certain degree of similarity. This similarity considered syntax, semantics and context. If documents reached a certain threshold of similarity, they were considered near duplicates and could then be removed accordingly.

Sedhai and Sun (2017) introduced a semi-supervised approach to detect spam tweets which takes blacklisted URLs, spam words and pre-labelled tweets into account. Hence, they were also guided by some rules and similarity measures just like Viswanathan et al. (2019) who referred to a duplicates as a pair of sentences conveying the same meaning and tested six machine learning algorithms to identify such similar sentence pairs. Before performing sentiment analysis on newspaper articles, Dehler-Holland et al. (2022) also removed duplicates via the Levenshtein distance in order to reduce biases.

For such newspaper articles, which are longer than tweets, it makes sense to consider the similarity as a criterion for duplicates as the possibility of two or more articles to be copy-pasted is very small. Also, near duplicate detection seems plausible for the purpose of finding similar threads in a question answering forum as done by Viswanathan et al. (2019). However, for this thesis, the data retrieval based on keywords already led to a rather high degree of similarities within a corpus of a certain energy source. Although similarity measures would allow the identification of spam tweets that vary only slightly which would increase the precision, the recall would suffer from this approach due to textual resemblance of the tweets and their short lengths. Due to the subsequent

bot detection step, it was expected to catch users that operate with slightly varying spam tweets. So, the recall was deemed more important than the duplicate precision in this first preprocessing step. Hence, near duplicate detection was considered as unsuitable in this case.

Consequently, I decided to conduct a strict approach and removed the tweets with the exact same wording. However, as it was found that some users posted the same tweet with just a different user tagged or different URLs mentioned within the tweet, tags (e.g. @user1) and URLs have been removed beforehand to avoid the inclusion of such slightly variable spam tweets. Hence, this approach is a bit less strict than the one of Feizollah et al. (2019) as their MD5 hashing treats every single character, capitalization or punctuation variation as a new document. The first idea was to remove only duplicate tweets sent by the same user. As there are reports of extensive spam campaigns involving a network of different users (Chu, Widjaja, & Wang, 2012), this idea was discarded and duplicates were ruled out ignoring the origin author. Still, I decided to implement each spammed tweet exactly once as it is still an opinion which should be considered but should not be overweighted. So, for each duplicate tweet, the most recent one was kept while the others were dropped. Thus, in this thesis, it's referred to a *duplicate* or a *spam* when the main tweet text exists more than once in the exact same wording, while additional user tags and included URLs may still differ.

This approach has strenghts and weaknesses. While it's rather unlikely that two tweets with 280 characters have the exact same wording solely by chance, this probability increases with decreasing tweet length. For example, short slogans like "Atomkraft nein!" were definitely posted by many different users independently without the intention of spamming. Moreover, many Twitter users frequently share newspaper articles via the share function of the newspaper site, resulting in duplicate tweets although users operated independently but just considered the same article as interesting. Although the majority of newspaper articles is of neutral sentiment, all those tweets were removed though they should have accounted for the sentiment calculation. This negatively affected the recall. On the other hand, one single character variation (with exception of user tags and URLs) was enough to bypass the duplicate detection algorithm, decreasing the precision of the spam detection methodology. Nevertheless, after weighting the drawbacks of the fuzzy and strict approach, the latter was assumed to be more suitable, mainly due to the textual similarities of the tweets and the subsequent cleaning steps.

After the duplicates were removed, a total of 4'245'701 tweets were left. Large amounts of duplicates were found for solar energy, wind energy and gas/oil. By analysing the Twitter accounts responsible for the duplicates, the most extensively spamming user accounted for not less than 92'178 duplicate tweets (over all six energy categories). This is almost 2% of the initially downloaded tweet volume.

3.3.2 Bot detection

It's probable that the aforementioned Twitter user who accounted for such a large portion of all spams used an automated approach to post tweets, making him/her a bot. Literature estimates

that between 9% and 15% of all Twitter users are bots (Ahlgren, 2023; Edry et al., 2021). They can be motivated by an ideological agenda or use Twitter for political propaganda, among other reasons. Hence, a single human can purposely influence an online discussion. For the purpose of this thesis, this is not desirable. For instance, a pro-gas activist (or an entire network of activists) could set up a code which automatically posts plenty of tweets about the advantages of gas energy, maybe even containing wrong information. In addition, as stated by Edry et al. (2021), bot tweets usually contain more emotions than human tweets and are much more stubborn, showing less sentiment variations (Dickerson et al., 2014). Thus, the sentiment towards gas as an energy source will be skewed although this sentiment is not even human-made, at least not in its full extension. Therefore, it’s crucial to identify bot accounts and remove their tweets to clean the data before the sentiment is analysed.

To detect bots, different approaches exist. While crowdsourcing via manual annotation and graph-based methods inspecting the social network between accounts are rather rare, machine learning techniques have been used by various scholars. Two famous and widely used online bot detection algorithms are *Botometer* (formerly *BotOrNot*) (Davis et al., 2016) and *Bot Sentinel* (Bot Sentinel Inc., 2022). Both models take different features into account to predict the probability that a certain account is a bot. For Botometer, more than 1000 features (from the user’s network, user metadata, firends, temporal activity, tweet content and sentiments) are considered. Since these algorithms require access to the Twitter API, they had to be shut down due to the API restrictions taken in early 2023. Hence, a custom bot detection model had to be built.

Bot characteristics

To create a suitable bot detection model, a range of features retrieved during the data downloading process had to be defined which should be able to predict whether an account is automated (bot) or not (human). Herefore, knowledge about the characteristics of bots was necessary. While most researchers differentiate between humans and bots, some further include the term *cyborg* defined as ”a human assisted by bots, or a bot assisted by humans” (Yan, 2006, p. 191). Several character differences between humans and bots/cyborgs were identified in the literature. It was found that bots and cyborgs follow more accounts than they’re followed by (friends) whereas this ratio is close to 1 for human actors (Chu, Gianvecchio, et al. (2012), Dickerson et al. (2014), and Lundberg et al. (2019)). Tweet frequency and temporal regularity were observed to be additional relevant differences between bots and humans (Chu, Gianvecchio, et al., 2012; Dickerson et al., 2014; Ferraz Costa et al., 2015). Generally, (partly-) automated accounts were found to post more tweets than humans although bots usually have longer tweet breaks (Chu, Gianvecchio, et al., 2012). Once active, the number of tweets posted by bots exceeds that of humans, resulting in a shorter inter-arrival time (IAT) (time gap between two consecutive tweets) (Chu, Gianvecchio, et al., 2012). Moreover, Ferraz Costa et al. (2015) observed that many bots are triggered by timers, leading to periodic tweets and, thus, to a larger IAT standard deviation of humans. Compared to humans, bot tweet activity doesn’t decrease during the weekends (Ferraz Costa et al., 2015). The density

of URLs, user tags and hashtags mentioned in the tweet texts were further found to represent important factors to distinguish automated from non-automated accounts (Chu, Gianvecchio, et al., 2012; Dickerson et al., 2014; Ferraz Costa et al., 2015; Lundberg et al., 2019). Including a diversity component containing these three densities, Kosmajac and Keselj (2019) found consecutive tweets of humans to be significantly more diverse than such of bots. Moreover, the type of tweets (normal, reply or retweet) (Lundberg et al., 2019), the profile age (Lundberg et al., 2019) and the device type the tweet was sent from were considered for bot detection (Chu, Gianvecchio, et al., 2012; Lundberg et al., 2019). Especially tweet frequency and URL, tag and hashtag density measures in combination with tree-based models like Random Forest and J48 or an ensemble of classifiers turned out to be successful at separating humans from bots (Chu, Gianvecchio, et al., 2012; Dickerson et al., 2014; Lundberg et al., 2019). Random Forest was also initially used by Davis et al. (2016) to train Botometer.

Training a bot detection model

Most of the aforementioned authors used semantic information (tweet content), tweet metadata and account metadata. However, since the respective user field was not included during the downloading process, the third type of data couldn’t be used for building the model. Consequently, feature selection for the bot model had to be adjusted on the available data (see Table 1). Moreover, while algorithms like *Botometer* inspect the most recent 300 tweets of an account (Davis et al., 2016), only the retrieved tweets about energy could be used here. These two main limitations restricted the model building process.

Making use of the available features, different indicators were created considering bot characteristics found in the literature. Additionally, some more features were introduced and tested in order to increase model performance. Table 3 lists all features and their description.

Indicator	Description
duplicates	checks how often the user is found in the previously extracted duplicate dataset, considers the mean length of the user’s duplicate tweets as short tweets are more likely to be unintentional duplicates
retweet average	how often all user tweets were retweeted by other users (normalized)
impression average	how often all user tweets were viewed by other users (normalized)
likes average	how often all user tweets were liked by other users (normalized)
reply average	how often other users replied to the user tweets (normalized)
number of tweets	the number of user tweets (about energy)
inter-arrival time	proportion of consecutive tweets with an inter-arrival time near 24 hours, set to 1 if only one user tweet was available (based on assumptions: propaganda bots post several tweets about energy; bots post more tweets than humans in their active periods leading to shorter inter-arrival times)

posting time	proportion of tweets sent between 6 a.m. and 11 p.m. (German timezone)
temporal diversity	Shannon diversity index showing the tweet type variability over time (how similar are consecutive tweets in terms of containing URLs, hashtags, mentioned users, see Kosmajac and Keselj, 2019)
URL density	how many URLs per tweet length (in words)
hashtag density	how many hashtags per tweet length (in words)
mentions density	how many tagged users per tweet length (in words)
is reply ratio	how many user tweets are replies (normalized)
tweet length	the average tweet length (in words)

Table 3: The indicators used to train the bot detection model.

To create training data, Botometer, that achieved an overall bot classification accuracy of 86% (Varol et al., 2017), was used before its shutdown to label 2131 random Twitter accounts found in my data. Botometer assigns a score between 0 and 5 to each account whereas a high value means a high probability of the account being automated (Botometer, 2023). However, finding a suitable threshold to distinguish between bots and humans in real-world scenarios is controversially discussed in the literature. When applying the most commonly used threshold of 2.5, Gallwitz and Kreil (2022) found 47% of all all US congress members being misclassified as bots in 2018. Hence, Keller and Klinger (2019) suggest that thresholds should be higher than 2.5. Analysing German political party members on Twitter, they came up with a suitable threshold of 3.8. However, when Gallwitz and Kreil (2022) applied Botometer on the presidential account of Joe Biden (POTUS) in 2022, a bot score of 3.8 was returned. Since my data contains quite a lot of political actors, a threshold of 3.8 would probably lead to major errors. Thus, the threshold value was increased up to 4.3. Based on the this threshold, the 2131 Twitter accounts were classified into bots and humans. These were split into training and testing data using a 70% to 30% split.

Model performance

First of all, the model was trained using all features shown in Table 3. By checking the model performance metrics and the feature importances, different feature combinations were tested. The retweet average, impression average and inter-arrival time features were omitted since they couldn’t increase the performance. On the contrary, the tweet length, the URL and mention densities and the duplicate occurrence were the most decisive features. In accordance to Chu, Gianvecchio, et al. (2012) and Dickerson et al. (2014), a Random Forest classifier was found to perform best. Eventually, the most accurate model achieved an overall accuracy of 0.90 and a macro average F1 of 0.80 (see Table 4). However, the model had significantly more difficulties to predict bots. While 78% of the model’s bot suggestions were also bots as per Botometer, the model only caught 57% of all bots found by Botometer (and the respective threshold). Hence, the model was less strict

regarding bots as it labelled only 82 accounts as bots while Botometer labelled 112 as such.

Class	Precision	Recall	F1-score	Support
human	0.91	0.97	0.94	528
bots	0.78	0.57	0.66	112
accuracy	-	-	0.90	640
macro avg	0.85	0.77	0.80	640
weighted avg	0.89	0.90	0.89	640

Table 4: Performance of the used bot detection model on the Botometer labelled testing set.

As stated before, one main restriction of the trained model is its sole focus on tweets about energy. This data sparsity led to a limited amount of tweets that could be used to create the features/indicators shown in Table 3. From the 640 accounts of the testing set, only 372 had two or more tweets. When the relevance of the number of tweets per account on the bot classification was tested, it was found that the accuracy of the model increased rapidly as the number of *energy tweets* increased. This was especially the case for the bot class. While the bot recall was low (0.57) for all users, it already reached 0.71 when only considering users with at least three and 0.74 for users with more than four tweets.

After calculating the features for all 354’242 users, the trained bot detection model was finally applied on the whole dataset. Overall, 9.5% of all users were labelled as bots by the model which corresponds to the literature (Edry et al., 2021). In consequence, the tweet load was reduced by 17.9%. The *typical human* and *typical bot* were calculated by averaging all features used during the model building process. It was found that bots had almost four times as many tweets as humans but tweets from humans had significantly more retweets, replies, likes and impressions, a bigger temporal diversity, more mentions and were more often replies to other users. Moreover, tweets from humans had more words (22) compared to bots (13).

Due to the aforementioned limitations of Botometer, especially regarding a strict bot-human-separation, a manual evaluation is highly recommended by Gallwitz and Kreil (2022). Hence, all Twitter users Botometer and my model didn’t agree on were selected and manually evaluated using the Twitter website. This was done by inspecting the username, the follower-friends-ratio, the followers’ profiles, the interaction of the user (replies, retweets, likes), the tweet texts and the biography. From the 66 users with disparities between Botometer and the model, Botometer was found to be correct 38 times while the model seemed to be more accurate 28 times. Furthermore, 98 Twitter accounts from the test and train set were randomly selected, manually evaluated and then compared to the model predictions and the Botometer labels. Interestingly, the overall performance is similar (macro average F1 of 0.79 for the model, 0.80 for Botometer) with both models reaching their limits when it comes to the bot precision (0.57 and 0.53) but Botometer scoring a much better F1 for bot recall (0.73 vs 0.91). Accounts that were wrongly classified as bots by the model were manually inspected by selecting their energy tweets. It was found that most of these

users had only very few tweets while most of them contained an URL. As the URL density was one of the top three features used for the bot/human predictions, the model classified these users as bots although they likely seem to be humans.

However, it's important to note that distinguishing bots from human accounts is very challenging even for human-beings and became nearly impossible in times of generative AI. So, despite well-trained machine-learning algorithms, massive data loads and manual evaluations, identifying bots is connected to a lot of uncertainties. For instance, there are accounts who seem to be non-automated according to the semantic information in their profiles but tend to share many newspaper articles via Twitter. All these tweets then contain an URL. And although it's very likely that a person manually shared these articles, every sharing process was done semi-automatically as most newspaper websites provide the option to convert the article to a standardised Tweets. Hence, it was decided to evaluate them as bots as long as they don't show much interaction with other users (replies, likes, retweets). Still, it's possible that a normal human shared these articles but it's impossible to verify it via the retrieved Twitter data. Moreover, there are accounts that seem to post automated tweets but also have human-like interactions and tweets that are likely to be non-automated. A clear assignment for such users is impossible, either. Both these problems intensify if the user only has one or two tweets about energy as the prediction is then based on less data.

Apart from these evaluation uncertainties, the model building process had limitations as well, mainly induced by Botometer which was the labelling algorithm of the training data. Just like the trained model, Botometer also had to deal with the problem of tweet rarity. As Gallwitz and Kreil (2022) found by manually evaluating Botometer, more than 20% of the non-automated accounts that were misclassified as bots only had one tweet. The authors reproduced such situations and found all accounts with one single tweet to be labelled as bots with a very high confidence score.

Conclusively, it can be noted that the trained bot model reaches a decent performance. While precision is nearly at 80%, the model is less strict on labelling a user as a bot, leading to a lower bot detection recall. However, this tolerance is acceptable for this thesis as the sentiment analysis task will include a daily sentiment cap per user. Further uncertainty comes from Botometer itself which is error-prone when it comes to users with few tweets and when a separation threshold is defined. Hence, neither could this bot detection task eliminate all bots nor could it catch all human users. The accuracy of the model could have been increased if user metadata, profile information and all tweets from a user's feed were retrieved.

3.3.3 Aspect separation

The problem of multiple aspects within one tweet as described in [subsection 2.3.3](#) is also present in the downloaded data. Especially during the energy crisis, Twitter users tend to mention more than one energy source within a tweet. They might compare different types of energy to argue for their preferred solution. In the fictional example "Es ist völlig absurd, den ganzen Strom mit gefährlichen Atomkraftwerken zu produzieren, wenn es dafür die sichere, umweltfreundliche und

erneuerbare Lösung der Photovoltaik gibt”, two aspect terms (‘Atomkraftwerken’, ‘Photovoltaik’) are mentioned. It’s crucial to assign the opinion terms (‘absurd’, ‘gefährlichen’, ‘sichere’, ‘umweltfreundliche’, ‘erneuerbare’) to the corresponding aspect terms. To perform a sentiment analysis on such data successfully, a very powerful aspect-based sentiment model comparable to Aßenmacher et al. (2021) would be necessary. Training such a model would take an enormous amount of effort, resources and time. Hence, the initial goal was to simplify the data by separating the occurring aspects and divide the sentence into two parts. Then, a sentiment analysis model should be run on the document level for each clause. However, due to the challenging grammatical structure of the German language in combination with the informal writing style of social media posts, most of the separated clauses couldn’t properly be interpreted by the multi-lingual pre-trained sentiment model (XLM-T) from Barbieri et al. (2022). Though further rule-based ideas of including conjunctions, other signal words or machine learning algorithms to separate the tweets in a reasonable way were considered, the success of these methods was estimated to be too low compared to the expense. Due to the error-proneness of the aforementioned method, it was decided to rule out all tweets containing more than one energy source via keyword matching. Here, it was found that such single energy aspect tweets represent the majority of the data. While 61% of all water tweets didn’t contain a second energy source, this value reached 79% for tweets about nuclear energy. All in all, about 75% of the non-bot tweets were still left after those tweets with multiple energy aspects were excluded. Hence, this simple method was evaluated to be the most effective one when compared to the complexity, time-intensity and uncertainties of the other options. However, this simplified aspect-separation method still reduced the amount of data that could have made Twitter sentiment more representative. For the excluded tweets, it was further found that especially renewable energy sources often co-occurred.

3.3.4 Thematical relevance

After the data was cleaned regarding duplicates, bots and multi-aspect tweets, the question of a relevancy classification filter arose. Due to the retrieval methodology based on keyword matching, even tweets that weren’t necessarily relevant for the energy discussion were retrieved. A fictional example would be ”Heute wurden im Wald beim Windrad zwei Bären gesichtet”. While the keyword ‘Windrad’ led to the retrieval of this tweet, the tweet doesn’t focus on the windmill as a form of energy. Hence, it’s not a relevant tweet regarding the research questions.

Although different machine learning algorithms like Support Vector Machine, Random Forest classifier or Neural Networks in combination with word embeddings and n-grams were used to classify tweets about disaster events into relevant and irrelevant ones (Habdank et al., 2017), attempts of training a RF classifier were not successful. Using SpaCy (spaCy.io, n.d.) to build vector representations of the tweets led to a decent overall accuracy but a very bad F1 score for non-relevant tweets. Errors are suggested to origin from the massive amount of tweets compared to the small manually labelled training and testing set and the high vector dimensionality of SpaCy vectors. Although more tweets could have been labelled, the timely effort of doing so was estimated to be

disproportionate regarding the outcome. Reason for this assessment was the data itself. While skimming through thousands of tweets and labelling hundreds of them, it was found that most retrieved tweets are directly linked to the energy topic.

To quantitatively test this observation, a sensitivity analysis was performed investigating 100 random tweets from the duplicate-, bot- and aspect-cleaned dataset. It was found that in 85 cases, the main underlying topic was energy. Not surprisingly, there's still an uncertainty in this assessment as it's not always that clear what the main intention of the author was, especially for tweets that were replies to other tweets. If unclear, the tweet was also assigned to the not-relevant class, making up the rest of the tweets (15). Since building a powerful relevancy classifier could be a thesis itself, this value is still considered to be sufficient for the upcoming sentiment analysis. Moreover, while applying the final sentiment analysis model (see next chapter), 14 of the 15 unclear or irrelevant tweets were labelled as neutral. Thus, the error for the sentiment analysis is expected to be minor. However, extrapolated on the entire dataset, even such neutrally labelled tweets incorrectly influenced the final sentiment results, limiting the validity of the findings.

Conclusively, after the removal of duplicates, bots and multi energy aspect tweets, a total of 2'610'290 tweets were left for the three upcoming main tasks (Table 5), namely Sentiment analysis, content inspection and geospatial separation.

Energy source	Nuclear	Coal	Solar	Wind	Water	Gas/oil
tweets left	1'086'646	234'171	336'591	525'749	44'869	382'264

Table 5: The total amount of remaining tweets after the pre-processing tasks.

3.4 Sentiment Analysis

This first main task was time-intensive as two different methods were set up to analyse the Twitter users’ sentiments about the defined energy sources. Reason for this was the unsatisfying performance of existing models for this thesis’ specific sentiment analysis task.

As found by Schmidt et al. (2022), BERT models are the most promising approach when it comes to sentiment analysis. Due to the complexity of the German language and the irregularities of social media posts, it was clear that such cutting-edge models should be applied to answer RQ 1.

3.4.1 Unsuitable document-level sentiment analysis models

The fully pre-trained and fine-tuned models of Guhr et al. (2020) and Barbieri et al. (2022) (XLM-T) can be accessed for free via the Hugging Face online repository (Hugging Face, 2023) and applied for sentiment analysis on German documents. This has been done for the pre-processed (duplicate-, bot- and aspect-cleaned) Twitter data after tags and URLs of the tweets have been removed as they didn’t provide valuable information but rather interfered the sentiment analysis.

Despite the encouraging F1 score found by Guhr et al. (2020), their model’s performance for these energy tweets was insufficient. While the model was able to identify simple statements and assign a correct sentiment polarity with certainty, it failed as soon as multiple clauses were included. Although the performance of the XLM-T model by Barbieri et al. (2022) was slightly better, it started to struggle when it came to more complex tweets of which there were a lot of. The model especially had major problems detecting positive tweets, resulting in a very low recall. Accordingly, these models were not found suitable for the tweets used in this thesis.

One of the main characteristics of the analysed tweets that led to so many errors was the aspect. Despite having ruled out multiple aspect energy tweets in the pre-processing phase, the tweets could still contain other, non-energy aspect terms. For example, the sentence ”Es nervt mich, dass die Regierung einfach nicht versteht, dass Solarenergie uns retten würde” only contains one energy aspect (solar energy) but, additionally, mentions the government as another aspect. Since both previously mentioned models analyse the sentiment on the document level, they can’t differentiate between these two aspect terms and their corresponding opinion terms. Hence, the overall sentiment of the whole tweet is evaluated which leads to errors since it’s not of interest what the Twitter users think about the government but only about the energy source.

Due to the insufficient performance induced by the complex tweets and the too coarse analysis level, I decided to fine-tune my own custom model which should be able to perform an aspect-based sentiment analysis (subsection 2.3.3) on the Twitter data.

3.4.2 A custom aspect-based sentiment analysis model

Next to fully user-ready sentiment analysis models, Hugging Face (2023) also provides a variety of pre-trained BERT models and respective libraries to fine-tune them for a specific tasks by own data.

Different BERT models exist that mainly differ regarding parameter training size, language trained on and case sensitivity. While there are *large BERTs* and *base BERTs* (Aßenmacher et al., 2021), Sanh et al. (2019) came up with the smaller but still competitive *DistilBERT*. While fine-tuning different BERT models for German aspect-based sentiment analysis on the GermEval17 dataset of Wojatzki et al. (2017), Aßenmacher et al. (2021) found the *bert-base-german-dmbdz-uncased* from the Munich Digitization Center (MDZ) of the Bavarian State Library (dbmdz, 2022) to be the most accurate one. Hence, I also opted for this dataset.

To fine-tune the model, the methods of Aßenmacher et al. (2021) were closely followed. Moreover, the official Hugging Face manual for fine-tuning a pre-trained model (Hugging Face, 2021) was considered since it provided more detail of the technical side than any research paper. For a deeper understanding of the processing steps, the free preview of Rothman and Gulli (2022) was consulted as well.

Opposed to Aßenmacher et al. (2021), who made use of 26'000 labelled documents from Wojatzki et al. (2017), I had to create my own training data. For each pre-processed energy source dataframe, 10 sets of 500 tweets each were randomly created. Since the downloading process via keyword matching also retrieved tweets when the respective keyword was found in the attached link but not in the tweet text itself, those *link-only tweets* were ruled out since only the text will be passed to the model later on. Via a python widget using *Tkinter* (Shipman, 2013), the tweet sets were manually labelled as either *negative*, *neutral* or *positive* towards the respective form of energy. This time-intensive process also allowed microreading (Purves et al., 2022) which provided lots of insights into thousands of tweets and allowed a first glimpse at re-occurring topics, problems and challenges evoked by the data.

Labelling tweets was challenging. For example, the tweet "Merkel sagt, dass Atomkraft sicher sei" contains a positive statement towards nuclear power but it's referring to the former chancellor Angela Merkel. So, the Twitter user's sentiment is not retrievable. Such irrelevant and doubtful tweets were labelled as neutral to avoid major errors. To prevent confusion, sarcastic tweets were trained as non-sarcastic ones as models have great difficulties recognising sarcastic statements (Maynard & Greenwood, 2014) and would easily get confused.

After a certain amount of tweets was labelled, a first run to fine-tune the chosen BERT model was performed to get an glimpse at the performance. This was done via the *Pytorch Trainer* from the *transformer* library (Wolf et al., 2020) and the Hugging Face manual (Hugging Face, 2021). Therefore, the labelled data was split into training data (70%) and testing data (30%) and converted into a special *Dataset* dictionary. Then, the tweet texts were tokenized via the tokenizer of the chosen BERT model from dbmdz (2022). Here, the transformers library (Wolf et al., 2020) also allowed to add padding tokens to get equal lengths for the inputs. During tokenization, the sentences are split into words or subwords (tokens) which are supplemented with special tokens marking the begin and end of a phrase (Usuga-Cadavid et al., 2022), resulting in the token embedding layer (Hoang et al., 2019). Additionally, a segment embedding layer (perceiving

the affiliation of a word to a sentence) and a position embedding layer (storing the position of the word within the sentence) are added (Hoang et al., 2019). The tokens are then converted into numerical representations (input IDs) depending on the vocabulary size of the pre-trained model. In the case of the used BERT model, this added up to more than 30’000 different tokens (dbmdz, 2022) serving as the input for the encoders of the BERT model which will then convert these inputs to high-dimensional vector representations and decide which sequences are important to make a final sentiment prediction (in this case: *negative*, *neutral* or *positive*) once the model is fine-tuned for aspect-based sentiment analysis (Hoang et al., 2019). To fine-tune the BERT model for such a specific task, some hyperparameters need to be set. This requires experimenting with different combinations of these hyperparameters (namely batch size, learning rate, number of epochs and weight decay) as there is no standard rule guaranteeing the perfect setup since this highly depends on the data (Aßenmacher et al., 2021). However, Devlin et al. (2018) recommended some values for the hyperparameters in their initial BERT paper which acted as an orientation for me.

Despite working on the powerful Ubuntu instance, the batch size could not be increased since the processing would stop when larger batch sizes were chosen. Hence, batch size was kept at 8 which was the default value as per Hugging Face (2020). Regarding the learning rate, Sun et al. (2019) suggested to use a rather small value like 2e-5. By doing so, the model takes longer to learn but can overcome the *catastrophic forgetting problem* (Sun et al., 2019). Another problem observed after a few runs was *overfitting*, shown by decreasing accuracy but increasing evaluation loss (Salman & Liu, 2019). Hence, a weight decay parameter was set to reduce overfitting (Bos & Chug, 1996). Lastly, epoch size was varied between 3 and 7. Since the F1-score did not increase significantly with more epochs, epoch size was kept at 4 as suggested by Devlin et al. (2018). Just like S. Y. Kim et al. (2021), the *AdamW* optimizer was used to minimize the cross entropy loss.

Overall, 16 models were fine-tuned. After the first eight training runs, the hyperparameters weren’t changed anymore due to promising results and the time-intensity of the fine-tuning process. Table 6 shows the final hyperparameter choices made.

Hyperparameter	Number of epochs	Learning rate	Weight decay
value	4	2e-5	0.01

Table 6: Final hyperparameters chosen to train the aspect-based sentiment analysis model.

Once suitable hyperparameters were found, the performance of the model could primarily be improved by additional training data. The more data the model was fed with, the better it learnt to put its attention on the *energy words* and the corresponding opinion terms. While the first runs only contained a few hundred training samples, the final model was trained on more than 6400 tweets. Due to the uneven distribution of tweets per energy source, training data wasn’t equally distributed per energy source either. To portray the most important discussions, energy specific subsets of the top three days when tweet activity peaked were created and labelled. Additionally, artificial data was set up to mimic tweets that were found to be misclassified after the first model

runs (e.g. very difficult double-negations like "Ich finde den Kohleausstieg eine schlechte Idee"). Finally, due to an imbalance in the data (the *neutral* class was prominent), an amount of the negatively and positively labelled training data was oversampled. Although oversampling can lead to overfitting and a loss of precision (Yoosuf & Yang, 2019), it was found to improve the model without leading to a greater evaluation loss than without oversampling. So, the total amount of training data was increased to 8681 samples. The fine-tuning process via the transformers library (Wolf et al., 2020) incorporating the hyperparameters took more than 16 hours. The performance of the final model will be described in the result section.

3.4.3 Applying the fine-tuned ABSA model

Once the aspect-based sentiment model was fine-tuned, it could be applied on the pre-processed data. This was done on the Ubuntu instance via a loop which analysed the sentiment of all tweets successively. Since they don’t contain valuable information about the sentiment of a person, URLs and tagged users were removed before. Hashtags, however, are a common method to make statements (e.g. '#antiatom') and were thus kept. Thanks to the pipeline approach provided by the Hugging Face transformers library (Wolf et al., 2020), application of the fine-tuned model was straightforward once the folder containing the model, its configuration and the corresponding tokenizer was given. Not only did the model assign a sentiment label (*negative*, *neutral* or *positive*) to the tweet, it did further include the probability score (between 0 and 1) (T. M. Lai et al., 2020). The label (numerically represented as -1 for *negative*, 0 for *neutral* and 1 for *positive*) was multiplied with the probability score which resulted in a final *sentiment score* ranging between -1 (very negative) and +1 (very positive).

According to Bashir et al. (2021), retweets serve as some kind of acknowledgement for the initial tweets or, as S. Y. Kim et al. (2021, p. 4) put it, are a way to "express an individual’s support". Hence, it’s assumed that the more often a tweet is retweeted, the more people share the sentiment stated within the tweet. Thus, frequently retweeted tweets should be weighted more heavily. I decided that the sentiment of the tweet should be doubled at maximum if a tweet gets many retweets. However, this produces a bias in favour of publicly famous people since they have a bigger reach on Twitter. Hence, the retweets were included via a non-linear weighting process to tackle this overrepresentation of famous personalities.

Since the goal of the sentiment analysis was to get a timeline of the sentiment changes over the study period, a time span had to be defined summarising the sentiments per energy source. Due to the large amount of tweets, especially in recent years, it was decided to summarise sentiments per day. Although duplicate tweets were already removed, the same user could post multiple tweets with slightly different wordings but the same content and sentiment which would result in an overweighting of that user’s voice. Hence, a mechanism to avoid overvaluing the sentiment of very active users was integrated by summarising the sentiments per day *and* per user. Thereby, the sentiment of a user posting more than one tweet about a certain energy source per day was only counted once, taking the average of the user’s sentiment scores at that day.

3.5 Content Analysis

While the results of the sentiment analysis only show how sentiments about different energy sources changed over time, possible reasons for these changes should be investigated to understand what positively or negatively influenced the Twitter user’s opinions. To achieve this, a variety of different methods is used, including term frequency, word pairs, Latent Dirichlet Allocation (LDA) and micro-reading. These methods are described in detail in this section.

3.5.1 Topic modelling

Despite being a more advanced form of content analysis than term frequency methods, topic modelling is useful to ”explore a corpus” and ”identify different forms of discourse” (Purves et al., 2022, p. 77). It aims to ”find hidden semantics in document collection and cluster the themes as topics” (Kherwa & Bansal, 2019, p. 2). Hence, it served as the initial approach to get a thematic overview into the discussions behind the sentiment timeline. Concretely, LDA (Blei et al., 2003) was used as it probabilistically defines different categories (topics) of frequently co-occurring words (Jelodar et al., 2019). Since predominantly conjunctions and prepositions are not meaningful enough for content analysis, pre-processing was necessary. Apart from URLs and user tags, stop words were removed from the tweets while the remaining words were further lemmatized just like Dehler-Holland et al. (2022) did. Lemmatization transforms words into their root form (Purves et al., 2022). As per Dehler-Holland et al. (2022), it is more suitable than stemming when working with German text due to the large variety of richly inflected terms. Stop words removal and lemmatization was both done via the SpaCy library (spaCy.io, n.d.) on the Ubuntu instance. To further make the LDA process as efficient as possible, all lemmatized words were uncapitalised and all *energy words* that served as keywords during the downloading process were removed since the sentiment timeline of each energy source was analysed separately anyway. Additionally, special characters like ’rt’ or ’—’ which were not caught as stopwords, were also ruled out. Moreover, some irrelevant words which were observed during the sentiment model training phase were removed as well. For nuclear energy for example, there were many tweets about the famous song called *Kernkraft 400* by *Zombie Nation*. Since these tweets didn’t have anything to do with the energy source, the respective words were expelled if they co-occurred in a tweet.

Then, the *Gensim* library for python (Řehůřek & Sojka, 2010) was used to build a dictionary with all unique lemmatized words and corresponding token IDs which numerically represented each word. For each tweet, the tokens and the corresponding token count (the number of times the same token was found within a tweet) were stored in an energy-specific corpus. To successfully build an LDA model, the dictionary and the corpus are taken as input, accompanied by an integer representing the desired number of topics the data should be categorized in. While Karami et al. (2020) chose not less than 40 topics, this number was found to be too fine-grained for an initial overview. Hence, it was decided to let the LDA model find 10 appropriate topics while this number was later increased to find more specific themes.

Once the model was built, the characterising terms per topic could be accessed, sorted according to their statistical importance for the specific topic. Now, a list of words passed to the model would output the probability of each topic to fit the given words. To compare this with the sentiment timeline, the topic variations had to be temporally represented. Thus, a custom function was built which would successively take all lists of lemmatized tweet words per energy source, return the most probable topic per tweet and then summarise all this on a monthly basis over the whole study period. Finally, the function calculates how often a topic occurred per month, normalized by the respective monthly tweet volume. Eventually, a dataframe with the percentual proportions of each topic for each month was created which could be plotted to detect temporal variations per topic. Based on the LDA model and the assigned main topic per tweet, *positive* and *negative* topics were derived for each energy source. So, the temporal development of topic shares could then be compared to the observed sentiment variations.

3.5.2 Term frequency methods

Since topics consist of a bunch of characteristic words, topic modelling falls short at analysing specific words. Hence, a simple term frequency analysis was implemented as well to monitor the temporal usage pattern of words of interests. According to Purves et al. (2022), such frequencies of words can be effective in analysing large corpora.

Thus, a python code was set up to inspect the temporal abundance of chosen terms derived from the topics. While a topic of wind energy might have included the term 'Wald' among 9 other terms, the sole inspection of this term provides more insight into how central it was for the observed topic variations. Since tweet volume drastically increased over time, term frequencies were normalized on a monthly basis to allow a suitable comparison over time. Eventually, via the a custom function and a list of desired words, the temporal variations of these terms found within the corpus of a energy source could be visually plotted to detect ongoing dicussions. While word pairs as applied by Vrana et al. (2023) were found to be too computationally intensive, most frequent bi- and tri-grams (Purves et al., 2022) were inspected for a defined time frame to get further insights into prominent debates.

In combination with the sentiment assigned to each tweet, further observations were possible. First of all, the word cloud approach of Schmidt et al. (2022) was applied. Here, subsets of negative, neutral and positive tweets were built. For each subset, a word cloud was generated showing the most frequently contained words. As for the LDA model, the pre-processed lemmatized words were chosen, leaving out stopwords and energy words as they would have distorted the analysis.

3.5.3 Micro-reading

Although previously described methods that statistically make use of a BOW are useful for an initial analysis of text corpora, such macroanalysis approaches fail to tell a detailed story about

the content (Purves et al., 2022), or as Jockers (2013, p. 4) put it: "[...] these staples of the digital humanist's diet hardly satiate the appetite for more." Hence, so-called *microreading* is necessary to further uncover detailed information. According to Purves et al., 2022, p. 64, microreading "involves [...] reading and interpreting individual passages or texts that have been identified as potentially of interest computationally [...]." Conclusively, while computational text analysis approaches provide a useful overview of the rich volume of tweets and potentially interesting events which could have influenced sentiments, the manual inspection of tweets is an inevitable task to see behind this heavy curtain of data. Consequently, thousands of tweets were read in order to discover reasons for prominent sentiment changes.

3.6 Geospatial separation

In accordance to the research question, the users tweeting about energy need to be assigned to either Germany, Austria or Switzerland in order to monitor country specific sentiments about energy sources. Hence, the users’ home locations are of interest. To predict home locations, the tweet content (text), the network (friends and followers of the user) and context information (such as timestamp, GPS tags or location provided in the profile) are conventionally considered (Zheng et al., 2018). Just as for the bot detection model, continuous access on the Twitter API would have been required to make use of the network and most context information. So, the tweet text was the main available source that could be used to infer the home location from as tweet-specific GPS location is traditionally used for less than 1% of all tweets (Morstatter et al., 2013). Compared to other studies, this thesis posed one advantage: The precise location is not needed as – due to mostly country wide energy policies – only a country level distinction is necessary. On the other hand, the terminated access of the Twitter API didn’t allow to retrieve further, non energy related tweets of a user which usually simplifies location prediction from tweet content (Zheng et al., 2018). For example, Hecht et al. (2011) and Han et al. (2014) took tweets from users from a certain region to infer the home location based on the chosen vocabulary via machine learning algorithms. As the amount of energy tweets downloaded per user was very limited and no additional tweets could be queried anymore, such an approach didn’t seem to be auspicious. Consequently, concluding the home country of users only from the available energy tweets posed a major challenge.

3.6.1 Assumption to infer the home location from the tweet content

Although Hahmann et al. (2014) found that toponyms mentioned in a tweet rarely correspond to the exact location the tweet was sent from, Mahmud et al. (2012) tested a set of classifiers to predict the home location of a user. Concretely, the prediction accuracy of nouns, hashtags and place names found in the tweet text were examined. Furthermore, two heuristic classifiers were assessed. Results showed that the mentioned toponyms had the best recall performance, followed by a local place frequency heuristic building on the hypothesis that users mention near-to-home places more frequently. Based on these findings and – with regard to the coarse spatial resolution necessary for the geospatial separation (country level) – the following assumption was defined:

”People from country C predominantly post tweets about energy sources/projects/debates going on in their home country C and, thus, mainly mention toponyms from country C.”

3.6.2 Geoparsing, disambiguation and inclusion of GPS tags

Based on this assumption, the tweets had to be scanned for toponyms. This toponym recognition part was done via the *Stanza* library (Qi et al., 2020). Stanza was run on all tweets and extracted words it identified as toponyms. The found toponyms then had to be assigned to the respective country. *Nominatim*, a geocoding service based on *OpenStreetMap* (Maier, 2014), turned out to be the best option to do so as most other geocode services like the Google Geocode API were not free

of cost. Nominatim took a toponym as an input and returned a dictionary containing thorough spatial information such as coordinates or the corresponding country code. In case of multiple meanings per toponym, a list of dictionaries was returned. However, *Nominatim* had difficulties interpreting some spatial adjectives. For example, tweets containing the word 'deutschen', which in most cases for the used data refers to the country Germany, was always assigned to a bus station called 'Teutschen' near Bozen in Italy. As *OpenStreetMap* relies on voluntary contributors, the "quality and completeness of the information" is not guaranteed as per Maier, 2014, p. 3. Hence, it seems like the bus station was accidentally named as 'Deutschen', leading to *Nominatim* mapping all terms 'deutschen' to the country Italy. Thus, closely reading the results of the geocoding process was necessary to find misclassified toponyms. These misclassifications were then manually corrected via rule-based approaches to assign the correct country code. This was a time-intensive process whereby – due to the data load – not all misclassifications could be identified.

Nominatim was further used to disambiguate place names occurring in various countries. So, in case of multiple occurrences, the *more important* one was returned as it was assumed that people rather mention, for instance, the city of Baden in Germany instead of Baden in Switzerland since the latter is way smaller. Nominatim defines the importance of a place either based on the place type (city more important than village, etc.) or – if equal place type – on the extent (Nominatim, 2023b). However, as per Nominatim (2023a), the Wikipedia ranking system is applied for most places. Here, Nominatim scans the Wikipedia articles of all toponyms with the same name and counts the inlinks. Based on this, the places are ranked, leading to Baden in Germany (rank 10) to be more important than Baden in Switzerland (rank 16).

Although the home location of most users was derived based on the tweet text, the georeferenced tweets were also taken into account, similarly to what De Bruijn et al. (2017) did. It was assumed that people generally send tweets from their home country and that deriving the home location (on a country level) from GPS referenced tweets was more precise than deriving it from toponym mentions. Hence, if a user had georeferenced tweets, this location was prioritized. However, even if the share of georeferenced tweets was higher (1.53%) than found in literature (Carley et al., 2016; Hecht et al., 2011; Morstatter et al., 2013; Ryoo & Moon, 2014), the large majority of these tweets had Twitter encoded place IDs which couldn't be decoded. Thus, only those tweets with real coordinates could be handled by Nominatim (Maier, 2014) to derive the country location from.

4 Model performance

Hereafter, the performances of the aspect-based sentiment analysis model and the geospatial home location separation approach are documented.

4.1 Sentiment analysis model

Since the model fine-tuning was processed via four epochs, the model performance after each epoch could be evaluated by the PyTorch Trainer (Hugging Face, 2021). Due to its robustness towards unbalanced classification data (Vågerö et al., 2023), the macro average F1 score was selected as the evaluation metric. Additionally, the evaluation loss is returned which is a more complex measure for the difference between true value and model predictions (Yacouby & Axman, 2020).

Overall, the final model achieved a macro average F1 score of 0.77 on the testing data. Since the F1 score is only the harmonic mean of the precision and recall of all classes (Vågerö et al., 2023), the latter two measures were separately calculated, giving more insights into the true model performance. Table 7 shows the full classification report of the model. On the same data, the XLM-T model by Barbieri et al. (2022) for document level sentiment analysis scored an overall F1 of 0.48 which underlines the necessity of a aspect-based approach. The detailed performance per epoch is shown in section 9.

Class/metric	Precision	Recall	F1-score	Support
negative	0.76	0.82	0.79	775
neutral	0.84	0.79	0.82	1267
positive	0.71	0.72	0.72	562
accuracy	•	•	0.79	2604
macro avg	0.77	0.78	0.77	2604
weighted avg	0.79	0.79	0.79	2604

Table 7: Sentiment model performance on the testing data.

Apart from the testing data used during the fine-tuning phase, an additional set of 100 tweets containing all energy sources was manually labelled. Then, the custom model was applied while its predictions were compared to the manually added labels. Interestingly, the performance is worse compared to the testing data although the model has neither seen the testing data nor these 100 tweets. Overall, a macro average F1 of 0.74 was achieved with a slightly better performance for the neutral class (F1: 0.86) but worse results for the negative (F1: 0.72) and positive tweets (F1: 0.63).

Apart from the little training data, I expect that this might lie in the nature of the labelled data which was supplemented by oversampled and artificially created tweets as well as posts from dates with a high tweet frequency. Hence, the labelled data used for training and testing wasn’t completely random. Consequently, there was a slight tendency towards more important events and

– due to the massive data load – there were still numerous grammatical structures to express a sentiment that the model hasn’t encountered during training phase. Inspecting the misclassified tweets didn’t show distinct patterns as the errors were very diverse.

4.2 Geospatial separation approach

Assessing the performance of the chosen approach to predict a user’s home location is complex due to the time-intensive acquisition of ground truth data via the Twitter homepage. Hence, the testing set used to assess the performance was smaller. In total, 34 users were selected from the home location predicted datasets. Table 8 shows that the chosen approach could achieve decent results for German and Swiss users. Most of the user the predictor suggested to live in Germany or Switzerland were actually found to be from these two countries when inspecting their Twitter profiles. A surprisingly large number of users stated their home country in the profile location. However, the predictor had difficulties to catch all German and Swiss users. The other way around, the predictor caught all Austrians but additionally misclassified five users as Austrians although they were Germans or Swiss. This error originated from *Stanza* that identified words like ‘Supergau’ or ‘Ukrainer’ as toponyms which were then geocoded to Austria by *Nominatim* since a company and a street there were found to contain these character sequences. Both missclassifications were not found beforehand and, thus, couldn’t be manually corrected as done for several other words.

Class/metric	Precision	Recall	F1-score	Support
Germany	0.92	0.79	0.85	14
Switzerland	1.00	0.77	0.87	13
Austria	0.58	1.00	0.74	7
accuracy	•	•	0.82	34
macro avg	0.83	0.85	0.82	34
weighted avg	0.88	0.82	0.83	34

Table 8: Performance of the home location predictor.

Based on the predefined assumption that georeferenced tweets were likely to represent the user’s home location, the toponym derived home countries were compared to the tweet referenced home countries. Overall, about 1100 users had georeferenced tweets as well as tweets containing toponyms. Here, the classification report showed slightly worse results, again predominantly for Austria (F1: 0.60) while German and Swiss users matched better (F1: 0.84 and 0.78). In contrast to the comparison with the ground truth, the recall was lower than the precision. However, it’s not entirely guaranteed that the GPS referenced tweets represent the user homes accurately. Maybe, users predominantly activate the GPS when being on holidays for navigation purposes. Hence, the manual assessment seems more reliable.

5 Results

In this section, the findings of the previously described methods will be presented. In accordance to the three research questions, it will be shown *how the sentiments about the chosen energy sources changed over the study period* (RQ 1), *which events or circumstances were found to contribute to these variations* (RQ 2) and *how the sentiments differed between users from Germany, Austria and Switzerland* (RQ 3). First of all, a broad overview of the sentiment results – on a supranational as well as on a country-specific level – will be presented, followed by subsections explaining important sentiment changes by carving out relevant events and circumstances.

5.1 Sentiment of all German-speaking Twitter users




Figure 4: The final sentiment timeline showing daily sentiments about the chosen energy sources and respective uncertainties visualised via buffers. Graphs are smoothed over 365 days, hence not starting in 2007.

Figure 4 shows the final *sentiment timeline* for the six chosen energy sources between 2007 and 2023 with the year on the x-axis and the sentiment score on the y-axis. Values larger than 0 represent a positive sentiment about the energy source, those below 0 a negative one. However, values between -0.05 and +0.05 were defined as a *neutral* sentiment. The sentiment scores were smoothed over 365 days to make long term variations visible. Due to the smoothing and Twitter's limited popularity pre 2010, most of the depicted graphs only start sometime in 2009, depending on the publicity of the energy source. The buffers symbolise the uncertainty of the sentiment scores which depend on the number of tweets per day. Uncertainty was defined to be larger if the sentiment was derived

from a few tweets only, resulting in a wider buffer area, whereas sentiment scores inferred from numerous tweets were assumed to be more robust.

On the first sight, it becomes apparent that solar power and hydropower were positively perceived by German-speaking Twitter users while wind energy changed from slightly positive until 2015 to rather negative afterwards with a noticeable downfall in 2019. Sentiment about gas and oil was in the neutral range throughout the study period but registered a negative trend starting in 2019 resulting in a slightly negative perception of the energy source. Nuclear and coal energy on the other hand were constantly viewed as negative. While coal energy fluctuated until 2018, nuclear power reached a low by the end of 2011 and has since been more or less on a rise. After a remarkable negative trend, coal energy reached its most negative sentiment score in 2019 and recovered in a steep rise since 2021. Overall, sentiment about coal energy was the most negative while solar energy was perceived as the most positive one during the whole study period. At least for the temporally smoothed plot, coal, nuclear and wind energy showed the largest sentiment fluctuations with solar, gas/oil and especially water remaining more stable.

5.2 Country-specific sentiment differences

Figure 5 displays the country-separated sentiments of German, Swiss and Austrian users. Due to a strongly reduced tweet number of Swiss and Austrian users, the monthly average sentiment instead of a daily smoothed version (Figure 4) is shown here to avoid large temporally inaccuracies. Again, sentiment varies more strongly in early time periods since inferred from a smaller number of tweets. Due to the dominance of German users, their respective sentiment timeline strongly resembles the non-country-separated sentiment curve. At certain points in time, the sentiment peaks or lows are aggravated or attenuated by sentiments from Swiss and Austrian users. Overall, Twitter users from all three countries had similarly positive feelings about solar power. Hydropower scored remarkably positive sentiments in Switzerland, predominantly post 2018. Austrians, however, were less keen on water energy. On the contrary, they viewed wind power more positively than their neighbours, especially in the past few years when Germans and Swiss became increasingly negative about this energy source. Similarly, despite a general aversion, Austrian users perceived coal energy as less negative than their neighbours. Conversely, they were more critical about nuclear energy which Swiss users had an only slightly negative or even neutral sentiments on. No significant differences for gas and oil were registered. Yet, these energy sources seem to have the most difficult status in Switzerland.

In the following sections (starting on the page after the next), the focus is successively placed on the individual energy sources. Based on the sentiment variations and the tweet frequency, several time periods are more closely inspected to uncover underlying reasons for the sentiment changes.

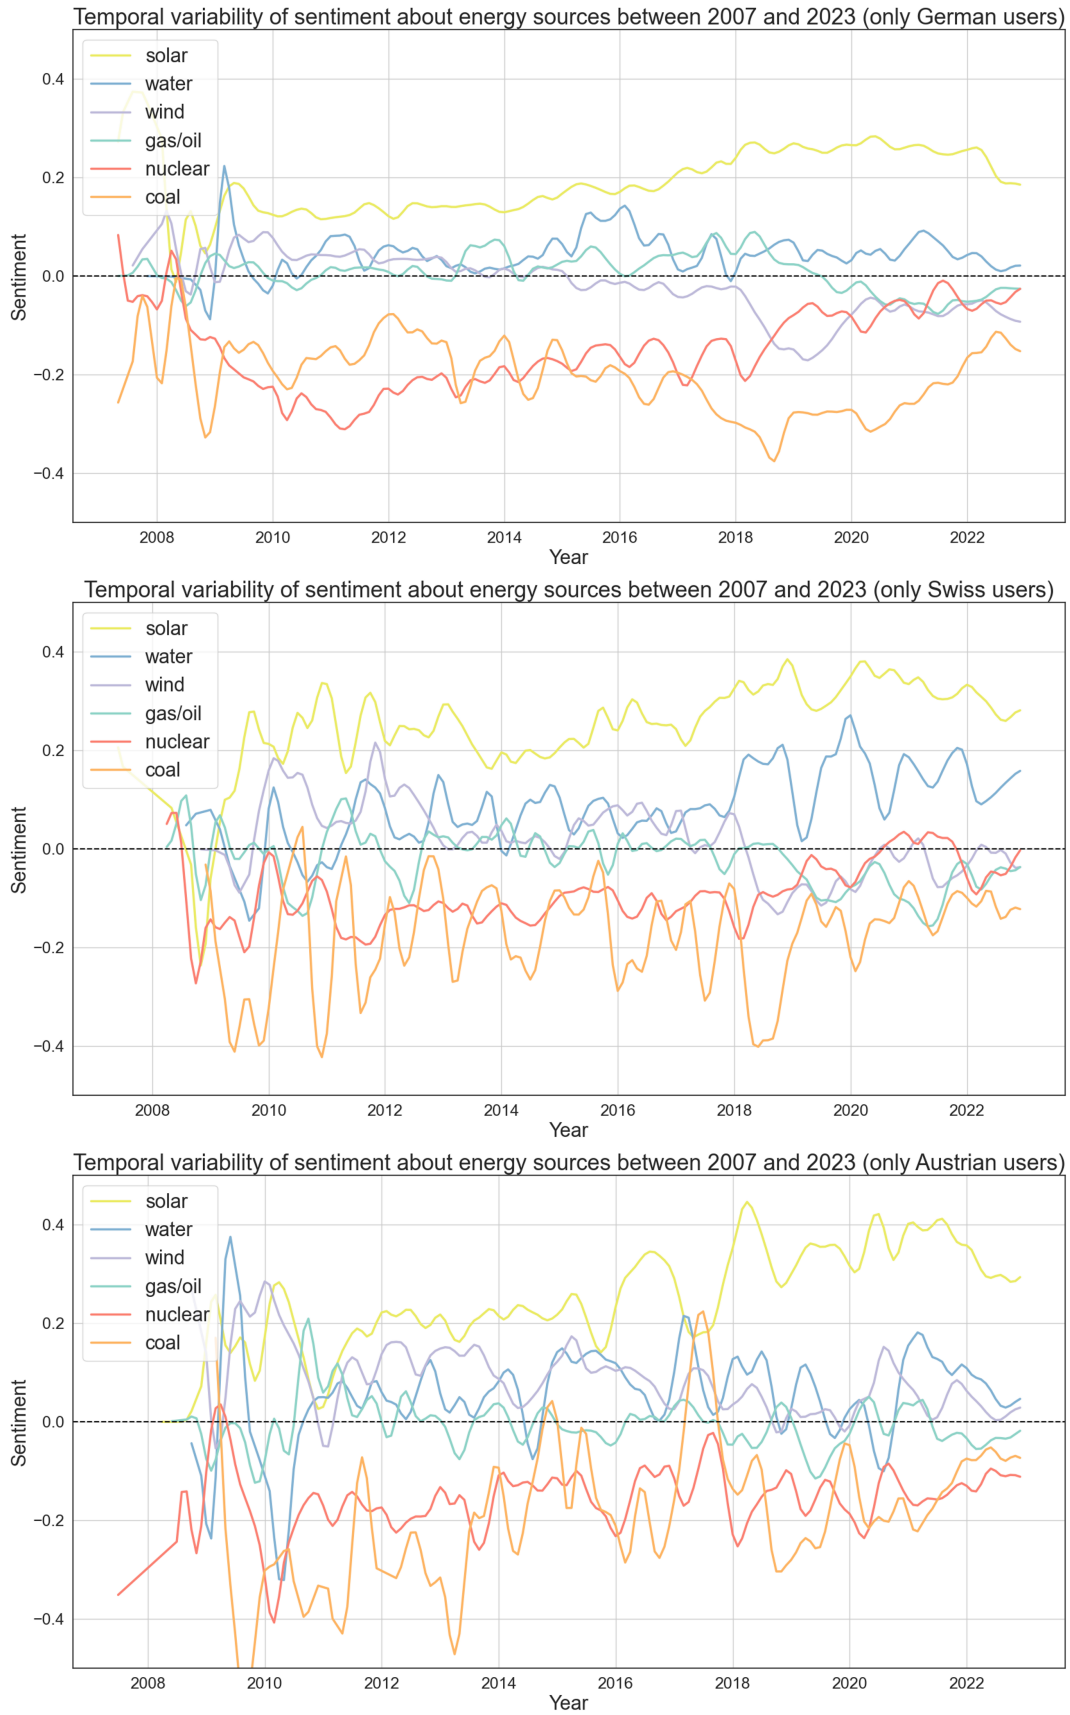


Figure 5: The country-specific monthly average sentiments between 2007 and 2023 (slightly smoothed by a Gaussian filter).

5.3 Nuclear energy

The graph shown in the conclusive [Figure 4](#) is plotted in a more smoothed variant in [Figure 6](#), showing more sentiment variations regarding nuclear energy. The general sentiment trend can be divided into mainly two phases: A strongly decreasing sentiment until summer 2011 when the sentiment reached an all-time low, followed by an increasing trend getting close to a neutral sentiment in winter 2021 before suffering a setback in the year after. As symbolised by the buffer range, [Figure 7](#) illustrates the tweet volume per month via transparency effects. While the tweet load was smaller in the first years of Twitter, it significantly grew in the past years, eventually leading to a sentiment that was based on a massive amount of tweets in 2022. This pattern applied for all other energy sources as well. Slightly deviating from the smoothed graph, this plot contains more precise temporal information showing that the the lowest sentiment score of 2011 occurred in May. An interesting inter-annual sentiment pattern could be observed in 2014, 2015, 2017, 2018, 2020 and 2021 as tweets from the first few months of these years were found to be more negative about nuclear energy than such from the rest of the year. Out of these years, February 2017 was the month with the largest share of anti-nuclear tweets, slightly interrupting the long-term rise towards the neutral sentiment range. What further stands out is the surprisingly large load of tweets posted in March 2011.

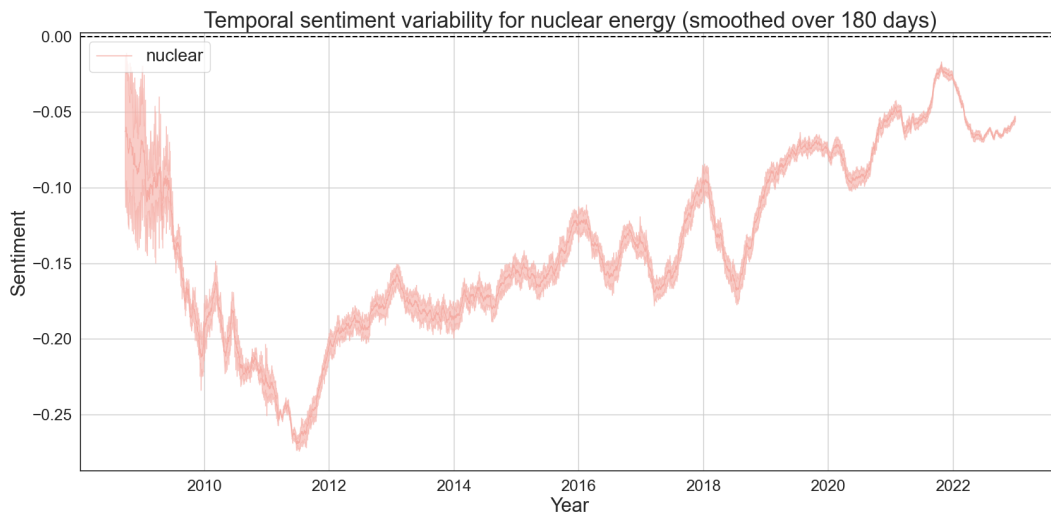


Figure 6: The sentiment timeline for nuclear energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

Due to this unexpected tweet volume in combination with a negative sentiment about nuclear energy, the content of Twitter posts from the first half of 2011 was specifically analysed to detect underlying reasons. Additionally, tweet content of the year 2022 was inspected as tweet load peaked and the sentiment was much less negative than in 2011, allowing to examine reasons for this sentiment shift.




Figure 7: The monthly average sentiment about nuclear energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).

2011: Anti-nuclear movement culminates after Fukushima disaster

Already by the start of the year 2011, the sentiment about nuclear energy was clearly negative. In the first two months, the word 'antiakw' was the most frequently tweeted one, usually via a hashtag. It was a characterizing term for topic 0 derived from the LDA model which recorded a drastic increase as more than 15% of all nuclear tweets were best represented by this topic in early 2011 (Figure 8). Not surprisingly, all tweets allocated to this topic were anti nuclear energy. Bi- and tri-gram analysis revealed that Twitter users specifically mentioned the two German nuclear power plants *AKW Neckarwestheim* and *AKW Grafenrheinfeld*. Via microreading, it quickly became clear that people were planning a large demonstration against the former by creating a 'Menschenkette' (human chain) and organising reoccurring 'Spaziergänge' (strolls). The demonstration was set to take place on March 12. At the same time, people were opposing against 'castor', referring to the transportation of nuclear material from German power plants to temporary storage facilities organised by the *Castor* company (Götz, 2011). Such anti-nuclear demonstrations had taken place before, for example in April 2010 when a 120 km long human chain was formed to protest against German nuclear power plants (Möhl, 2010). This event was frequently commented by anti-nuclear Twitter users resulting in the dominant sentiment low shown in Figure 7. The *AKW Grafenrheinfeld* was mentioned in connection with a possible crack found in a tube of the nuclear power plant in summer 2010 which had only been reported to the federal department of environment months later (Spiegel.de, 2011). This led to further anti-nuclear tweets by Twitter users.

Surprisingly, even more prominent than the both German nuclear power plants, the *AKW Mühleberg* was found to be the most frequently mentioned bi-gram in the first two months in 2011. Most of

such tweets were posted on February 13 as a popular vote took place in the canton of Berne (Berner Zeitung, 2011). Hence, most people were reporting results and other neutrally perceived information while many users stated their negative sentiments about this project, fearing negative health consequences from nuclear waste. From 46 non-neutral tweets, only a single one was positive towards the planned nuclear power plant. Although, 51.2 % of the population voted in favour of an upgrade for the *AKW Mühleberg* (Berner Zeitung, 2011).

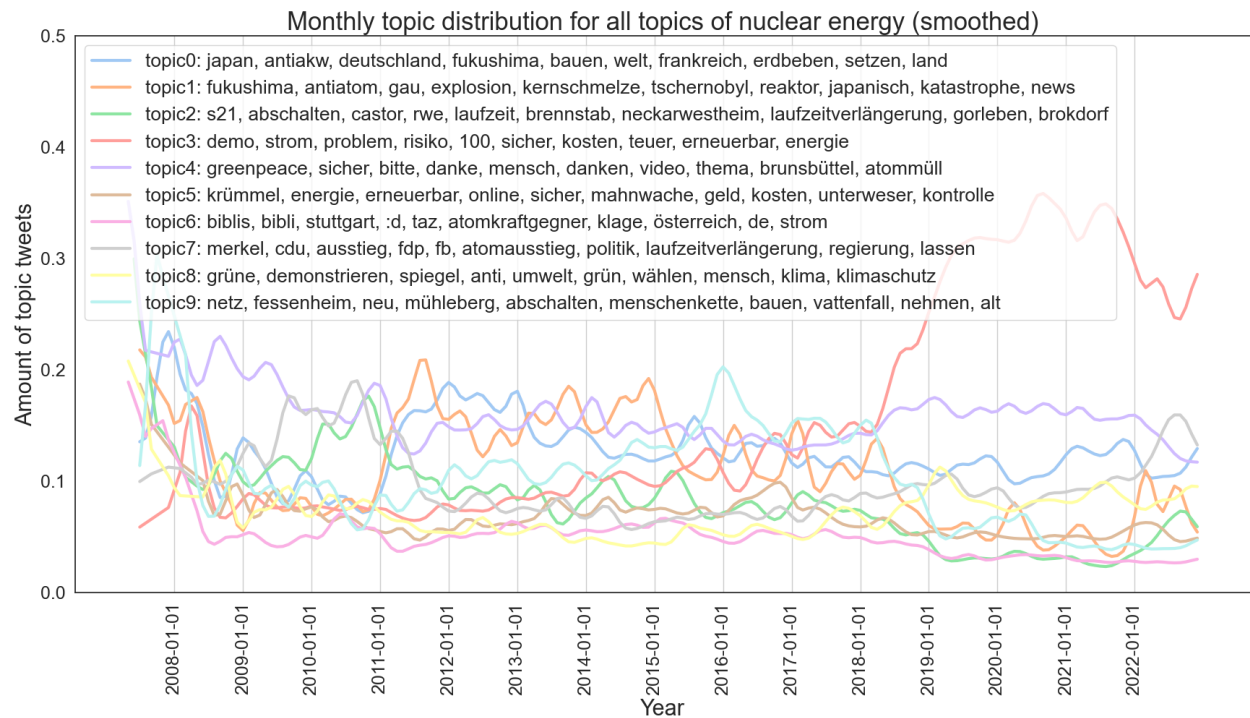


Figure 8: The temporal evolution of the ten different topics for nuclear energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.

While there were 2600 tweets about nuclear energy in the first two months of 2011, the tweet volume in March increased twentyfold. A glimpse at the characterising terms of the skyrocketing topics 0 and 1 and the most prominent tri-gram 'explosion, akw, fukushima' reveals the well-known reasons behind it: A massive earthquake of magnitude 9.0 occurred near the Japanese coast on March 11, triggering a tsunami which heavily damaged parts of the *Fukushima Daiichi* nuclear power plant (Baba, 2013). Due to the damage, the emergency cooling system couldn't be activated which led to a nuclear meltdown in reactor cores and the discharge of nuclear material to the environment (Baba, 2013). The whole world was shocked by this disaster. People on Twitter were posting thousands of tweets about it in March and beyond. While there were many neutral reports of the events, even more users expressed their negative sentiments about nuclear power technology using this disaster as an example for the danger it can cause. According to the strongly negative sentiment of topic 1 (Figure 10), the perception of nuclear energy reached values from almost -0.3

in the following months until the share of that topic finally decreased a little by the end of the year. Once again, calls for demonstrations and a nuclear phase-out could be observed very frequently. Consequently, the term 'abschalten' was one of the most prominent ones found in anti-nuclear tweets (Figure 9). In all three countries, March 2011 stood out in terms of tweet volume as well as a clear anti-nuclear sentiment stated by Twitter users. Despite still being very critical, Swiss Twitter users were less negative than German and Austrian ones (Figure 11). Sentiment about nuclear energy even dropped further in the same year, at least from German and Swiss users. On September 12, Twitter users and media outlet wrongly reported an explosion in the French nuclear power plant in *Marcoule*. However, a furnace for scrape metal exploded, not located in the power plant itself (IRSN, 2011). Hence, no radioactive material was set free (IRSN, 2011). Still, this further intensified the demand for a nuclear phase-out among users who were afraid due to the geographical proximity of the accident.



Figure 9: The word clouds for positive and negative tweets about nuclear energy.

2022: Health versus climate-friendliness and energy supply

After the years 2010 and 2011 when negative sentiments about nuclear power dominated on Twitter, this sceptic perception recorded – despite reoccurring months with very negative sentiments – a rising long-term trend (Figure 7). In the last three months of 2022, the sentiment score was at about -0.05, derived from more than 90'000 tweets (October). Topic 3 showed an impressive rise over the years, primarily starting in 2018 (Figure 8), which correlates well with the overall sentiment. Despite losing share in 2022, it could recover by the end of the year and remain the most important topic by a landslide as every fourth tweet was allocated to it. Term frequency analysis found that this topic was mainly characterised by the words 'Strom', 'Energie', 'Problem' and 'sicher'. Closely reading tweets allocated to this topic found Twitter users mainly referring to two problems: On the one hand, the problem of nuclear waste as well as safety (*sicher*) and health concerns led to negative sentiments about nuclear power. This anti-nuclear argument significantly shaped the negative sentiments in 2010 and 2011 after the Fukushima disaster and momentarily returned in spring 2022 when the Ukrainian nuclear power plant *Saporischschja* was under attack (Raiwa et al., 2023). On the other hand, the problem of a lacking energy supply security (*sicher*) due to nuclear phase-outs and the energy crisis resulted in positive sentiments about nuclear power

in 2022. Even more, other pro-nuclear users further claimed that nuclear power was relatively secure (*sicher*) and especially clean or climate-friendly compared to alternative energy sources. Such arguments were barely present back in 2010/2011, but started to be mentioned in tweets more frequently since 2018.

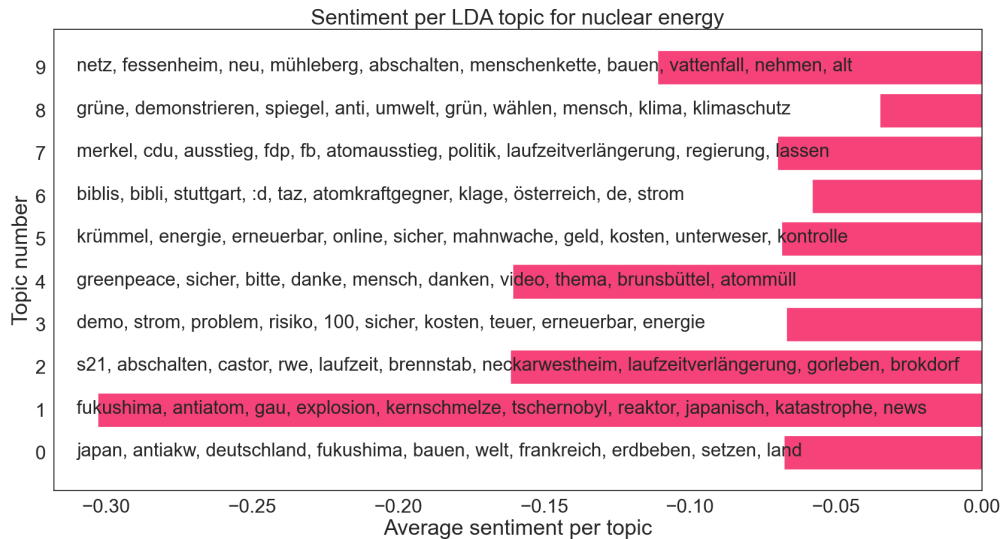


Figure 10: The average sentiment per nuclear topic of the whole study period.

For the last three months in 2022, there are almost as many pro-nuclear as anti-nuclear tweets containing the term 'Klima'. A lot of users described nuclear power as an indispensable measure to mitigate climate change and maintain energy supply at the same time. Here, some people referred to the European Union which included nuclear energy in the taxonomy of environmentally sustainable economic activities to facilitate energy transition for member states (European Commission, 2022a). Impressively, the tri-gram 'akw, laufen, lassen' was the most common one found between October and December 2022. The majority of these tweets came from German users whereas most of them were arguing for the usage of nuclear energy. This was during a time period when the German government and parliament were debating about the extension of their nuclear power plants which should have been shut down by the end of 2022 (Der Bundestag, 2022). Despite this rising support, there was still a majority of Twitter users arguing against nuclear power in all three countries of interest. However, a glimpse at Figure 11 reveals that nuclear sentiment of Swiss users was more positive than those of their neighbours. Even monthly positive sentiment scores could be observed in 2020 and 2021, but also in December 2022. Users particularly argued in favour of the climate-friendliness of nuclear energy and its reliability as a supplement to renewable energy sources, which they believe are too weak to provide sufficient energy. Those arguments were debilitated by users who pointed out that the construction of new nuclear power plants is too time-intensive, too expensive and poses high risks to health and environment. The latter were already the main arguments of Swiss people in autumn 2016 when nuclear power gained attention on Twitter prior to a public vote about the nuclear phase-out. Although the sentiment of Swiss users was negative

about nuclear energy, the people and cantons didn't want a fixed date for the phase-out and, thus, rejected the initiative (Bundeskanzlei, 2023).

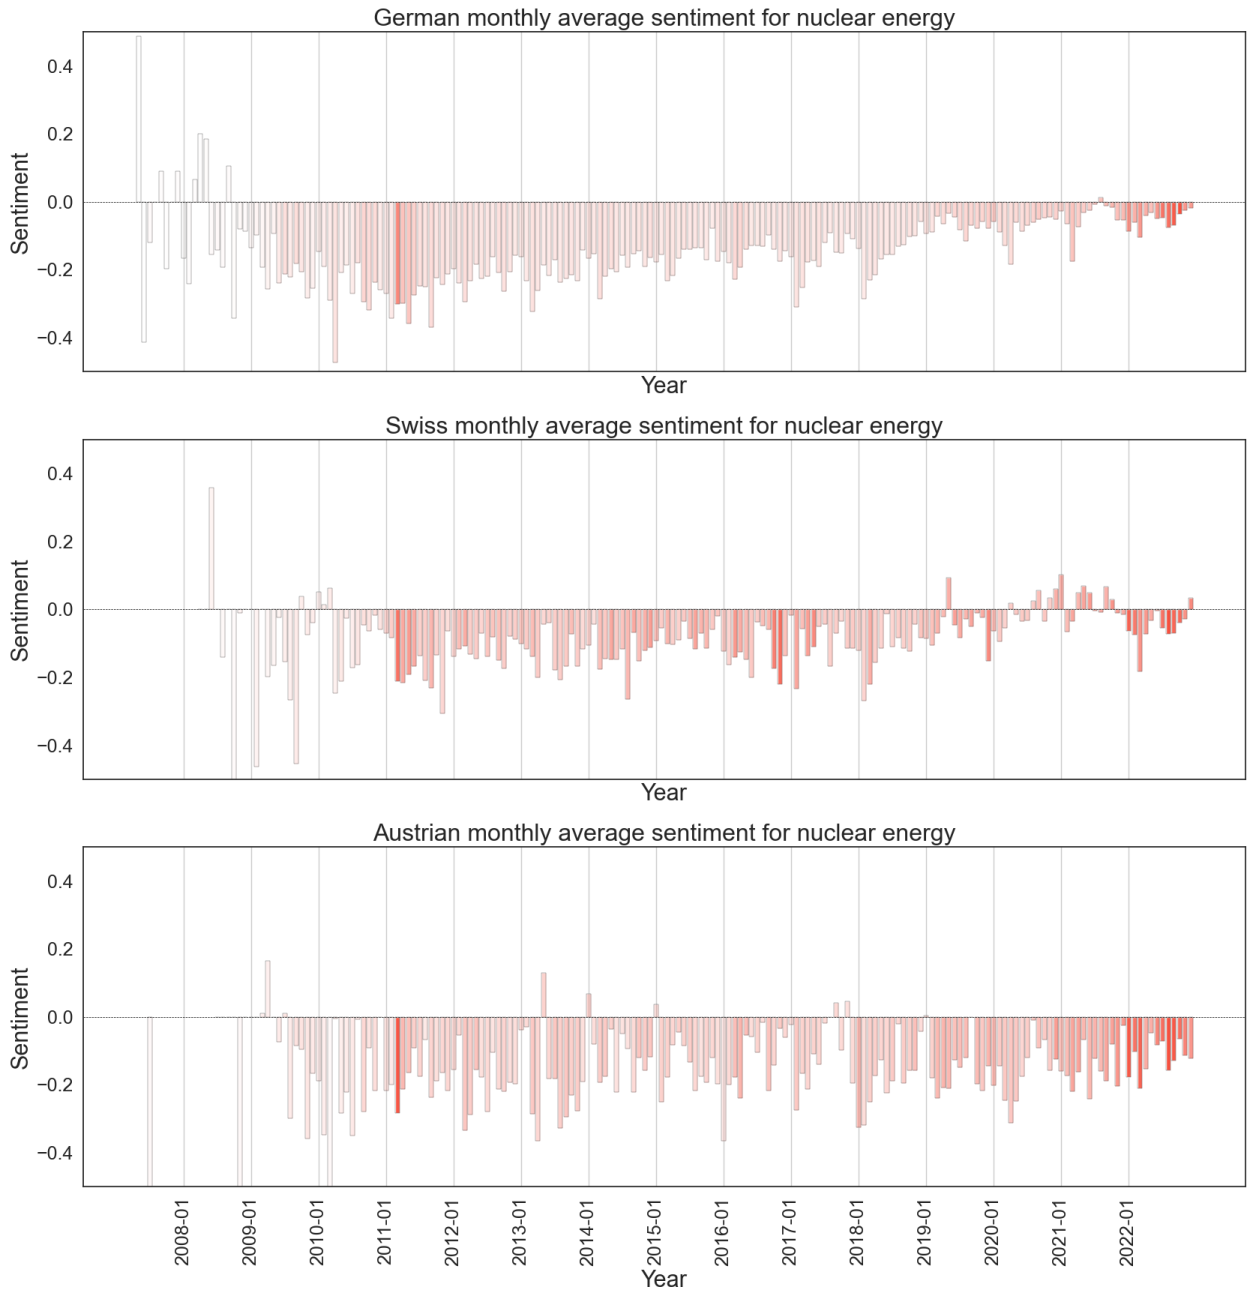


Figure 11: The country-specific sentiment timeline for nuclear energy per country.

5.4 Coal energy

Similarly to nuclear energy, the sentiment of Twitter users about coal energy was constantly negative when smoothed over 180 days as shown in Figure 12. In contrast to nuclear power, the sentiment was closer to being neutral in 2011/2012 and then started to decrease to reach a monthly minimum value of almost -0.5 in September 2018 (Figure 13). This was also the month with the highest tweet volume ever recorded. However, sentiment quickly recovered from this low and recorded an impressive uptrend that was only interrupted by an interim relapse in spring 2020. Afterwards, the sentiment has risen from about -0.35 and almost reached the neutral range scoring a monthly average of about -0.06 in summer 2022. In the last few months of that year, sentiment slightly dropped again but remained higher than in the past six years. Despite this upward trend, the overall sentiment only managed to reach the positive side in only four months of the study period while most of them were characterised by a small tweet volume and, thus, a large uncertainty.




Figure 12: The sentiment timeline for coal energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

Based on the monthly sentiments in combination with the tweet volume, I decided to have a closer look at the anti-coal months of September and October 2018, the almost neutral sentiments in June 2022 and the last few months of that year when Twitter users were more negative again. These temporal sentiment differences should allow to give insights into arguments and shifting discussions.




Figure 13: The monthly average sentiment about coal energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).

Autumn 2018: The fight for forest conservation leads to a strong opposition to coal energy

The term frequency analysis of September and October 2018 found that the words 'hambacherforst' and 'rwe' were prominently used in tweets, mostly accompanied by the hashtag 'hambibleibt'. 'Hambacherforst' has recorded several smaller term frequency peaks since 2012 but then especially gained a lot of attention in autumn 2018 which is also shown by the dominant topics 7 and 5 (Figure 14). The term refers to a forest near Europe's largest open-pit coal mine *Hambach* that is located next to Cologne, in a region that is famous for coal mining (Voss, 2022). As a form of protest against the *RWE Power AG*, the coal extraction company, anti-coal activists have been occupying the *Hambacher Forst* since 2012 (Kaufer & Lein, 2018). In mid-September 2018, the government of *Nordrhein-Westfalen* (NRW) tried to clear this occupation (Kaiser, 2020). Since this eviction happened just before RWE planned to clear the wood to be able to extract the underlying brown coal, the *Hambacher Forst* became the statewide symbol of the coal opposition movement (Kaiser, 2020). Hence, the large majority of tweets containing at least one of these terms expressed clear negative sentiments about coal energy. Most users pointed out the climate- and environmental damage caused by the combustion of coal, calling for a rapid phase-out and the continued existence of the forest and the transition to renewable energy sources instead. This strong opposition against the coal-extraction company led to 'RWE' being a prominent word among anti-coal tweets (Figure 15). On the contrary, there was a small minority of users defending coal for its secure and affordable energy supply or others who viewed the occupation as a futility considering that the "rest of the world" keeps constructing additional coal power stations. As

displayed by Figure 14, topic 5 registered an impressive peak in autumn 2018. However, this was found to be due to a spam of the *WWF Aktion* Twitter account which made up for about one fifth of all coal tweets in September and October 2018. In the tweet, *WWF* wrote "Braunkohle ist einer der klimaschädlichsten Energieträger. NRW verfeuert mehr davon als alle US-Bundesstaaten zusammen. Verlassen Sie diesen irRWEg! #StopptdenWahnsinn #kohlefrei" which resulted in most of the mentioned words to be found in the word clouds (Figure 15). After #kohlefrei, the first name and last name of a person was mentioned. This person varied in all tweets. Thus, the duplicate removal couldn't identify this spam. However, thanks to the author weighting restriction, only one sentiment per day from the *WWF Aktion* account (the daily average sentiment) was included in the assessment. Still, not all 3800 tweets were sent on the same day and, in addition, there were other Twitter users who copied this text and posted it in on their own account as well, linking the *WWF* website (link not available anymore). As the latter were strict duplicates, they could be removed in preprocessing. Still, the three most prominent topics during these months were also the most negative ones (Figure 16), resulting in this salient anti-coal period of autumn 2018.




Figure 14: The temporal evolution of the ten different topics for coal energy and their characteristic words derived from the LDA model, slightly smoothed by a Gaussian filter.

Summer 2022: The reactivation of coal-fired power plants in Germany arouses pro-coal users

After the groundbreaking events regarding Hambacher Forst declined, the dominant topics 5 and 7 significantly lost share. In June 2022, when the sentiment about coal energy reached the least negative value in years, topics 2 and 4 were dominant. While the former predominantly contained tweets discussing about the environmental implications of coal-fired power plants, the latter further involved a political dimension by frequently mentioning the Green party and Germany's *Minister for Economic Affairs and Climate Action* Robert Habeck (Die Bundesregierung, 2023). Just like in 2018, many users still argued against coal energy with references to its massive CO₂ emissions and air pollution. Some people also called coal-fired power plants "economic nonsense" claiming that coal-based electricity was more expensive than such of renewable energy sources. However, climate-related arguments predominated the discourse. A lot of them were triggered by an announcement of Robert Habeck who reacted to the gas shortage by announcing the reactivation of coal-fired power plants on June 19 (Rzepka, 2022). A lot of anti-coal users were shocked and disappointed by this sudden change of policy. On the contrary, pro-coal users who had criticised the government for being ideologically motivated or acting with double standards as coal-based electricity was imported from foreign countries now supported this decision. They described coal energy as a reliable energy source that is necessary to provide energy security during the ongoing crisis. While anti-coal tweets still outnumbered pro-coal ones, the gap was narrower than it has been for a long time.

While it's no surprise that tweets from German users prevailed, the ratio was extreme for coal energy as tweets from Austrian and Swiss users only made up for about 1% each. In comparison to Germany, their neighbours were less sceptic about coal energy (Figure 17). Both even recorded positive sentiments in June or July 2022, however, derived from a small amount of tweets. Swiss users were commenting on the German coal plant reactivation by expressing overall neutral sentiments about coal energy, contrasting its climate-implications with its reliable energy supply. In comparison to Swiss users, Austrian Twitter users registered a significant increase in tweet volume in June 2022. Reason for this was the dismay of many users about the governments decision to reactivate the coal power plant *Mellach* in order to become independent of Russian gas in case it should get scarce (ORF.at, 2022a).



(a) positive tweets

(b) negative tweets

Figure 15: The word clouds for positive and negative tweets about coal energy. Note that 'braunkohl' appeared twice as it was part of connected words like 'braunkohl-abbau' which couldn't be handled by the word cloud library.

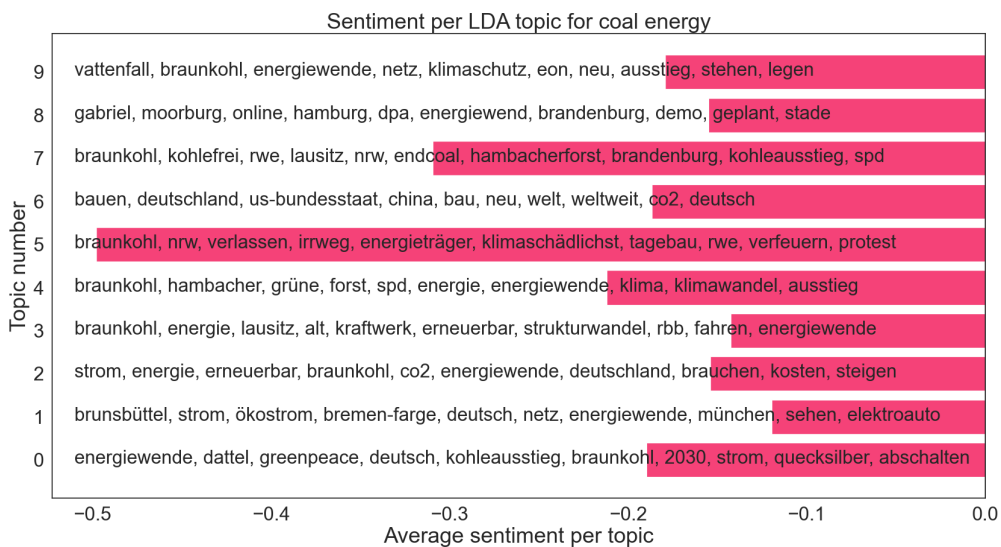


Figure 16: The average sentiment per coal topic of the whole study period.

October until December 2022: Ongoing consternation and a comparison to China

After the interim sentiment high in summer 2022, the sentiment about coal energy dropped again in the more negative range. The previously dominant topics lost share while topic 6 advanced to become the second most prevalent one (Figure 14). Between October and December 2022, the most frequent term found in tweets allocated to that topic was 'China'. For all tweets in that time range, topic independent, the bi-gram 'China, bauen' was the third most mentioned. The term frequency of 'China' was four times higher than in previous months. While no specific reason for this surge was detected, microreading uncovered lots of Twitter users from all three countries criticising efforts of ending coal energy. Users justified their critique by referring to China as the world's largest producer and consumer of coal whose coal consumption only peaked in 2020 (B. Zhang et al., 2023). Those people feared that their countries will face a tragic energy shortage and even higher energy prices during the energy crisis by phasing out of coal while other countries like

China keep building new coal-fired plants. Many of such tweets were addressed to climate activist movements, primarily the *Letzte Generation* who were advised to stick themselves to Chinese coal power plants instead of German roads. However, most of these tweets were found to be misclassified as 'negative about coal energy'. While there were only a few people who tried to weaken the China arguments by referring to the historic responsibility of middle-European countries or by calling it a poor excuse, anti-coal users were primarily found in topics 2 and 4 again. The consternation about the reactivated German coal power plants still existed. So, tons of tweets criticized the government by calling out the climate goals which are heavily endangered by the coal reactivation. On the other side, pro-coal Twitter users were still glad about the reactivation, once again pointing out the reliability and effectiveness of the technology in times of a potential gas shortage.

Similar to the discussion of nuclear energy, 2022 was characterized by a dilemma between being climate-friendly by reducing coal energy and tackling the energy crisis by using pre-existing facilities, in this case, coal-fired power plants. Yet, like in the previous years, the majority of anti-coal tweets led to an overall negative sentiment. However, due to the misclassification of tweets calling out China's coal production, I expect the real sentiment to be less negative.

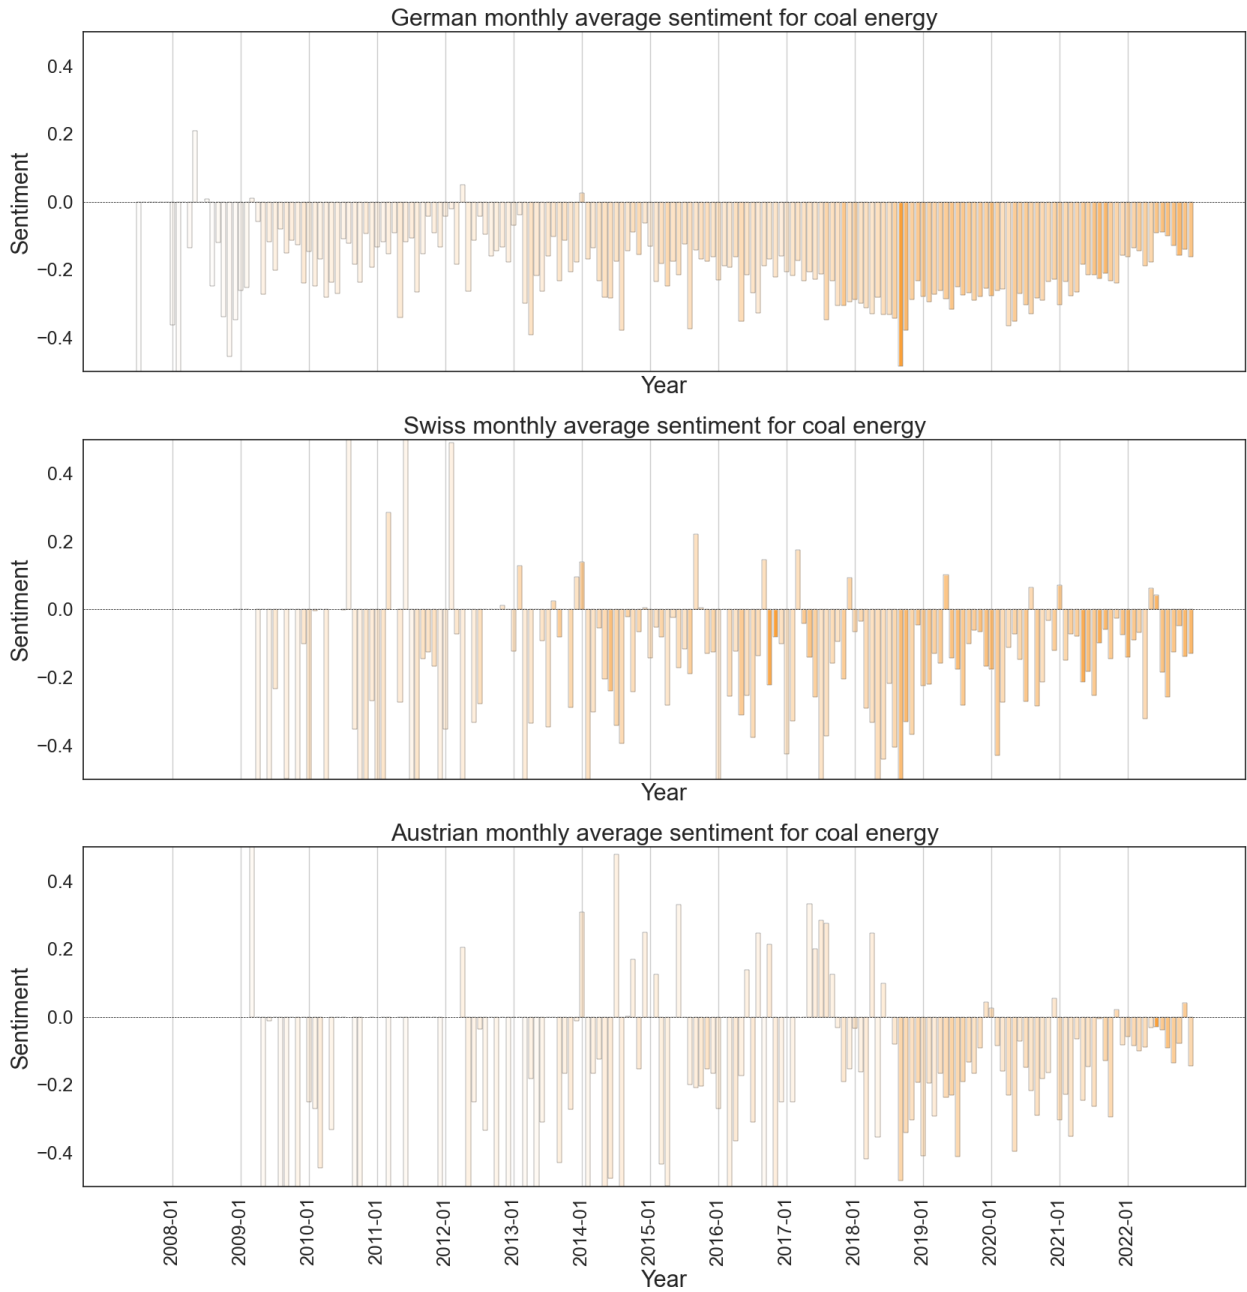


Figure 17: The country-specific sentiment timelines for coal energy.

5.5 Solar energy

In contrast to nuclear and coal energy, solar power was positively perceived by Twitter users over the whole study period (Figure 18). While the sentiment slightly decreased between 2009 and 2010, it has since been on a long-term rise, interrupted by some minor variations over the course of the years and an abrupt downfall starting in summer 2022. The most positive sentiment was reached in September 2020 when the sentiment score reached 0.28, making it the highest sentiment score observed from all energy sources. As shown by the transparency effects of Figure 19, the tweet volume gradually increased over the years. As opposed to nuclear power and coal energy, months with an unexpected, prominent increase in tweet volume cannot be found. But in December 2015 and 2017, the average monthly sentiments significantly deviated from the surrounding months. However, this pattern couldn't be observed in the other years. The plot also documents the sentiment decline in summer 2022 which was found to be almost steady, except for October and November. In that year, the monthly average sentiment fell from about 0.28 to 0.19 and registered the most negative value since 2017. Despite this decline, solar power was still by far the most positively perceived energy source in 2022 (see Figure 4). Over the whole study period, there were only three months, all before 2009, where a negative sentiment about solar energy was recorded.




Figure 18: The sentiment timeline for solar energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

Based on above observations, it was decided to closely inspect the sentiment valley in December 2017 in order to find underlying causes for this abrupt change. On the opposite, insights on people's perception in spring 2022 will be given to understand what made people appreciate solar power. Additionally, the last month of the same year will be given insights to find reasons for the sentiment decline.




Figure 19: The monthly average sentiment about solar energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).




Figure 20: The temporal evolution of the ten different topics for solar energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.

December 2017: Two anti-solar articles have a negative impact on the sentiment

In comparison to the sentiment setback in December 2017, no salient changes in the LDA topic shares were registered in that month (Figure 20). Only slight increases in topic 0 and topic 2 could be observed. However, the tri-gram analysis found very specific words like ‘überteuert’, ‘ineffizient’ and ‘verheerend’ to occur frequently in that month. Microreading revealed that those words belonged to tweets which referenced articles from the *Frankfurter Allgemeine Zeitung (FAZ)* and the *Basler Zeitung (BaZ)*. Both stories called out the allegedly energetic and economic inefficiency of solar power (Ferroni & Reichmuth, 2017; Mihm, 2017). The FAZ article with the headline *Solarstrom ist überteuert und ineffizient* compares photovoltaics with wind energy and reports that the former costs a lot more than the latter when it comes to avoiding CO₂-emissions as operators of solar power stations are better paid than such of wind power plants although wind power produces much more electricity (Mihm, 2017). The article gained a lot of attention and was bookmarked 300 times. For comparison, other FAZ articles about solar power have less than 20 bookmarks. The BaZ article named *Die verheerende Bilanz von Solarstrom* comes to an even more devastating verdict, claiming that the installation of photovoltaic panels in Switzerland consumes more energy than they can produce in their lifetime (Ferroni & Reichmuth, 2017). Both articles were frequently shared on Twitter. From 101 anti-solar tweets in that month, 43 referred to one of these two articles which led to a less positive sentiment than observed in other months. However, the majority of users still felt positive about solar energy reporting how they installed their own PV modules, how solar power had great potential and was desperately needed for the energy transition or – in answer to the FAZ article – how it is still the cheapest option to produce electricity for private consumption.



(a) positive tweets



(b) negative tweets

Figure 21: The word clouds for positive and negative tweets about solar energy.

Spring 2022: Personal reports and independence in geopolitically unstable times

As shown in Figure 20, the LDA topics 0 and 7 dominated in spring 2022 when the sentiment about solar energy reached a high, especially in March. In that year, a good correlation of the topic 7 – containing mostly pro-solar tweets (Figure 22) – and the sentiment variations is observable. It was striking to see how many tweets which were allocated to topic 7 had a positive sentiment towards solar energy but were at the same time complaining about the bureaucratic hurdles for

house owners. These users endorsed solar energy but demanded to make the process of installing photovoltaic panels more straightforward as it seemed deterrent and, thus, hinders people from installing PV modules. There were also some users who had already installed solar power on their roofs and proudly presented the amount of energy it produced. Consequently, the term 'Dach' was one of the most frequent words found in pro-solar tweets (Figure 21). Some people further used real-life examples by writing how they were able to charge their electric car with their own electricity. Here, arguments of climate-friendliness, economic advantages and an independent electricity production resulted in the positive perception of solar power. The aspect of independence from the electricity market and geopolitical tensions was mentioned remarkably frequently, often referring to discussed import stops of Russian gas (Castanho Silva et al., 2022). On the contrary, anti-solar tweets often contained another dimension of independence as users criticised solar power for being dependent on weather conditions which would make constant energy supply impossible due to a lack of storage techniques and central Europe's subpar climatic conditions. Other users complained about the financial outlay which simply wouldn't make it feasible for a lot of people.




Figure 22: The average sentiment per solar topic of the whole study period.

Winter 2022: No solar energy without solar insolation

A negative sentiment trend in 2022 was registered for Twitter users from all three countries (Figure 23). While the yearly minimum sentiment of Swiss users was recorded in September, Austrian users felt the least positive in November. For all three countries together and for Germany specifically, the sentiment reached its yearly minimum in December. Just like the sentiment, the share of topics 0 and 7 declined as well while smaller topics gained some share. Nevertheless, topics 0 and 7 remained the most prevalent ones. Compared to the other investigated time periods, it was notable that the term 'scheinen' occurred much more frequently. The same applied to the most common bigram 'sonne, scheinen'. By inspecting anti-solar tweets, it was found that an unexpectedly large

amount of users argued against solar energy due to its dependence on weather conditions. Some of them connected their critique to the actual season stating that solar power was useless when the sun doesn't shine or insufficient when the days are shorter during winter months. This argument was further strengthened by users tweeting that solar power was not only unreliable but additionally very expensive. Despite this increase of seasonal and weather-dependent arguments, the sentiment about solar energy remained positive. Once again, people shared their positive experiences with photovoltaic, describing how they could cover a certain amount of electricity usage via their solar panels and praising this market-independent energy supply. Moreover, a discourse about nuclear fusion emerged. Origin of this debate was a scientific breakthrough at the *Lawrence Livermore National Laboratory* where scientists managed to perform a controlled fusion experiment that generated more energy than it consumed for the first time in history (Tollefson & Gibney, 2022). While some Twitter users depicted this as the universal solution for humanity, this situation was leveraged by numerous individuals to advocate for solar energy which already made usage of the sun as a huge nuclear fusion reactor. Lastly, it was the most pro-solar power topic 3 (Figure 22) which registered a slight increase. This increase primarily based on the word 'Wirkungsgrad' (level of efficiency). In pro-solar power tweets, Twitter users reacted to a new photovoltaic cell developed at the *Helmholtz-Zentrum Berlin* which reached an unprecedented efficiency of 32.5% (Mariotti et al., 2023). The German newspaper *Tagesspiegel* wrote an article about it (Tagesspiegel.de, 2022) that was frequently shared on Twitter, accompanied by positive sentiments about the technology. Still, it couldn't lead to a turning point to break the sentiment decline.

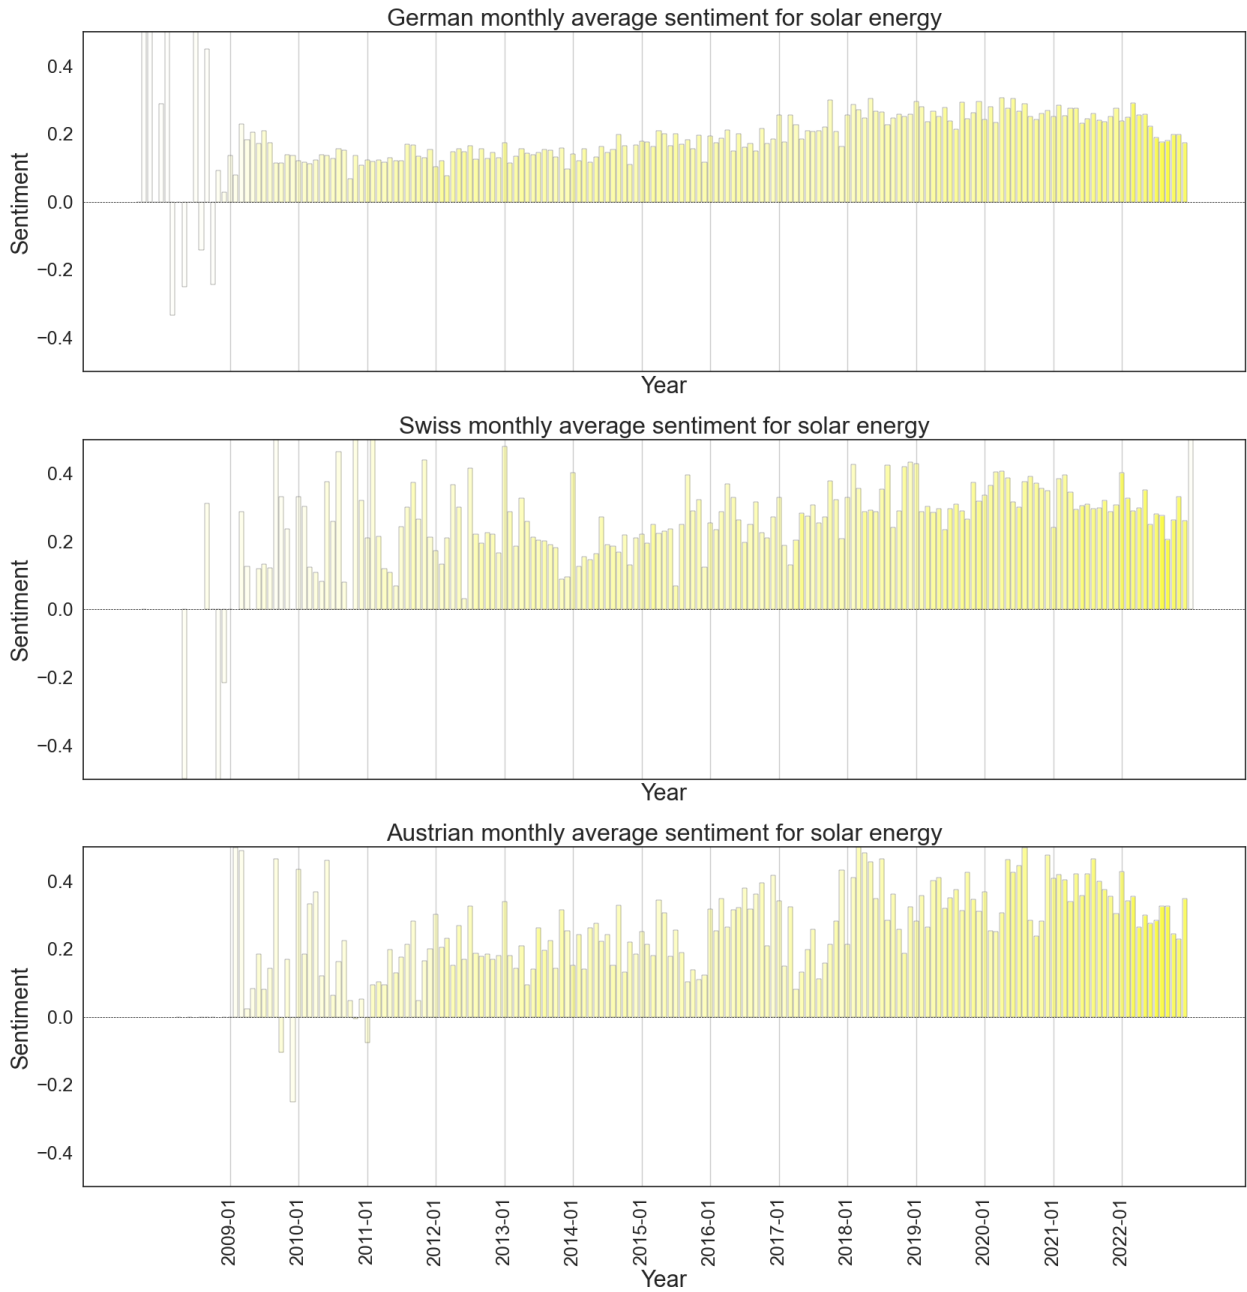


Figure 23: The country-specific sentiment timelines for solar energy.

5.6 Wind energy

In [Figure 24](#), a clear negative sentiment trend about wind power was observable on the long term. While Twitter users perceived wind energy as slightly positive or neutral until 2018, the sentiment drastically declined afterwards and eventually reached its most negative point in March 2019 ([Figure 25](#)). However, the sentiment recovered as quick as it dropped. Although, it never reached the positive side anymore but remained slightly negative throughout the rest of the study period. In July 2021, the sentiment registered another local minimum that was characterised by a large tweet volume. Similar to the previous observations, the number of tweets posted about wind energy peaked in 2022, predominantly towards the end of that year. In the same time – just as for solar power – sentiment saw another decrease. Compared to German Twitter users, the perception of Swiss and Austrian users varied more widely as [Figure 29](#) shows. In general, Austrians were more positive about wind energy than their neighbours.




Figure 24: The sentiment timeline for wind energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

It's of great interest to understand underlying causes for the abrupt sentiment downfall observed between 2018 and 2019. Hence, the following content analysis will put a focus on this time period. Furthermore, spring 2022 will be closely inspected to monitor reasons for the neutral sentiment as a contrast to 2018 and 2019.




Figure 25: The monthly average sentiment about wind energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).




Figure 26: The temporal evolution of the ten different topics for wind energy and their characteristic words derived by the LDA model, slightly smoothed by a Gaussian filter.

2018/2019: Great concern about forest areas causes support for wind energy to collapse

The temporal distribution of LDA topic shares provides great insight into the reasons for the sudden increase of negative wind sentiments in 2018 and 2019. As shown in [Figure 26](#), that negative trend temporally correlates very well with the massive increase of topic 4 which was found to unite mostly negative tweets ([Figure 27](#)) and became extremely dominant ([Figure 28](#)). In accordance to the characterising words of topic 4, term frequency analysis revealed that Twitter users suddenly mentioned 'Wald' and 'roden' much more often since summer 2018. As microreading revealed, in September 2018, a lot of users were annoyed by the a planned deforestation project in the mythenshrouded *Reinhardswald* which should allow the construction of large wind power plants (Rapp, 2019). Further evoked by the concurrent protests against coal energy in the *Hambacher Forst* (see [subsection 5.4](#)), these plans were widely and controversially discussed, also on Twitter. Thus, the large majority of Twitter users argued against wind energy as they prioritized the legendary forest ecosystem which would also have a positive impact on the climate. The German conservationist Hermann-Josef Rapp and specially founded conservation activist alliances heavily criticised the project due to its negative impact on biodiversity (Rapp, 2021). Although the agitation about these forest areas diminished by the end of 2018, topic 4 rose again in spring 2019 and reached its peak in summer when almost every third tweet was assigned to that topic. In March, when the monthly sentiment hit another minimum, the term 'Insekt' was mentioned in almost every seventh tweet, most of them posted on March 25th. This was triggered by numerous German newspaper articles published on that day, including one in Spiegel (Spiegel.de, 2019). These articles referred to a study of the DLR (*Deutsches Zentrum für Luft- und Raumfahrt*) from 2018 where the authors calculated that approximately 5 to 6 billion insects per day would be killed by wind turbines (Trieb et al., 2018). The sudden focus of several news outlets on this topic led to an enormous amount of anti-wind power tweets although the authors of the study and the news articles pointed out difficulties in interpretation due to a lack of quantified studies on other insect-damaging activities. In the following months, wind energy saw continuous critique which referred to deforestation and insects but was further complemented by numerous users bringing birds into the discussion. These users claimed that birds were also killed by windmills and disturbed by their noise and infrasound emissions, further resulting in negative sentiments about wind power.

Conclusively, the strong increase in negative sentiments emerged from several factors, most of them related to environmental drawbacks of wind power (deforestation, danger for insects and birds) in combination with plans for additional wind power stations. Additionally, topic 7 – the second most negative one ([Figure 28](#)) – registered a remarkable increase in 2018. Here, people came up with arguments regarding the reliability and efficiency of wind power, complaining about its inherent dependence on certain weather conditions. On the other hand, supporters of wind power countered with statistical causes of death of birds and insects where wind mills only play a small role. They argued that glass, housing infrastructure, cats or cars were more dangerous than wind turbines. Unfortunately, a lot of these tweets were misclassified as anti-wind although

they defend the technology. Still, such tweets were a minority that couldn't have compensated the negative ones. Surprisingly, there weren't many pro-wind power tweets referring to the technology's climate-friendliness. Those that contained climate keywords either made fun of ongoing projects that required deforestation (such as in the *Reinhardswald*) or referred to the positively perceived offshore wind parks. The latter were in focus in the early days of Twitter as part of the rather pro-wind topic 3 which saw a drastic decline over the years.



Figure 27: The word clouds for positive and negative tweets about wind energy.

Spring 2022: Twitter users are demanding more wind power

The term 'Bayern' (Bavaria) was a reoccurring theme found in wind tweets, predominantly included in topic 2. This topic saw a slight increase in spring 2022 and 'Bayern' gained a lot of attention as it was the most frequently mentioned term in March and April, prominently accompanied by 'Söder' and '10H'. Although the overall sentiment stayed below zero, people referring to Bavaria had positive sentiments about wind energy. The sudden rise in attention was found to be triggered by a political decision about the so-called *10H rule*. Introduced in 2014 by the Bavarian state, this rule specified the minimum distance of wind turbines to infrastructure which had to be at least ten times the height of the wind power station (Hehn & Miosga, 2015). Due to this rule, the approval of wind power projects has rapidly fallen by about 90% in Bavaria which was heavily criticised with regard to defined climate goals (Stede & May, 2019). It was in April 2022 when the Bavarian *Landtag* decided to relax this rule and defined a new distance of 1000 meters (BR24 Redaktion, 2022). However, most Twitter users were not satisfied with this decision as they wanted the rule to be fully cancelled. Thus, they still accused Markus Söder, the Bavarian governor, of curbing wind power development. Those users argued in favour of wind power as a means to mitigate climate change and a strategy to become independent of foreign fossil fuels. Consequently, 'Bayern' was primarily found in pro-wind power tweets (Figure 27). Especially due to the Russian invasion in Ukraine, microreading revealed a lot of Twitter users taking a stand for wind power instead of Russian gas. While most of these users saw more potential for wind power in Bavaria, a minority claimed that it was no 'wind state'. Moreover, it was remarkable how often the term 'Sonne' was mentioned within wind power tweets. A lot of people recognised similarities between wind and solar power and saw their combination as an optimal solution for the energy transition

while others argued against both. Such anti-wind tweets were mostly based on popular arguments, such as the intermittent energy supply due to its strong dependence on suitable wind speeds. Here, users reported how they saw a lot of wind turbines standing still most of the time. Especially in consideration of the energy crisis and environmental arguments such as deforestation, soil sealing or bird endangerment, this led to negative sentiments. Furthermore, an aesthetic aspect was brought into play as some people claimed these high wind turbines would ruin the visual depiction of the landscape. This was also observed by Rand and Hoen (2017) who found size and colour to be relevant for the assessment. Although anti-wind tweets were still more prevalent, the war-induced challenges had an initial positive impact on the perception of wind power. However, this effect declined towards the end of the year.

While the overall sentiment was still below zero in Germany, Swiss and Austrian users felt more positive about wind power in spring 2022 (Figure 29). The latter were less worried about environmental and visual implications but primarily argued in favour of wind turbines because they saw the importance of climate-friendly energy and independence. Like German users did with Bavaria, Austrian users pleaded for more wind power in potentially suitable regions. The same was observed for Swiss users who also referred to the energy crisis which was exacerbated by the Russian invasion. Hence, the claims for wind power expansions grew louder. Although people seemed to be aware of its limitations, they pointed out its strengths and perceived wind power as a piece to the puzzle towards an independent and climate-friendly energy supply strategy.




Figure 28: The average sentiment per wind topic of the whole study period.

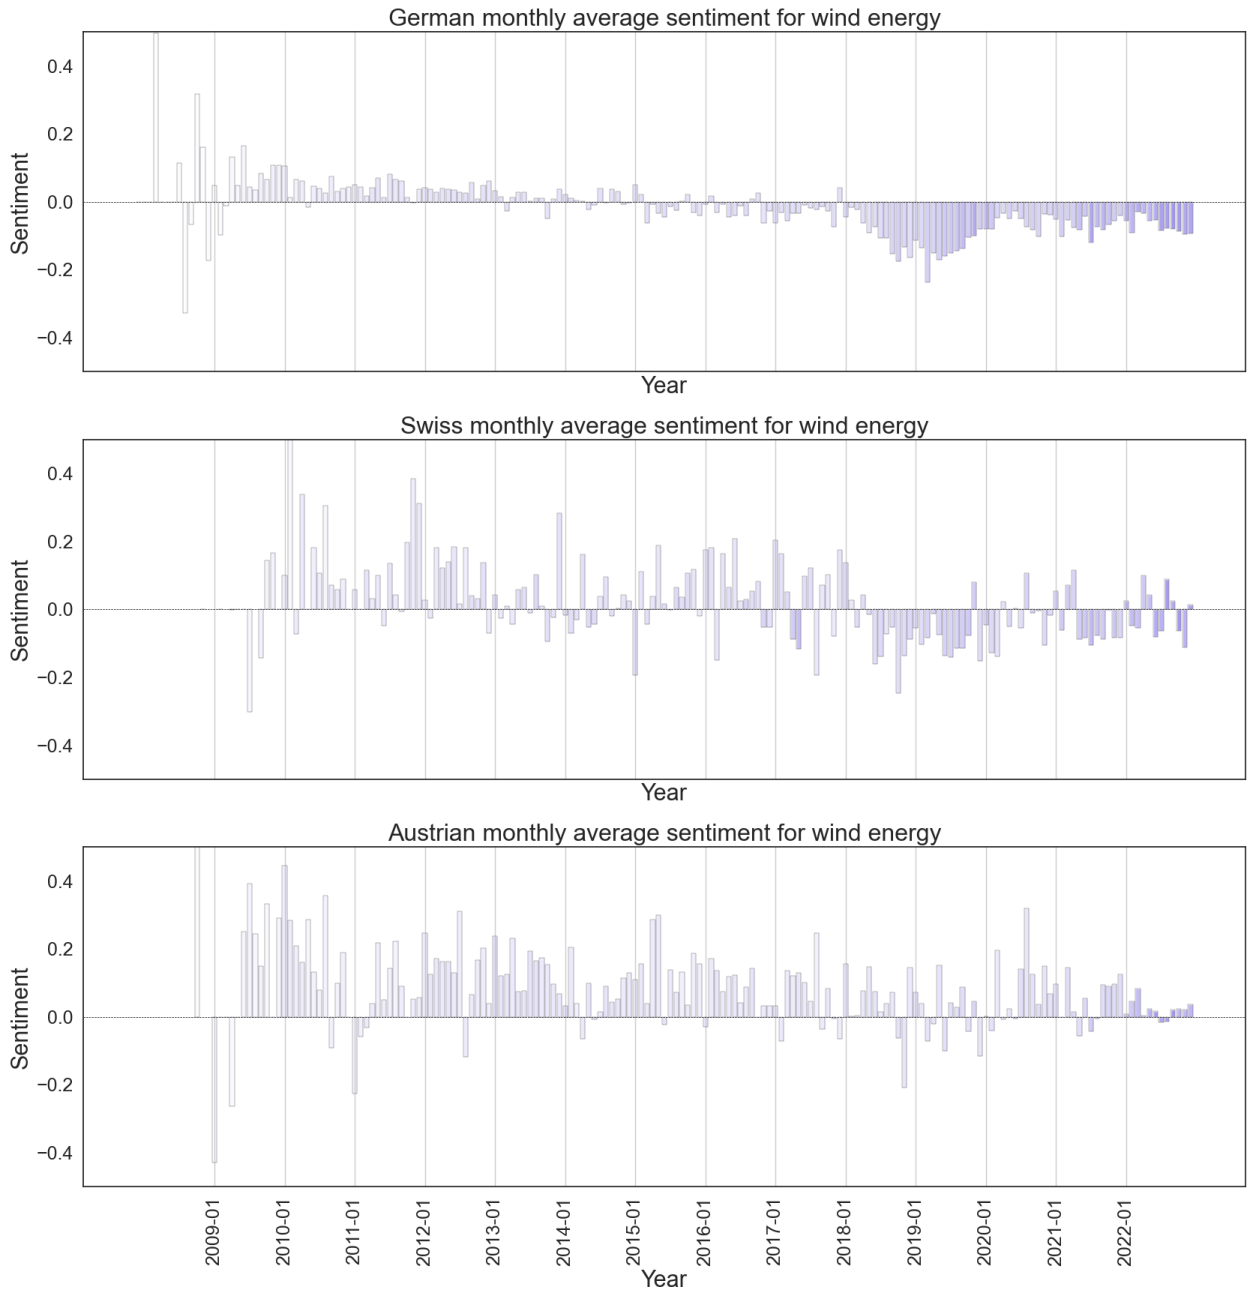


Figure 29: The country-specific sentiment timelines for wind energy.

5.7 Water energy

Compared to the other energy sources, the sentiment timeline of water energy is noisy and shaped by a large uncertainty due to the sparse amount of data. Only 44'869 tweets were left after pre-processing to derive the sentiment from which impedes the analysis. Most of the time, hydropower enjoyed a slightly positive sentiment expressed by Twitter users which was accompanied by many local peaks and lows (Figure 30). Over the whole study period, there were only a few months when the sentiment fell below zero, yet usually staying in the neutral range (Figure 31). In general, hydropower mainly enjoyed periods of positive sentiments in 2011/2012, 2015/2016 and 2021. In 2022, when tweet volume peaked, a negative trend was observed towards the end of the year. However, this didn't apply to Swiss users whose sentiment has risen during 2022 (Figure 34) and who were found to perceive hydropower more positively than their neighbours.




Figure 30: The sentiment timeline for water energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

Although the tweet volume seems to successively increase over time, data revealed that there was a slight unexpected tweet load surplus in May 2017 which was primarily visible in the sentiment timeline of Swiss users (Figure 34). Furthermore, interesting developments of topic shares were registered in spring 2018 (Figure 32). Consequently, reasons for these variations should be uncovered. Lastly, I will closely inspect tweets posted in December 2021 and August 2022 as those were examples of quite positives and neutral/slight negative sentiments, both characterised by a large tweet volume if compared to previous years.




Figure 31: The monthly average sentiment about water energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).

May 2017: Swiss users feel positive about hydropower ahead of a public vote about renewable energy support

In accordance to the country-specific sentiment timeline (Figure 34) and the slight peak of topic 8 (Figure 32), the term frequency analysis revealed that discussions in May 2017 were mainly due to a political event in Switzerland. In particular, users were discussing the public vote about the new *Energiegesetz (EnG)* as part of Switzerland's *Energiestrategie 2050*. This new law should further promote renewable energy sources and financially support large hydropower stations since they could hardly cover their costs due to the low electricity prices (Schweizerische Eidgenossenschaft, 2017). Most Swiss users were pleading in favour of the new proposal and, thus, also in favour of water energy. They argued that hydropower was an essential part of the renewable energy transition and could guarantee an independent energy supply. Some users even complained that the EnG wouldn't promote hydropower enough. Arguments against hydropower stations referred to its interventions into the natural environment or the price that couldn't compete with other energy sources. However, anti-hydropower voices were a small minority. German users, on the other hand, were neutral about hydropower as some of them claimed that can't be further developed in Germany while others referred to Switzerland's extensive hydropower infrastructure to demand a similar expansion in their country as well.




Figure 32: The temporal evolution of the ten different topics for water energy derived from the LDA model, slightly smoothed by a Gaussian filter.

Spring 2018: Support for hydropower in Switzerland and German fear of negative consequences for the natural ecosystem

In spring 2018, the sentiment about hydropower was predominantly positive but registered a slightly negative value in April. However, on a country level, the latter was only observed for German users (Figure 34). General discussions covered the *Swiss Federal Council's* decision of keeping the maximum water tax at the current level (Der Bundesrat, 2018) which wasn't well received among Twitter users who would have preferred less taxes to further strengthen hydropower. On the other hand, an impending dam burst at Colombia's largest storage power plant construction (*Hidroituango*) (see Bedoya and Cuellar (2018)) led to some negative voices regarding water power as it can pose great risk in case of unstable dams in combination with unexpected water masses. This resulted in an interim high of topic 6 which was generally the least positive topic derived from the LDA model (Figure 33). While the overall sentiment score still remained greater than zero, negative sentiments prevailed for German users. Reasons for this were different newspaper articles, some about thousands of planned hydropower stations in Greece and Slovenia which would threaten wild river ecosystems, others about the deadly consequences of hydropower stations for fish which was criticised by the *Hessian Fishing Association* (Süddeutsche Zeitung, 2018).

December 2021: Continuous Swiss support and reliability-arguments

In December 2021, the sentiments about hydropower were positive in all three countries. Users from Switzerland were primarily happy about the results of the *Runder Tisch Wasserkraft*, comprising different stakeholders including the Head of the *Federal Department of the Environment, Transport, Energy and Communications*, where all participants agreed on the importance of hydropower for the Swiss energy supply security and the need of its further expansion (EnDK, 2021). The large majority of Swiss Twitter users were convinced that additional support for hydropower was the right decision, calling it the most effective technology for a sustainable and secure energy supply. Again, German users were less positive about hydropower. While most people didn't dismiss the technology per se, they argued that an expansion wouldn't make sense due to the country's topography and its negative implications on river systems. On the opposite, Austrian users were well aware of the geographical advantages their country has for the usage of hydropower. In contrast to wind and solar power, people appreciated the reliability of hydropower, its continuous energy supply during night times and its opportunity to store energy in a simple way using pumped-storage power plants.




Figure 33: The average sentiment per water topic of the whole study period.

August 2022: Droughts are endangering hydropower generation

In August 2022, the sentiment about water energy, derived from a large number of tweets, fell below zero which hasn't been the case for almost two years. Quite specific words like 'Dürre', 'China' and 'Norwegen' were found to be part of anti-hydropower tweets. While energy generation via hydropower plants was usually appreciated since it's less dependent on weather conditions (as opposed to wind and solar power), heavy droughts in summer 2022 made people revise their statements regarding energy supply reliability. Most of the Twitter users referred to Norway and China, both countries with an intensive hydropower usage, as examples for the insecure energy supply of hydropower plants. Indeed, the massive drought in summer 2022 limited hydropower

generation in *Sichuan Province* which resulted in a serious power shortage as the province is heavily dependent on hydropower (X. Liu et al., 2023). Similarly, electricity production strongly declined in Norway, resulting in lowered energy exports and increasing electricity prices (ORF.at, 2022b). However, while some Twitter users exploited these situation to argue against hydropower, others drew increased attention on the risks of climate change and the need to support a climate measures, that include renewable energy sources. In combination with reoccurring doubts of German users regarding the expansive potential of hydropower in less mountainous regions, the sentiment resulted to be below zero, despite remaining in the neutral range.

Interestingly, the sentiment of Swiss users remained positive, even reaching up to 0.2 in the same month. Although some Swiss people also recognised the drought-induced limitations of hydropower, their support for the technology was preserved as long as there are temporary alternatives in cases of heavy droughts. Others claimed that further pumped-storage power plants need to be built, regardless of the latest events.

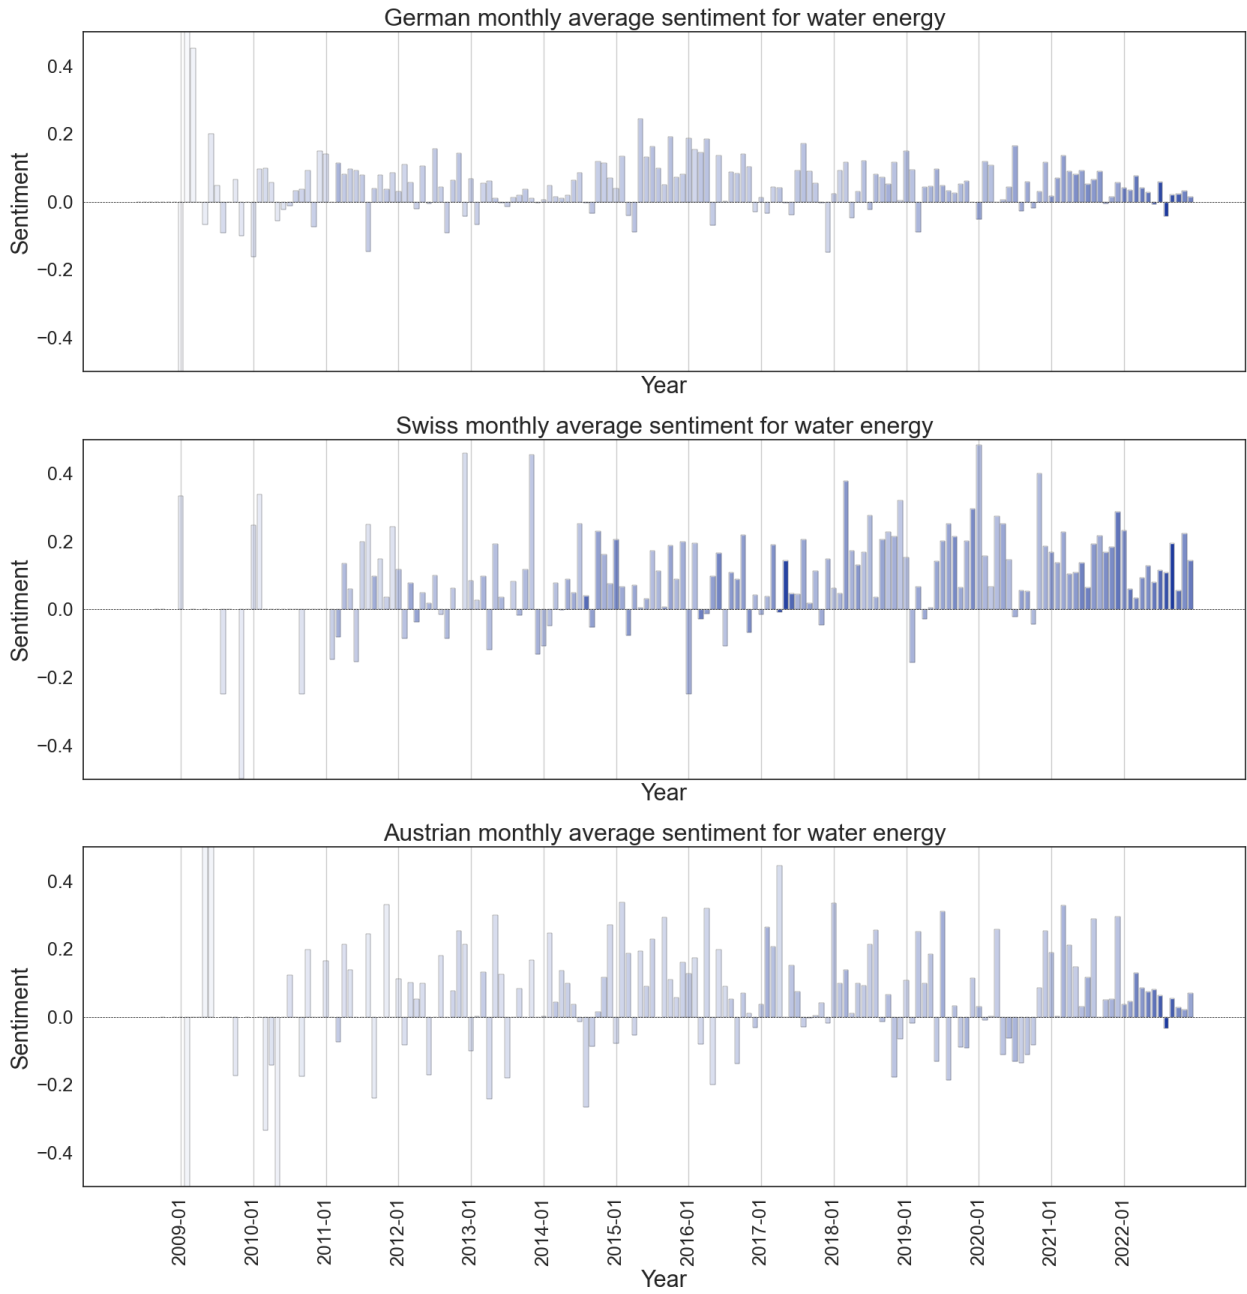


Figure 34: The country-specific sentiment timelines for water energy.

5.8 Gas/oil energy

Similar to wind energy, sentiment about gas and oil lightly oscillates in the neutral sphere with a minimal rather positive feeling between 2013 and 2019 before exiting this cycle to remain rather negative in the past few years. The most negative sentiment was reached in 2021 according to the smoothed curve of [Figure 35](#). The most positive value was either found in 2009 or late 2017, though rather remaining in the *neutral* scope. The reoccurring theme of decreasing uncertainty thanks to increasing tweet volume towards 2022 is shown as well. In the last year of the study period, the sentiment about gas and oil energy has risen again, approximately equalling the increase rate of that of wind energy. As per [Figure 36](#), the monthly average sentiment score shows a quite sharp stripline in January 2019 after which there has been only one single month (March 2020) when the sentiment value was above zero. While the slight positive sentiment period between 2013 and 2019 showed a minimal tendency of more positive sentiments during summer months, this pattern changed after 2019 when sentiments about gas and oil were more negative in summer, except for 2022. However, overall, sentiment values mostly lay in the neutral range, which also applied to 2022 when most tweets about gas and oil were posted.




Figure 35: The sentiment timeline for gas/oil energy, smoothed over 180 days with the buffer indicating the uncertainty based on the available tweet load taken to calculate the sentiment scores.

Based on the sentiment timeline for gas and oil energy in combination with the monthly tweet volume, I decided to specifically dedicate the content analysis to the time range between 2017 and the end of 2022 in order to find circumstances which could have been responsible for the sentiment changes. The chosen period can broadly be divided into three subperiods: 2017 and 2018 when the sentiment was still positive, 2019 until 2021 when the sentiment became more negative and 2021 onwards when it recovered to get close to zero again.




Figure 36: The monthly average sentiment about gas/oil energy between 2007 and 2023 with transparency indicating the tweet load per month (less transparency indicating more tweets).

2017 and 2018: Gas as the best alternative

In the first subperiod, the LDA topics 1 and 2 were dominant and the rise of topic 4 – reaching its climax in 2021 – started (Figure 37). The relevant terms explaining topic 1 mostly had to do with types of fuels whereas 'cng' was the most dominant one. This term referred to *Compressed Natural Gas* which is an alternate fuel for vehicles (Khan et al., 2015). Just like the slight positive sentiment of topic 1 (Figure 39) reveals, the large majority of users viewed CNG as the next big thing, the long-awaited alternative to petrol and diesel for the transportation sector, praising its environmental friendliness compared to the conventional fuels. Indeed, CNG is seen as a cleaner fuel than petrol or diesel, especially due to less emission output (Semin, 2008). In contrast to the developments in recent years, those users preferred CNG over electric vehicles, mainly due to the limited cruising range and the costly production of batteries. When closely inspecting topic 4, it was observed that a lot of people even viewed CNG as renewable ('erneuerbar') although this isn't correct as it's a form of fossil energy (Semin, 2008). While tweets of topic 4 were found to be in the *neutral* range (Figure 39), the word 'erneuerbar' was primarily found in posts which were positive about gas and oil energy (Figure 38). Moreover, in accordance to topic 2, a lot of users pointed out economic terms, predominantly complaining about petrol and diesel being expensive ('teuer') which further opened the gates for the auspicious CNG technology. Furthermore, it was feared that heating oil will reach its price peak soon, thus, the hashtag #peakoil was seen in many tweets. *Peak oil* refers to the theory of *Marion King Hubbert* who proposed that oil production cost rise after its peak was reached, leading to dramatic implications for economy and society (Bardi, 2009).




Figure 37: The temporal evolution of the ten different topics for gas/oil energy derived from the LDA model, slightly smoothed by a Gaussian filter.

2019 until 2021: Rising climate awareness and Nord Stream critics

The negative sentiment trend started in 2019 when the previously dominant 'CNG' topic 1 was overtaken by the 'erneuerbar' topic 4. This goes hand in hand with climate and environment related terms like 'co2' (often found in anti-gas tweets as per [Figure 38](#)) and 'fossil' suddenly appearing within the most frequently mentioned terms. While there were still tweets about gas as the desired alternate fuel, more people started to worry about nature and climate consequences due to the combustion of oil-based fossil fuels which led to more anti-gas and anti-oil tweets. In this context, words like 'Klimagerechtigkeit' or 'Klimakatastrophe' were mentioned. The first one hints to the *Climate justice movement* which – in Germany – was established in 2007 to address the historical responsibility of Northern hemisphere countries for the massive greenhouse gas emissions (Sander, 2016). However, negative climate impact was not the only reason for an increasing amount of anti-gas and anti-oil tweets. Other users posted tweets arguing against gas and oil criticizing the induced dependence on "oil countries" and the financial support of their cruel regimes. It was in late summer 2020 when topics 3 and 8 registered a sudden peak. Both prominently contain the word 'fracking', an unconventional shale gas extraction method enabled by horizontal drilling which predominately started booming in the US in 2013 (Jackson et al., 2014). As both topics contained tweets with a negative sentiment about gas/oil energy ([Figure 39](#)), the overall sentiment further decreased during their peaks. Microreading showed that users were complaining about fracking for being climate-unfriendly as it further contributes to greenhouse gas emissions (Staddon & Depledge,

2015). Moreover, the terms 'nord' and 'stream' were found to be very frequently mentioned during this topic 3 peak while the majority of tweets containing these terms showed a negative sentiment towards gas and oil energy. *Nord Stream* refers to a gas pipeline from Russia to Germany, announced in 2015 (De Jong, 2023). In fear of gas shortage and increasing gas demand due to the German coal phase-out, the project was supported by various central European countries (De Jong, 2023). On the contrary, it was harshly criticised by the United States of America (De Jong, 2023). Twitter users assumed that the critique arose as the US saw the cheap Russian gas a powerful competitor to their fracking gas exports. Users either criticised the fracking gas or the nord stream gas for its environmental risks or the rising dependence on global players, both leading to negative sentiments about the energy form.




Figure 38: The word clouds for positive and negative tweets about gas/oil energy.

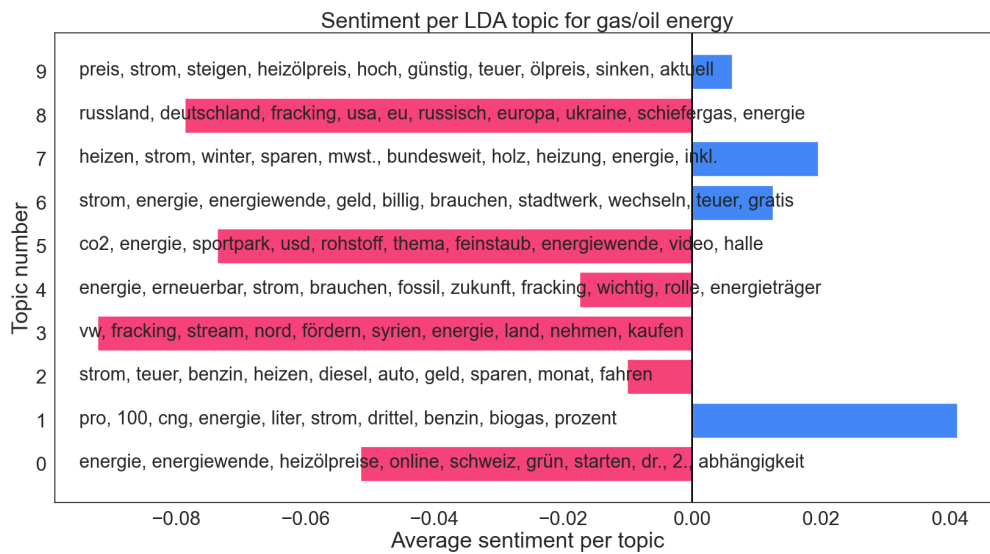


Figure 39: The average sentiment per gas/oil topic of the whole study period.

2021 and 2022: War and climate ethics versus affordable energy supply

While the sentiment about gas and oil was still negative until summer 2021, it started to recover to hit the neutral range afterwards, yet staying below zero. Congruously, the neutral/slight negative topic 2 started to dominate and the slightly positives topics 6 and 7 increased. However, as the

most frequent bi-gram 'russisch, gas' suggests, it was the Russian-Ukraine conflict in spring 2022 (Farghali et al., 2023) which led to a peak of topic 8, disrupting the sentiment recovery in the first half of the year. For many Twitter users, the Russian invasion intensified the urgency to get away from gas, become independent and stop to finance the Russian war. However, the peak flattened shortly after and topic 2 became the prevalent one again in the second half of 2022. In this subperiod, its most dominant terms were 'Strom', 'heizen' and 'teuer', which recorded an unprecedented surge. Microreading revealed that many users asked to get secure gas supply again as the rapid gas exit led to gas prices being six times higher than usually, making heating more expensive (Ruhnau et al., 2023). These voices also amplified the arguments of advocates for affordable energy prices against such alluding for sanctions against Russia and non-fossil energy sources. While a minority was still pointing out the necessity to replace gas and oil by renewable, climate-friendly energy sources, most users expressed their anger towards the government, claiming they weren't able to pay their heating and electricity costs during the upcoming winter months anymore. While the majority of the latter group seemed to be aware of the drawbacks of gas and oil, they found the EU's gas embargo (European Commission, 2022b) proving one thing: Under- and middleclass households simply can't finance their living expenses without an affordable energy supply. Consequently, personal destiny became paramount, leaving behind war and climate ethics. [Figure 40](#) uncovers the sentiment differences between users from Germany, Austria and Switzerland. Since the large majority of Twitter users was located in Germany, the general sentiment timeline strongly resembles the one for German users. On the contrary, sentiments about gas and oil stated by Swiss and Austrian users highly fluctuates and was rather negative in 2017 and 2018. However, sentiments of German and Swiss users were found to show the same tendencies since 2019 with Swiss users being more negative about gas and oil overall. As opposed to their neighbours, there were two months in 2020 and even 2021 when Austrian users expressed clearly positive sentiments about this energy source. When reading the respective tweets, no clear reason for this sudden outlier was found. However, it has to be noted that only 61 tweets were sent by Austrian users during the most positive month in June 2020.

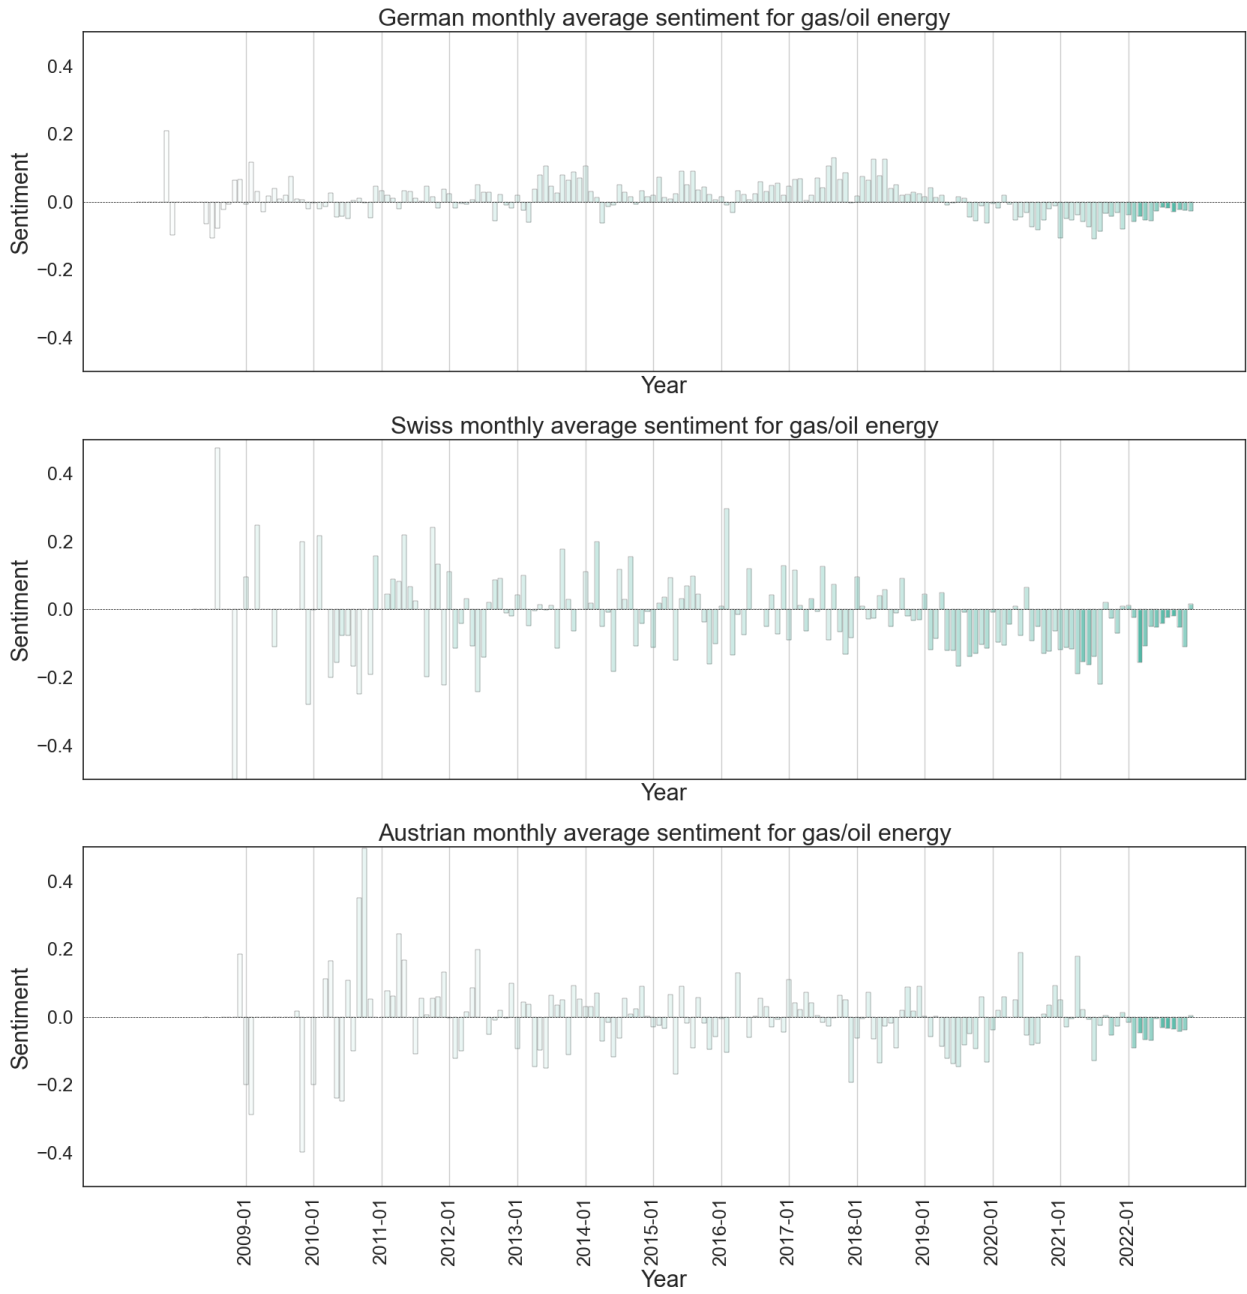


Figure 40: The country-specific sentiment timelines for gas/oil energy.

6 Discussion

In the first part of this chapter, the defined research questions are answered in summary by returning to the stated hypotheses. Then, the results presented in [section 5](#) will be linked together and connected to the existing literature. Afterwards, the circle to the introduction and background section will be closed by referring to existing theories and assessing implications for public support under the circumstances of the desired sustainable energy transition. In a final subsection, the chosen methods and data will be critically discussed by identifying limitations.

6.1 Main findings and evaluation of hypotheses

RQ1: How did the sentiments of German-speaking Twitter users towards different energy sources change between 2007 and 2023?

H1: *Due to the challenge of climate change, renewable energy sources and nuclear power gained popularity over the years while the sentiment about fossil fuels became more negative.*

As expected, the sentiment about nuclear power successively recovered after most users perceived it as very negative in 2011. However, a general positive trend of renewable energy sources was only registered for solar power which was interrupted in 2022. While the positive perception of hydropower mostly remained stagnant over the study period, wind energy went from slightly positive to controversial over the years and finally lost a lot of support starting in 2018. The sentiment about gas and oil showed a similar development as users felt undecided or even slightly positive about these energy sources while a scant negative trend started in 2020. In contrast to the expectations, coal energy did only lose support until 2019. Afterwards, people increasingly felt more positive about it than before. Although, it remained the most negatively perceived energy type.

RQ2: Which events or circumstances were responsible for these sentiment variations?

H2: *Energy availability, prices and projects, political decisions and popular movements, disasters or technological progression are driving forces behind sentiment variations.*

A variety of factors driving sentiment variations was identified. For instance, a dam burst in Colombia raised interim negative voices about hydropower, just like the attack on an Ukrainian nuclear power plant did with nuclear power. Although the disaster of Fukushima also had a negative impact on the perception of the latter, social anti-nuclear movements already led to a strongly negative sentiment of nuclear power prior to the catastrophe. Similarly, social climate movements starting in 2018 seemed to have an impact, especially when activists occupied the *Hambacher Forst* before its announced clearance. In combination with such long-term movements, distinct ventures or political decisions about projects could shape sentiments, for example observed when the construction of a new coal-fired power plant or its reactivation was announced or when the legendary *Reinhardswald* was to make way for a new wind park. Energy availability and price were especially found to be influential in 2022, caused by geopolitical tensions whose impact was underestimated.

The fear of energy scarcity ultimately led to weakened support for weather-dependent renewable energy sources. Moreover, it favoured geographically independent energy types. While new technologies were only found to have minor short-term effects (e.g. on solar power perception), the influence of media articles was unexpectedly large. Several articles led to an interim decline of solar power sentiments while those about birds endangered by wind turbines and incorrect news about an explosion near a French nuclear power plant had similar impacts on the respective energy sources.

RQ3: How did sentiments differ between Twitter users from Germany, Austria and Switzerland?

H3: Country-specific differences of sentiments mainly base on the respective energy mix of the country and national or regional events. Hence, more positive sentiments about hydropower are expected for Swiss and Austrian users while a more positive perception of coal power is expected for German users.

Overall, sentiments from German, Swiss and Austrian users were observed to be similar but heavily influenced by the numerical dominance of German Twitter users. While hydropower enjoyed clear support from Swiss users, Austrians viewed it as quite controversial between 2018 and 2020 which contradicts the hypothesis. Similarly, the expectations about coal power could not be vindicated as Germans viewed coal power as even more critical than their neighbours did. Hence, sentiment can't simply be inferred from the energy mix and a country's dependence on an energy source. As described in the previous subsection, national or regional events, however, were influential. Although, once regional events in Germany reached national sensation, they also influenced the online debate of Austrian and Swiss users. While a final explanation of observed differences is extremely difficult, Hornung (2023) claims that historic events and traditional convictions of social groups and political parties also play a role for energy perception.

6.2 The clash of interests

To progressively enable a widened perspective on the final implications of the thesis in the following chapters, the findings of the three research questions are interconnected and compared with results of existing literature, while the temporal focus lies on the periods of interests defined in the the result [section 5](#). Hence, slightly deviating from the previous structure, energy sources are grouped based on similar findings to allow a sight onto the *bigger picture*.

The basis for this *bigger picture* is shown in [Figure 41](#) which represents four superordinate groups of arguments (hereafter referred to as *dimensions*). They have been identified by summarising the main points Twitter users made to justify their energy perception. It was found that the large majority of users referred to at least one of these dimensions to express their positive or negative sentiment about the respective energy source. The dimensions correspond well to the advantages and disadvantages of energy sources described in [subsection 2.1](#) while no energy source can accomplish all four points: A secured, affordable and reliable energy supply, no negative health

implications, environmental protection and climate-friendliness. The following subsections will frequently return to these dimensions by uncovering dominant argumentation patterns and their temporal changes to derive underlying drivers of the sentiment variations.

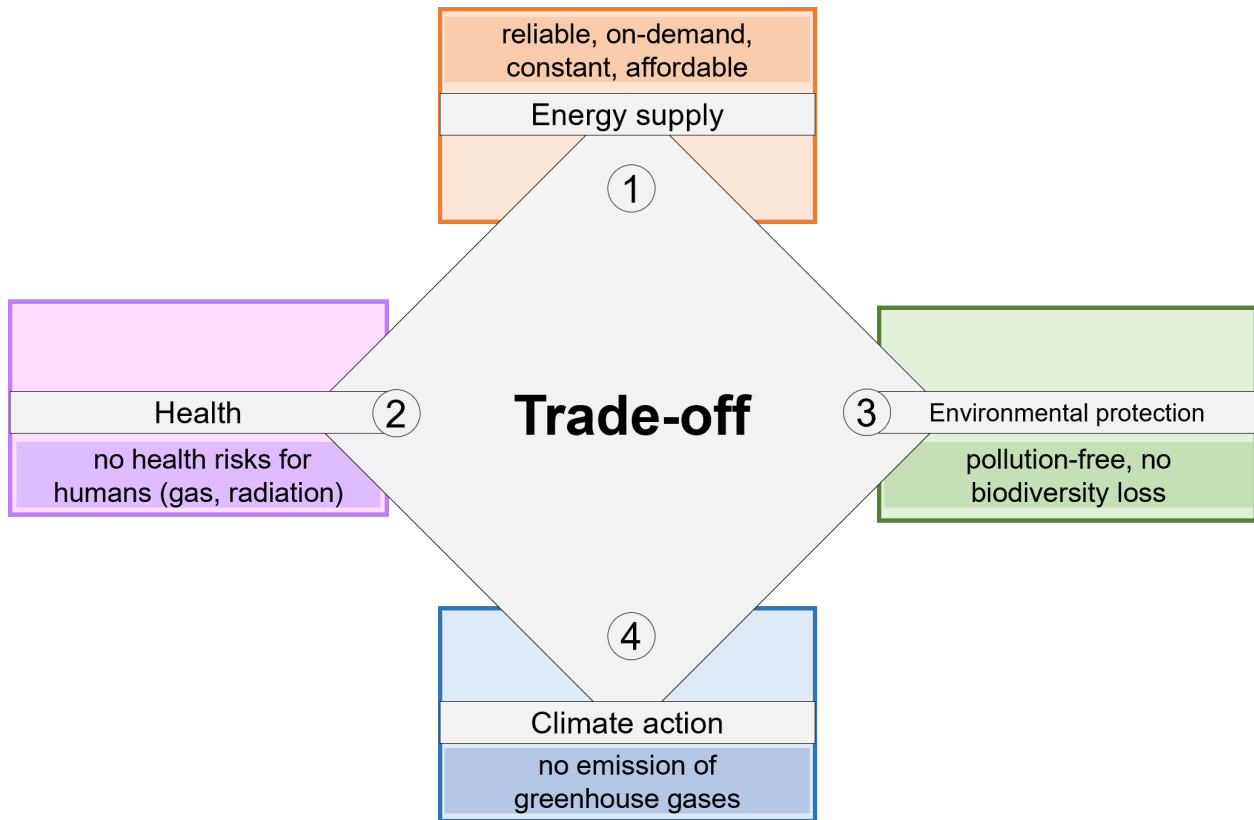


Figure 41: The four typical dimensions of arguments made by people to express their sentiment about energy types (with example keywords). Since no energy source can fully cover all four dimensions, a trade-off emerges.

6.2.1 Fossil fuels: Constant energy supply as the carrying argument

Although it was found that German-speaking Twitter users felt less negative about gas and oil than they did about coal, the argumentation mainly referred to the same dimensions shown in [Figure 41](#). Negative sentiments about fossil fuels were predominantly based on their implications for global climate, strengthening the call for contemporary fossil fuels phase-outs. Indeed, coal, gas and oil are major drivers the climate change (Ang et al., 2022), especially due to massive carbon dioxide emissions (Rahimnejad et al., 2012).

The difficult state of coal energy

In accordance to the results of Nuortimo and Härkönen (2018), coal energy was found to be dominated by negative sentiments. In comparison to gas/oil, German-speaking Twitter users further criticized its negative impact on the environment and human health, primarily evoked by coal

particles polluting the air which can cause respiratory disfunctions, cancer or cardiac diseases (Gasparotto & Martinello, 2021; Munawer, 2018). In accordance with the results of Markard et al. (2021) who found a strong and persistent anti-coal discourse coalition referring to climate impacts as the dominant delegitimation point of coal, my sentiment data have shown that these climate arguments largely outvoted the positive ones, at least until 2022. However, studies by Müller-Hansen et al. (2022) and Markard et al. (2023) observed an increasing pro coal phase-out tendency until 2020. This slightly deviates from my results which suggest the *Hambacher Forst* in 2018 to be the nadir of coal sentiment. However, the influence of the associated *WWF spam* (see section 5.4) was significant and sentiment remained very negative until 2020 which could hint this pro phase-out tendency. Though, Markard et al. (2023) took German newspaper data which doesn't seem to be directly comparable to Twitter data. Thematically, while Müller-Hansen et al. (2022) focused on the impact of the established *Coal Commission* that should plan the German coal phase-out, the authors found the coal debate on Twitter to be mainly influenced by this commission and social climate movements. While especially the movement *EndeGelaende* was prominently found to criticize the clearing of the *Hambacher Forst*, I haven't encountered the commission at all. This is probably due to Müller-Hansen et al. (2022)'s data acquisition method as they defined a separate keyword for the commission which artificially increased the recall of such tweets. Compared to the aforementioned studies, my thesis allows to enlighten the impact of the energy crisis intensified by the fear of energy shortage following the sanctions against Russia in 2022. Here, the findings coincide with those of Wiertz et al. (2023) who also found a shift of priorities away from climate arguments to energy supply security and even deduced from their results that "climate protection is no longer the clearly dominant goal of the energy transition" (Wiertz et al., 2023, p. 7). With reference to my sentiment results, I would not fully agree on this as the climate awareness was still a major argument of anti-coal tweets in 2022 which significantly contributed to the fact that sentiment about this energy source didn't reach positive values.

The war brings energy supply in focus

Still, coal energy's independence of environmental conditions as a undeniable strength for energy security (Roggenkamp et al., 2021) impressively demonstrated the pragmatic relevance of this dimension. The reliable energy supply argument was accompanied by the point of local availability that is especially given Germany's large coal reserves (Voss, 2022). Both arguments were also found to hinder a faster coal phase-out in studies by Markard et al. (2021) and Hermwille and Kiyar (2022). Consequently, the re-arrangement of weights towards coal's advantages led to an impressive revival of coal sentiment during the energy crisis. According to a survey conducted by *ARD* (German state television), the majority of participants from all parties except the Green Party voted in favour of a stronger usage of coal energy to become independent of Russian gas (infratest dimap, 2022). Still, even in 2022, the drawbacks of coal energy – negative climate, environmental and health implications – resulted in a slight negative perception of coal energy. However, the misclassification of tweets referring to China's extensive coal usage (see section 5.4) suggests that

the positive sentiment trend about coal, driven by the energy crisis, remained untamed. Since coal energy is not present in Switzerland and only a small part of the Austrian energy mix (Azarova et al., 2019; IEA, 2020), my data was strongly dominated by German events with the slight exception of Austria's coal power plant reactivation announcement (ORF.at, 2022a). As compared to Germany where the same decision strengthened pro-coal voices, this mainly evoked anti-coal Austrian users who reacted disappointed regarding the climate implications.

Gas and oil sentiment gets to feel the climate zeitgeist

While research lacks studies about the perception of oil and gas energy in Europe, a glimpse at the statistics proves their relevance, especially for heating systems. In 2022, almost 75% of all German residences were heated using gas (49.3%) and oil (24.7%) (BDEW, 2023). In Switzerland, 39.3% of all heating systems made use of oil and 17.5% of gas in 2022 (BfS, 2023) while these numbers were at approximately 12% (oil) and 19% (gas) for Austria in 2020 (Statistik Austria, 2021). Despite building on the same precondition as coal, namely the negative environmental and climatic impact of gas and oil combustion (Nicoletti et al., 2015), there were mainly two explanations for the neutral or slightly positive sentiments about gas/oil. First of all, the health implications were very rarely mentioned by Twitter users although the products emitted during oil combustion were found to increase childhood cancer risk (Knox, 2005). Secondly, especially *Compressed Natural Gas (CNG)* was found to be viewed as a valid alternative to petrol and diesel until 2019. As CNG emits less pollutants than these two oil derivatives (Semin, 2008), it was hyped in 2017 and 2018 when a lot of Twitter users praised its eco-friendliness. According to Singhal et al. (2017), environmental concerns regarding traditional fossil fuels led to the more frequent usage of CNG in the transportation sector. This corresponds to my results which further suggest that the fear of a drastic increase of diesel and petrol prices was a factor for the CNG hype as the latter was found to be less expensive than its competitors in the transportation sector (Khan et al., 2015). The shift away from a specific CNG focus towards the general perception of gas and oil as widespread energy sources – with intermin focus on the controversially discussed *fracking* practices and the Russian gas pipeline *Nord Stream* – temporally fell in a time period where climate awareness gained more attention which might have been influenced by climate movements like *FridaysForFuture* who emerged in central Europe in 2019 (X. Zhang, 2023). In consequence of this attention shift, the previously rather neutrally perceived climate and environmental dimensions increasingly became more negative.

The demand for independence of Russian gas and oil imports

The desire for an independent energy supply voiced during *Nord Stream* debates was intensified after the Russian invasion in Ukraine which corresponds to the findings of Wiertz et al. (2023). Their results suggest that the war has increased the urge of reduction of fossil energy usage, predominantly import-dependent ones like gas and oil among Germans who are strongly dependent on Russian

gas (Vrana et al., 2023; Wiertz et al., 2023). So, while the war boosted the support for coal energy, it had opposite initial effects on gas and oil. However, these primarily ethical concerns quickly vanished as Twitter user's demand for a secured gas supply increased after gas prices exploded (Ruhnau et al., 2023). This is consistent with the findings of Decker and Menrad (2015) who found economic aspects to be a crucial factor shaping the decision of heating system technologies. Similar tendencies were also observed in the political scape. As per Wiertz et al. (2023), even Green party members stated support for liquefied natural gas (LNG) from the US despite being climate-unfriendly. In Switzerland, Steffen and Patt (2022)'s survey found support for the ban of new gas and oil heating installations but only if an affordable alternative is ensured. Indeed, the federal council's suggestion of building several natural gas power plants to tackle possible power shortfalls in winter months also gained support after the war started (Steffen & Patt, 2022). While the data corresponds to the initial negative sentiment after the invasion and the recovering trend towards winter 2022, no explanation for Swiss users' sudden negative backslide in November 2022 was found.

Conclusively, just like for coal energy, the intensified energy crisis in 2022 led to a readjustment of weights assigned to the four dimensions in favour of energy supply at the expense of health, climate action and environmental protection.

6.2.2 Renewable energy sources: Positively perceived with one limiting factor

Solar energy, wind energy and hydropower are classic examples of renewable energy sources (Mohtasham, 2015). As found by Qazi et al. (2019) 17 of the 19 studies they retrieved from countries all around the world were found supportive towards renewable energy. However, despite their common characteristics, the sentiment of German-speaking Twitter users was found to vary between different types of renewable energy sources as per my results. A large portion of the variations can again be explained by the four dimensions shown in Figure 41. In a nationwide survey conducted by Sütterlin and Siegrist (2017) in 2012, Swiss people showed most support for solar power, followed by hydroelectric power and closely followed by wind power. This corresponds to my results, for Swiss people only as well as for the entirety of German-speaking Twitter users. While all three energy sources were praised for their climate-friendliness conditioned by their almost negligible emission output (Ang et al., 2022; Mardani et al., 2015; Mohtasham, 2015), people reported mixed feelings regarding the environmental dimension.

Climate-friendliness versus environmental implications

In the survey of Sütterlin and Siegrist (2017), most participants didn't associate solar power with any negative environmental impacts as they were not aware of the toxic waste formed during the solar panel production. Accordingly, I only remember having read about toxic waste in one single tweet which could explain the very positive sentiment observed for solar power throughout the study period. Similarly, Nuortimo et al. (2018) found Germans to not only be positive towards

solar power as an energy source but also towards respective policies and photovoltaic expansion. I'd suggest that the popularity of solar power also builds on the opportunity that individuals can start their own small energy projects by adding photovoltaic on their roofs.

Interestingly, Boudet (2019) described renewable energy sources to be perceived controversially due to "conflicting conservation priorities, pitting local harms to wildlife, landscape and so on against global benefits from reduced carbon emissions" (2019: 450). Here, wind energy was found to be a classic example of this clash of interests between environmental and climatic implications. As found in an Austrian study conducted by Scherhauser et al. (2017), especially conservationists and ecologists had strong negative views about wind energy and weren't found to be persuadable. They pointed out the noise emissions and the fatal danger for birds and bats crashing with rotor blades of the wind turbines. While noise emission was hardly noticed within my data, the latter matches my observations. However, even more common than birds, deforestation was used as anti-wind arguments, joined by worries about insects. Just like found by Nuortimo and Härkönen (2018), my data showed about as many negative as positive tweets about wind power in 2015/2016 before the perception of that energy source became significantly worse, driven by aforementioned environmental concerns massively covered in media (this strong attention of wind power in media was also observed by Nuortimo and Härkönen (2018)). Such an increasing negative trend could also be monitored by Dehler-Holland et al. (2022) who assessed changes in public sentiment about wind power in Germany between 2009 and 2018. The authors concluded: "Our results show that regional issues with health, environment, and landscapes have increased in prevalence over the past years, challenging wind power's legitimacy on normative grounds." (Dehler-Holland et al., 2022, p. 16). They also mention the geographical scale claiming that individual projects would conflict with societal values and regional planning laws (Dehler-Holland et al., 2022). Especially regarding the attention-grabbing worries about the legendary *Reinhardswald*, I found local arguments like the loss of a piece of culture or changing aesthetics to be brought into the discussion. The *Reinhardswald* is a classic example of a physical space that is associated with emotional meanings assigned by the local population which was often observed to cause resistance against energy projects (Devine-Wright, 2011). This corresponds to the findings of Sütterlin and Siegrist (2017), Müller et al. (2020) and Vuichard et al. (2019) who found Swiss people's high approval for renewable energy to clash with issues related to the acceptance of specific wind power projects on a local level. This is known as the *Not In My Backyard Phenomenon* which refers to people accepting energy projects as long as they're not directly affected by their drawbacks (Krohn & Damborg, 1999). Among others, Vuichard et al. (2019) view this phenomenon as a possible reason lessening public support of wind energy on a local level which I would agree on based on my microreading experiences of wind tweets and mostly positive descriptions of offshore plants.

My results of Swiss and Austrian users' positive sentiments towards hydropower were also documented in studies from Sütterlin and Siegrist (2017) and Klinglmair et al. (2015). Although Germans also felt rather positive about energy generation by water power, it could be Switzerland's and Austria's extensive usage of hydropower – in both countries, more than 60% of electricity is

generated by hydroelectric plants (Sütterlin & Siegrist, 2017; Wagner et al., 2015) – that led to higher approval rates than in Germany. Again, the findings of Venus et al. (2022) and my own suggest that the high positive sentiments of solar power couldn't be reached for any country due to hydropower's environmental implications. Venus et al. (2022)'s study found that Germans endorsed the climate-friendliness but at the same time stressed the importance of fish protection and river system conservation. In comparison to these findings, the fish argument only prominently popped up once in my data although I would have expected to see more people complaining about the negative consequences for the fish population. While Austrian users further pointed out the relevance of hydropower for jobs, they were found to be willing to pay extra for the expansion of hydropower but demanded for extra measures like fish protection (Klinglmair et al., 2015). In the same study, the *Not In My Backyard Phenomenon* was observed as well (Klinglmair et al., 2015). Such arguments were not found in my data. Though literature lacks long-term studies about the perception of hydropower, I explain the quite constant sentiment found in my data to be evoked by the lower dependence of hydropower on external factors (Dujardin et al., 2017) and the therefore evoked immunity to incisive events, the sophistication of the technology and its solidified pros and cons regarding the four argumentative dimensions (Figure 41).

The war and the renaissance of energy supply appreciation

Just like for fossil fuels, a very influential event for renewable energy perception was the Russian invasion in Ukraine in February 2022 which further emphasized the energy crisis (Farghali et al., 2023). However, it had a bidirectional impact. Immediately after the invasion, the increasing support for solar and wind power documented in the data was also observed by surveys of Steffen and Patt (2022) and Ochsenbein (2023) conducted in April and March 2022 respectively. In both surveys, Swiss people appreciated an intensified usage of renewable energy sources – predominantly solar and wind power – as a replacement for Russian gas and oil (Ochsenbein, 2023; Steffen & Patt, 2022). Furthermore, Steffen and Patt (2022) even found all political parties agreeing on this. Similar to coal power, a fifth argumentative dimension gained more attention as evermore people argued in favour of solar power as a private, local and economically feasible energy supply method that's independent of other nations. Though this argument was already noticed by Wolske et al. (2017), I found it to become more prevalent after the invasion. However, the quick turnaround of this development based on an inspection made by the salient study of Wiertz et al. (2023) who found the war to have conflicting impacts on the perception of renewable energy. While some Germans – as found for Swiss people – saw renewable energy sources as the solution to becoming independent of gas imports, others viewed renewable energy as the cause of the gas import dependency as they weren't yet a standalone option to cover all energy supply (Wiertz et al., 2023). In my data, renewable energy was increasingly confronted with critique of Twitter users building on that one dimension fossil fuels were praised for: Reliable energy supply. As much as renewable energy sources were approved for their climate-friendliness, their dependence on natural forces is also a major drawback since it leads to an unreliable, intermittent energy supply (Ang et al., 2022;

Azarpour et al., 2013; Behabtu et al., 2020; Leonard et al., 2020). Interestingly and in contrast to my findings, studies about renewable energy hardly found people who used this drawback for the argumentation. In my data, however, this energy supply dimension was especially used to argue against the installation of solar power and wind power plants due to their greater intermittency compared to hydropower (Dujardin et al., 2017). While the growing focus on the climate dimension and the emergence of climate movements might have been the reasons for the increasing positive trend of solar power over the years, the energy crisis finally caused the interruption of this trend. Just like for fossil fuels, the weights assigned to the different dimensions were revised in favour of energy supply.

6.2.3 Nuclear power: A special case

The results of the sentiments towards nuclear power match the findings of Paul Slovic's pioneering paper about risk perception (Slovic, 1987). Just one year after the devastating Chernobyl accidents (Mez, 2012), Slovic (1987) found lay people to perceive nuclear power as significantly more risky than experts did. The technology scored high values on the public perception of *dread risks* and *unknown risks*, meaning that its potential for fatal consequences and new, previously unknown risk was perceived as large (Boudet, 2019).

Fukushima's emphasize on the health dimension and its quick recovery

Just like the study by Kristiansen et al. (2018) in Switzerland and equivalent studies by Arlt and Wolling (2016) and Bernardi et al. (2018) conducted in Germany, I also found people strongly emphasizing the health risk of nuclear power after the Fukushima event in March 2011. So, it was that health dimension that dominated back then. Although, opposed to the Swiss *Angstbarometer* 2011 conducted by ENSI (2011), which found the perceived risk of an atomic contamination to be risen drastically compared to the previous year, the sentiment of Swiss Twitter users was partly already quite negative in 2010. However, this was only the case for four months characterised by a small tweet load. If comparing the yearly averages, 2011 registered a more negative sentiment overall while especially the difference between January/February 2011 and March 2011 was indeed enormous. Interestingly, such a drastic sentiment change was not the case for German users who had already very negative attitudes towards nuclear power in 2010 as well as in the first two months of 2011. This corresponds to the findings of Arlt and Wolling (2016) who also documented already negative perceptions towards nuclear power in 2010 (also found by Bernardi et al. (2018)) that further increased by a surprisingly small margin. This contradicts with observations made by Y. Kim et al. (2013) who found fundamental changes in public perception after the disaster, regardless of the previous level of acceptance. Moreover, when inspecting effects of the *Fukushima* disaster on public acceptance in 42 countries, the same authors found the public acceptance to decline more strongly for people living far away from Japan (distance effect) and to recover more quickly in countries with a greater energy production dependence on nuclear power (Y. Kim et al.,

2013). The former observations contradicts with my data as the *Marcoule* explosion (IRSN, 2011) in autumn 2011 caused even more concern and more negative sentiments than *Fukushima* did. The latter observation of Y. Kim et al. (2013), however, was also inspected in my data. Here, the negative sentiment of Twitter users from Austria (which has no nuclear power as per IEA (2020)) only slowly recovered from the disaster in 2011. On the other hand and in accordance to Kristiansen et al. (2018) and GFS Zürich (2012), the sentiment of Swiss people was found to recover quite quickly in the years after and even recorded an interim positive score in 2020 and 2021. In addition to a certain dependence on the energy source, the barely quantifiable health consequences in the aftermath of Fukushima could have been a reason for this. As Hasegawa et al. (2015, p. 482) wrote: "Notably, most injuries or illnesses were not related to radiation exposure."

Climate and energy crisis put focus on the advantages of nuclear power

The sentiment of German users also recovered after Fukushima, though by a slower rate than the Swiss one. For both countries, argumentative patterns throughout this long recovery period were similar and found their peak in the eventful year of 2022. In Germany, Hornung (2023) found the debate around nuclear energy to be shaped by diverging opinions between people who continuously viewed nuclear power as an unprecedented health risk and those who saw it as the optimal solution for the climate *and* the energy crisis at the same time. In accordance to these findings, these opposing arguments and an increasing weighting of the supply and climate dimensions as strengths of nuclear power were observed extremely frequently during microreading of tweets posted in 2022. Moreover, Wiertz et al. (2023) observed a less negative perception of nuclear power as more Germans further pointed out the possibility to become independent of Russian gas and oil. While I barely came across this point, the risk perception still outvoted the pro-nuclear sentiments. Consistently, Germans emphasized health and nuclear waste problems and didn't welcome the EU's declaration of nuclear power as *green energy* (European Commission, 2022a) as per Hornung (2023). Contradicting the still negative sentiments observed in 2022, when asked how a possible energy shortage following the Russian invasion should be solved, a majority of German and Swiss people saw nuclear power as a major solution as found by surveys of infratest dimap (2022) and Ochsenbein (2023). These results show the relevance of various circumstances and demonstrate the power of the energy supply dimension in crisis situations.

Consequently, nuclear power represents a special case. Compared to other energy sources, it was not characterized by a clash between the climate and the energy supply dimension as opposing argumentative spheres. Instead, it was predominantly the health argument that was responsible for the long ongoing scepticism which then clashed with energy supply and climate arguments during intensified energy crisis. Nonetheless, the health discussion flared up again for a short time when news about the attack on an Ukrainian nuclear power plant spread.

6.2.4 Implications for public support and the renewable energy transition

In this thesis, it was found that German-speaking Twitter users' sentiments about different energy sources varied over time. Reasoning patterns which were used to argue for or against an energy type primarily consisted of the four dimensions *energy supply*, *health*, *environmental protection* and *climate action* (Figure 41) as characteristic advantages and disadvantages of these energy sources. As described before, various events and circumstances led to changes in the weighting of these dimensions over the study period (*discourse shift*), which resulted in the observed sentiment variations. These temporally variable priorities were observed via microreading and assisted by term- and n-gram frequency analysis methods (see section 9 for a quantitative example).




Figure 42: The final sentiment timeline accompanied by the underlying dimensions and their temporally variable importance observed from the data. Solid lines represent a greater importance than dotted lines.

Figure 42 summarises the temporal development of the perceived importances assigned to the four dimensions by Twitter users. While environmental protection and human health (mainly regarding nuclear energy) dominated in the beginning of the study period, these dimensions were successively overlaid by prioritized climate and extraordinarily strong energy supply arguments due to the increasing awareness of the climate crisis and the challenges evoked by the intensified energy crisis in 2022. Hence the energy discourse not only shifted but got more complex due to this multidimensionality and the known trade-offs between the four dimensions. With the knowledge

of the dominant thematic priorities at a certain time (Figure 42) and the typical advantages and drawbacks of the six energy sources (described in subsection 2.1), the public sentiments of a specific energy source can be assessed.

For example, since both climate action and energy supply dimensions are advantages of nuclear power, its public sentiment increased in 2022 compared to previous years where disastrous events put emphasize on health and environmental protection which negatively influenced nuclear sentiment. Accordingly, the attack on the *Saporischschja* power plant put some additional weight on the health dimension again in 2022, which then led to a minor sentiment decline again.

For some time periods and energy sources, an additional *independence* dimension helps describing the recent sentiment trends of solar, wind and coal energy. While people put more weight on a geopolitically independent energy supply favouring solar, wind and coal after the outbreak of the war in spring 2022, this weight was later re-assigned to the supply dimension when fear of energy shortage arose. This again led to decreasing sentiments about the two renewable energy sources.

Public support for the sustainable energy transition exists but is challenged

While all studies about the public perception of energy sources were either limited by the amount of energy types investigated or a narrow study period, this thesis allowed the long-term sentiment observation and comparison of multiple energy sources. Therefore, protracted shifting processes like increased climate awareness as well as the impact of short-term events could be monitored, helping to estimate the effect of future energy projects, societal beliefs or specific incidents on the public perception of energy sources. Compared to longer time series like Dehler-Holland et al. (2022), my data also included the year 2022 and could therefore show the implications of incisive geopolitical events. This allowed to compare the impact of the Russian invasion to other events in order to assess the influence of different circumstances. Moreover, the thesis allowed further insights into country-specific sentiments to uncover impacts of local events or energy mixes. Hence, as no existing studies comparing energy perception of Germany, Austria and Switzerland were found, this thesis can be called up to suggest country-specific measures to policy-makers.

From the perspective of the Paris Agreement (United Nations, 2015) and the urgency of the sustainable energy transition (Fuso Nerini et al., 2019; IPCC, 2023; Markard et al., 2023), the latest positive sentiment trends of non-renewable energy sources appear threatening. The increasing prioritization of an on-demand, affordable and reliable energy supply was impressive, yet not completely unexpected. In a German survey in 2022, Siebel (2022) found a total shift of priorities as 48% of all participants named energy security as their main priority, while 40% chose low energy prices as the most important aspect. Only the remaining 12% ranked climate and environmental protection higher than the aforementioned issues.

On-demand and reliable energy supply is expected to remain a problem of renewable energy sources, mainly wind and solar power. The results of this thesis, with most people strongly criticizing the

dependence on weather conditions and the therefore induced intermittency, could be applied to Maslow (2000)'s famous *hierarchy of needs* pyramid, suggesting that people consider reliable energy supply as more important than implications on climate, nature or health. Coming back to Coyle and Simmons, 2014's quote in the introduction, research and development is challenged to further improve existing technologies and plan a suitable interplay of energy sources to tackle this problem in order to minimize the fear of power outage in our welfare-focused world.

While no data about the financial means of Twitter users exists, the findings of this thesis hint the *end of the world vs. end of the month* dichotomy mentioned by Martin and Islar (2021). During microreading, cost was often mentioned to justify a sentiment. Just like the findings of various studies (Lloyd & Nakamura, 2022; Sokołowski et al., 2023; Tatham & Peters, 2023), it is assumed that the financial situation of Twitter users is decisive for the prioritization of certain dimensions, with less wealthy persons to assign more weight on affordable energy supply than health, environment or climate. Understandably, they are more worried about how they're going to make it to the end of the month (financially) while people with enough financial means already think about the distant future, more worried about climate change implications. This would explain the results observed during the energy price burst in 2022 (Jayanti, 2022).

However, even for people with enough money, temporal effect are expected to play a role as well. Zhao and Luo (2021) found the *present bias* as a barrier to climate mitigation measures. Since a lot of people would "overvalue the immediate costs of climate mitigation policies (e.g., carbon tax) and undervalue the future benefits (e.g., greenhouse gas emission reductions)", they refuse to install energy-efficient devices due to higher upfront costs despite equally high or even higher future energy savings (Zhao & Luo, 2021, p. 3549).

Yet, the costs of renewable energy sources show a promising development as they have already registered a drastic decline in past years (Osman et al., 2023). At the same time, costs of non-renewable energy are projected to rise as they are expected to last another 200 years at maximum (Luderer et al., 2022; Okedu, 2018). This will again shift the discussion towards renewable-energy.

Despite these (for the energy transition) undesirable findings and the fact that a positive sentiment doesn't automatically mean public support for respective energy policies (Scharpf, 1999), it can't be drawn from my results, that the sustainable energy transition was unfeasible. Especially solar power and hydropower kept a broad public acceptance (on Twitter) even in 2022, predominantly in countries with already considerable usage of such. Most people were well aware of the huge advantages of these energy sources, even those who criticized them for their drawbacks. Since energy projects lacking public support are doomed to fail (Segreto et al., 2020; Sokołowski et al., 2023), policy-makers have to be careful in assessing appropriate measures as it was shown that too radical approaches like the wind park in the *Reinhardswald* can significantly alter public perception about an energy source. In all scenarios, it's crucial to improve the circumstances hindering public support for renewable energy, namely costs and supply security. Moreover, educative enlightenment of long term consequences of different energy projects must be continued to weaken the effects of

the *present bias* (Zhao & Luo, 2021) and the highly fluctuating individual risk perceptions (Slovic, 1987). Furthermore, governments are called upon to re-assess the role of nuclear energy as its global share will drastically increase in all modelled mitigation pathways aiming at the 1.5°C climate goal according to the IPCC (IPCC, 2018).

Still, energy sources will remain a pervasive dilemma that requires compromise and trade-offs until a magical source is found that continuously provides endless energy without having negative impacts on the environment, health and climate.

6.3 Methodological limitations and validity

The present thesis contains various sources of uncertainty. Since these uncertainties may sum up, the expressiveness of the results is impaired. Thus, major limitations of the three main tasks are explained while the validity of the results is critically assessed in the following sections. A detailed version of the limitations arisen from the data retrieval process and the various preprocessing steps can be found in the respective [section 3](#).

6.3.1 Data retrieval and preprocessing

The analysed data strongly depended on the keywords defined in the search queries. Due to ambiguity of some German words, the noisy nature of Tweets (Jackoway et al., 2011) like spelling mistakes or abbreviations and the implicit mentions of energy sources in comments and quotes, the precision and recall were impaired.

During the preprocessing phases which aimed to clean the data to increase the precision, multiple steps unleashed limitations. Due to the simple circumvention and the exclusion of accidental double creation tweets by the strict duplicates removal approach, the uncertain classification of Twitter users with little energy tweets by the predictive bot identification model and the omitted relevance classifier, preprocessing could not fully clean the data to exclude undesired tweets.

To tackle these limitations, the duplicate removal approach should be upgraded to at least consider the length of the tweet and the origin account. So, multiple users should be allowed to post the same tweet at least once, given that the tweet is not so long that accidental multiple postings are unlikely. To address the latter, some kind of similarity measures (such as introduced by Levenshtein et al. (1966)) could be defined. Moreover, an additional advanced relevance classifier could improve the precision. However, the usually short messages hamper this approach just like the cut access to the users' entire Twitter feeds impedes a more sophisticated bot detection. Nevertheless, Twitter policies seem to fluctuating, keeping the possibility of a less expensive research access alive.

6.3.2 Sentiment analysis

The fine-tuned aspect-based sentiment model was found to achieve an F1 score of 0.77, hence, misclassifying some tweets. I see mainly three reasons for these misclassifications.

Due to the time-consuming process (it took about 1 hour to label 180 tweets), it was decided to start the fine-tuning phase already after 6800 labelled samples. Due to the large number of tweets, this is a rather small proportion compared to S. Y. Kim et al. (2021), who used 9000 training samples but only had 266'686 tweets in total for their document-level sentiment analysis study of solar energy. In most studies, larger training sets were usually realized by multiple human annotators. Microreading revealed that there were still lots of scenarios to express a sentiment which weren't represented in the training dataset. Thus, the model had difficulties assigning the correct sentiment when applied on all 2.6 million tweets. Hence, more (diverse) training data would have been necessary to improve the model. Moreover, labelling the tweets wasn't always straightforward.

While some users made explicit statements, others used humour, sarcasm, news articles or the words of other people (politicians, experts, famous people) to express their sentiments in a more implicit way, leaving a great deal of room for interpretation. Furthermore, I caught myself being influenced by prejudices, pretending to know a certain sentiment of a tweet based on the user’s choice of words or the political undertone. Such assumptions could have led to inconsistent training data. Moreover, the sentiments between sarcastic statements couldn’t be considered although sarcasm is frequently chosen in political online discussions (B. Liu, 2012).

Lastly, the decision of the model checkpoint wasn’t optimal. Due to a missing parameter, the last epoch’s checkpoint was automatically chosen to represent the model. However, compared to checkpoint 3, this last checkpoint 4 had a slightly lower evaluation F1 score and a significantly larger evaluation loss (see section 9). Hence, the model became slightly worse at generalization, meaning it started to overfit Ying (2019). Thus, checkpoint 3 would have been the better choice as it managed to classify more tweets correctly while not being overconfident like checkpoint 4.

6.3.3 Content Analysis

The correct identification of important events which were responsible for prominent sentiment variations was mainly limited by the immense and heterogeneous number of tweets. While influential distinct events were easier to detect, especially long-term developments were challenging to uncover, even with support of LDA, term-frequency and n-gram analyses. Due to the 16 years time span in combination with the strictly probabilistic nature of LDA, the characteristic topic words were not universally explaining underlying discussions at any point in time. As implemented by Dehler-Holland et al. (2022), STMs could have been more suitable to deal with the time dimension. Moreover, semantic networks as created by Vrana et al. (2023) could have given more insights into the specific discourses. Lastly, only microreading could eventually allow a deep insight into the ongoing debates. The amount of data and its additional geospatial separation made it a really time-intensive task with a high likeliness of missing certain important events which could be (partly) responsible for sentiment variations.

6.3.4 Geospatial separation

Assigning Twitter users to one of the German-speaking countries was another challenging task, given the limited amount of available user data and the discontinued access on the Twitter API. Since only the energy tweets of a Twitter user could be considered to derive the home location from, the uncertainty of this approach was strongly dependent on the amount of tweets available per user. Hence, especially users without georeferenced tweets and only a few in-text toponym mentions were prone to be incorrectly assigned to a country. Especially attention-grabbing events in other countries induced misclassification. For example, the home location of users who only posted a few energy tweets whereas most of them mentioned *Fukushima* was assigned to Japan although this was highly doubtful. Thus, the underlying assumption of a tendency towards frequently men-

tioning toponyms of the home country was challenged by tweet scarcity. Moreover, 52% of users lacked georeferenced tweets or those containing toponyms and, thus, couldn't be included into the geospatial analysis. Here, access on the whole Twitter feed, network and context information as suggested by Zheng et al. (2018) would have been advantageous, especially since a large proportion of users actually mentioned their home town in the profile location or the profile description as it was found during the evaluation process. Furthermore, the manual correction of Nominatim's wrongly assigned country codes could not achieve completeness due to the immense amount of data. For instance, it was only uncovered after the whole process that words like 'Supergau' or 'Ukrainer' were misinterpreted as toponyms and then assigned to Austria because some objects in Austria contained these character sequences. However, since *Stanza* is a statistical learning approach using grammatical patterns to predict whether a word is a toponym or not, it misclassified some words as toponyms which *Nominatim* then tried to assign to respective coordinates. Since *Nominatim* was found to have extensive knowledge, most of these non-toponyms were actually assigned to a place. Additional errors probably happened due to Nominatim's importance-ranked disambiguation process. Here, again, access on the user's feed and profile could have helped identifying the respective place of interest.

Although the conformity between toponym-derived home locations and such inferred from georeferenced tweets reached more than 70%, the GPS data attached didn't necessarily had to represent the home country of a user. Yet, such errors are supposed to be minor and the comparison to the *true* user home locations manually derived from Twitter led to more accurate results. Lastly, it was predominantly the imbalanced dataset that aggravated the geospatial sentiment analysis process. Compared to Germans, Austrian and Swiss users only accounted for a small proportion of all tweets. Hence, the sentiment timeline of the latter two countries was inferred from less tweets, resulting in larger uncertainties, more severe effects of misclassified sentiments and, thus, larger limitations regarding results of RQ 3.

6.3.5 Validity

Irrespective of all methodological limitations, the data basis represents a decisive source of uncertainty. Although the research questions were specifically aimed at Twitter users, the thesis was intended to have a social benefit in terms of opinion research in order to show the level of public support which should help select appropriate, promising energy strategies. Just as with traditional surveys, however, representativeness poses a decisive challenge. The validity of the results must be treated with caution, as Twitter users are not representative of the general population (Barberá & Rivero, 2015; Mellon & Prosser, 2017; Wiertz et al., 2023). For example, it's supposed that younger people with interest in energy and politics are over-represented. Hence, sentiment on social media is not equal to public sentiment. This has been shown in 2011 when people of Berne voted in favour of *Mühleberg* despite the negative Swiss sentiment towards nuclear power. Furthermore, users are influenced by algorithms, exposing them to certain topics, shaping interaction patterns (Huszár et al., 2022). Based on that, while some individuals might change their sentiment about

an energy source over time, certain events could lead to some actor groups getting more active (e.g. by exploiting an event to justify their opinion) while others vanish. Although a daily sentiment cap per user was applied to avoid over-representation of such active users, the range of energy tweets per user remained large and allowed active users to contribute to the sentiment multiple times over the whole study period. So, the temporal resolution was an advantage and a drawback of this automated approach at the same time. Temporally repetitive surveys, on the other hand, don't allow such an imbalanced weighting of opinions. Hence, as Müller-Hansen et al. (2023) suggest, social media based studies should be supplemented by representative traditional surveys. However, due to the fairly high agreement with other study findings (see previous chapters), it's expected that at least long-term tendencies could be successfully inferred from the retrieved Twitter data.

7 Conclusion and further work

This thesis provides insights into the sentiments of German-speaking Twitter users about various energy sources (nuclear, coal, solar, wind, water, gas/oil) between 2007 and 2023. It further monitors these sentiments on a national level for German, Swiss and Austrian users. Compared to traditional surveys, this was achieved via a modern deep learning model that enabled a higher participant rate as well as a temporally extensive study period to uncover protracted variations.

In summary, it was found that several events, social movements, geopolitical incidents, political decisions and media focus led to long term shifts in priorities from health and environmental implications to climate effects and energy supply security which had impacts on the perception of energy sources. While sentiments about solar energy and hydropower remained positive throughout the study period, the negatively perceived coal and nuclear sources recently registered a strong positive trend, yet, remaining controversial. Opposed to the other renewable energy sources, wind power was an exception as it was viewed critically and lost attraction over the years, just like energy from gas and oil did.

Although the results are fraught with methodological difficulties and are not representative of society as a whole, they suggest that policy makers need to be cautious when proposing appropriate measures to accelerate the renewable energy transition, as renewable energy sources are not considered fully satisfactory and specific energy projects depend on various factors while all of them require public support. At the same time, limiting factors of renewable energy support need to be continuously addressed.

Based on the strong prioritization of an affordable and secure energy supply during recent crisis situations, it is assumed that – returning to Coyle and Simmons (2014)'s dramatic portrayal in the introduction – the impact of energy on our daily activities will remain indelible as it will continue to be everywhere and drive everything. The only question that remains is the price of this boundless imprint. And this ultimately seems to depend on the value the society outside of all social media networks places on factors apart from pure energy supply, namely health, nature and climate, among others.

Further work related to this thesis could focus fine-tuning a more sophisticated sentiment analysis model that allows the classification of tweets with multiple energy aspects as well as the extraction of aspect categories to quantitatively uncover argumentation patterns. Furthermore, the home location of Twitter users could be predicted on a more fine-grained scale to assess the impact of geospatial patterns like rurality or urbanity. Moreover, since the energy discourse was found to be heavily influenced by Germany, comparing the public support for energy sources of two or more countries with different languages and different energy strategies would be interesting as it could help predicting impacts of distinct energy projects and decisions to avoid reactions as seen during the *yellow vests* protests.

8 Bibliography

References

- Abdar, M., Basiri, M. E., Yin, J., Habibnezhad, M., Chi, G., Nemati, S., & Asadi, S. (2020). Energy choices in Alaska: Mining people's perception and attitudes from geotagged tweets. *Renewable and Sustainable Energy Reviews*, *124*, 109781.
- Acheson, E., & Purves, R. S. (2021). Extracting and modeling geographic information from scientific articles. *PLoS One*, *16*(1), e0244918.
- Ahlgren, M. (2023). 55+ Twitter statistics, facts & trends for 2023. Retrieved August 4, 2023, from <https://www.websiterating.com/research/twitter-statistics/>
- Ahmed, W., Bath, P. A., & Demartini, G. (2017). Using Twitter as a data source: An overview of ethical, legal, and methodological challenges. *The ethics of online research*, *2*, 79–107.
- Alothali, E., Zaki, N., Mohamed, E. A., & Alashwal, H. (2018). Detecting social bots on Twitter: A literature review. *2018 International conference on innovations in information technology (IIT)*, 175–180.
- Al-Rfou, R., Kulkarni, V., Perozzi, B., & Skiena, S. (2015). Polyglot-NER: Massive multilingual named entity recognition. *Proceedings of the 2015 SIAM International Conference on Data Mining*, 586–594.
- Amitay, E., Har'El, N., Sivan, R., & Soffer, A. (2004). Web-a-where: Geotagging web content. *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, 273–280.
- Amponsah, N. Y., Troldborg, M., Kington, B., Aalders, I., & Hough, R. L. (2014). Greenhouse gas emissions from renewable energy sources: A review of lifecycle considerations. *Renewable and Sustainable Energy Reviews*, *39*, 461–475.
- Ang, T.-Z., Salem, M., Kamarol, M., Das, H. S., Nazari, M. A., & Prabakaran, N. (2022). A comprehensive study of renewable energy sources: Classifications, challenges and suggestions. *Energy Strategy Reviews*, *43*, 100939.
- Arigo, D., Pagoto, S., Carter-Harris, L., Lillie, S. E., & Nebeker, C. (2018). Using social media for health research: Methodological and ethical considerations for recruitment and intervention delivery. *Digital health*, *4*, 2055207618771757.
- Arlt, D., & Wolling, J. (2016). Fukushima effects in Germany? Changes in media coverage and public opinion on nuclear power. *Public understanding of science*, *25*(7), 842–857.
- Arthur, R., Boulton, C. A., Shotton, H., & Williams, H. T. (2018). Social sensing of floods in the UK. *PloS one*, *13*(1), e0189327.
- Aßenmacher, M., Corvonato, A., & Heumann, C. (2021). Re-evaluating Germeval17 using German pre-trained language models. *arXiv preprint arXiv:2102.12330*.

- Attia, M., Samih, Y., Elkahky, A., & Kallmeyer, L. (2018). Multilingual Multi-class Sentiment Classification Using Convolutional Neural Networks. *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Aub, J. C., Evans, R. D., Hempelmann, L. H., & Martland, H. S. (1952). The late effects of internally-deposited radioactive materials in man. *Medicine*, *31*(3), 221–329.
- Azarova, V., Cohen, J., Friedl, C., & Reichl, J. (2019). Designing local renewable energy communities to increase social acceptance: Evidence from a choice experiment in Austria, Germany, Italy, and Switzerland. *Energy Policy*, *132*, 1176–1183.
- Azarpour, A., Suhaimi, S., Zahedi, G., & Bahadori, A. (2013). A review on the drawbacks of renewable energy as a promising energy source of the future. *Arabian Journal for Science and Engineering*, *38*, 317–328.
- Baba, M. (2013). Fukushima accident: what happened? *Radiation Measurements*, *55*, 17–21.
- Balet, T., Vo, Q., Salem, O., & Mehaoua, A. (2023). Cyberbullying Detection on tweets from Twitter using Machine Learning Algorithms. *2023 International Conference on Intelligent Computing, Communication, Networking and Services (ICCNS)*, 177–182.
- Barberá, P., & Rivero, G. (2015). Understanding the political representativeness of Twitter users. *Social Science Computer Review*, *33*(6), 712–729.
- Barbieri, F., Espinosa Anke, L., & Camacho-Collados, J. (2022). XLM-T: Multilingual Language Models in Twitter for Sentiment Analysis and Beyond. *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, 258–266.
- Bardi, U. (2009). Peak oil: The four stages of a new idea. *Energy*, *34*(3), 323–326.
- Barreto, R. A. (2018). Fossil fuels, alternative energy and economic growth. *Economic Modelling*, *75*, 196–220.
- Bashir, S., Bano, S., Shueb, S., Gul, S., Mir, A. A., Ashraf, R., Noor, N., et al. (2021). Twitter chirps for Syrian people: Sentiment analysis of tweets related to Syria Chemical Attack. *International Journal of Disaster Risk Reduction*, *62*, 102397.
- BDEW. (2023). Entwicklung der Beheizungsstruktur des Wohnungsbestandes in Deutschland. Retrieved October 20, 2023, from <https://www.bdew.de/service/daten-und-grafiken/beheizungsstruktur-wohnungsbestand/>
- Bedoya, N., & Cuellar, P. (2018). Colombia Hidroituango dam: 'The river took my house'. *BBC News*. Retrieved December 16, 2023, from <https://www.bbc.com/news/world-latin-america-44302566>
- Behabtu, H. A., Messagie, M., Coosemans, T., Berecibar, M., Anlay Fante, K., Kebede, A. A., & Mierlo, J. V. (2020). A review of energy storage technologies' application potentials in renewable energy sources grid integration. *Sustainability*, *12*(24), 10511.
- Benhardus, J., & Kalita, J. (2013). Streaming trend detection in twitter. *International Journal of Web Based Communities*, *9*(1), 122–139.

- Bernardi, L., Morales, L., Lühiste, M., & Bischof, D. (2018). The effects of the Fukushima disaster on nuclear energy debates and policies: a two-step comparative examination. *Environmental Politics*, 27(1), 42–68.
- Berner Zeitung. (2011, February). Berner sagen Ja zu neuem AKW in Mühleberg. Retrieved December 9, 2023, from <https://www.bernerzeitung.ch/berner-sagen-ja-zu-neuem-akw-in-muehleberg-926269284904>
- BfS. (2023). Energiebereich: Heizsysteme und Energiequellen. Retrieved January 10, 2024, from <https://www.bfs.admin.ch/bfs/de/home/statistiken/bau-wohnungswesen/gebaeude/energiebereich.html>
- Biesialska, K., Biesialska, M., & Rybinski, H. (2020). Sentiment analysis with contextual embeddings and self-attention. *International Symposium on Methodologies for Intelligent Systems*, 32–41.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan), 993–1022.
- Bos, S., & Chug, E. (1996). Using weight decay to optimize the generalization ability of a perceptron. *Proceedings of International Conference on Neural Networks (ICNN'96)*, 1, 241–246.
- Bot Sentinel Inc. (2022). Bot Sentinel - More than just bots... Retrieved August 12, 2023, from <https://botsentinel.com>
- Botometer. (2023). Botometer FAQ. Retrieved August 12, 2023, from <https://botometer.osome.iu.edu/faq>
- Boudet, H. S. (2019). Public perceptions of and responses to new energy technologies. *Nature energy*, 4(6), 446–455.
- BR24 Redaktion. (2022, April). 10H-Regel für Windräder bleibt – aber mit deutlichen Ausnahmen. Retrieved December 15, 2023, from <https://www.br.de/nachrichten/bayern/csu-ringt-um-lockerung-der-10h-regel-fuer-windraeder,T4DaMoV>
- Bruckner, T., Bashmakov, I. A., Mulugetta, Y., Chum, H., De la Vega Navarro, A., Edmonds, J., Faaij, A., Fungtammasan, B., Garg, A., Hertwich, E., et al. (2014). Energy systems. In *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group 3 to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- Bugden, D., Evensen, D., & Stedman, R. (2017). A drill by any other name: Social representations, framing, and legacies of natural resource extraction in the fracking industry. *Energy Research & Social Science*, 29, 62–71.
- Bundeskanzlei. (2023, December). Vorlage Nr. 608: Übersicht. Volksabstimmung vom 16.11.2023 "Für den geordneten Ausstieg aus der Atomenergie (Atomausstiegsinitiative)". Retrieved December 15, 2023, from <https://www.bk.admin.ch/ch/d/pore/va/20161127/det608.html>
- Bunting, A. (2004). Opposition to wind power: Can it be a catalyst for improving public understanding of energy usage. *Technologies, Publics and Power Conference*, 1–5.

- Cardis, E., Vrijheid, M., Blettner, M., Gilbert, E., Hakama, M., Hill, C., Howe, G., Kaldor, J., Muirhead, C., Schubauer-Berigan, M., et al. (2007). The 15-country collaborative study of cancer risk among radiation workers in the nuclear industry: estimates of radiation-related cancer risks. *Radiation research*, 167(4), 396–416.
- Carley, K. M., Malik, M., Landwehr, P. M., Pfeffer, J., & Kowalchuck, M. (2016). Crowd sourcing disaster management: The complex nature of Twitter usage in Padang Indonesia. *Saf. Sci.*, 90, 48–61.
- Castanho Silva, B., Wäckerle, J., & Wratil, C. (2022). *Determinants of public opinion support for a full embargo on Russian energy in Germany* (tech. rep.). ECONtribute Discussion Paper.
- Celik, A. N. (2021). Analysis of Energy Supply, Installed Power and Renewable Capacity in the World, the EU and Turkey. *Düzce Üniversitesi Bilim ve Teknoloji Dergisi*, 9(3), 500–519.
- Chalamkate, T., Tinnachote, C., Rutherford, A., et al. (2023). Developing Toponym Classification and Geolocation Estimation from Thai Tweets. *International Journal of Geoinformatics*, 19(7).
- Chen, X., Elmes, G., Ye, X., & Chang, J. (2016). Implementing a real-time Twitter-based system for resource dispatch in disaster management. *GeoJournal*, 81, 863–873.
- Chevalier, J.-M. (2009). The new energy crisis. In J.-M. Chevalier (Ed.), *The New Energy Crisis: Climate, Economics and Geopolitics* (pp. 6–59). Springer.
- Chu, Z., Gianvecchio, S., Wang, H., & Jajodia, S. (2012). Detecting automation of twitter accounts: Are you a human, bot, or cyborg? *IEEE Transactions on dependable and secure computing*, 9(6), 811–824.
- Chu, Z., Widjaja, I., & Wang, H. (2012). Detecting social spam campaigns on twitter. *Applied Cryptography and Network Security: 10th International Conference, ACNS 2012, Singapore, June 26-29, 2012. Proceedings 10*, 455–472.
- Chyong, C. K., Corbeau, A.-S., Joseph, I. B., & Mitrova, T. (2023). Future Options for Russian Gas Exports. *Center on Global Energy Policy*, 19.
- Cieliebak, M., Deriu, J. M., Egger, D., & Uzdilli, F. (2017). A twitter corpus and benchmark resources for german sentiment analysis. *5th International Workshop on Natural Language Processing for Social Media, Boston MA, USA, 11 December 2017*, 45–51.
- Compton, R., Jurgens, D., & Allen, D. (2014). Geotagging one hundred million twitter accounts with total variation minimization. *2014 IEEE international conference on Big data*, 393–401.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Cotton, M., & Devine-Wright, P. (2013). Putting pylons into place: a UK case study of public perspectives on the impacts of high voltage overhead transmission lines. *Journal of Environmental Planning and Management*, 56(8), 1225–1245.

- Coyle, E. D., & Simmons, R. A. (2014). *Understanding the global energy crisis*. Purdue University Press.
- Dai, J., Li, S., Bi, J., & Ma, Z. (2019). The health risk-benefit feasibility of nuclear power development. *Journal of cleaner production*, *224*, 198–206.
- Dave, K., Lawrence, S., & Pennock, D. M. (2003). Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. *Proceedings of the 12th international conference on World Wide Web*, 519–528.
- Davis, C. A., Varol, O., Ferrara, E., Flammini, A., & Menczer, F. (2016). Botornot: A system to evaluate social bots. *Proceedings of the 25th international conference companion on world wide web*, 273–274.
- dbmdz. (2022). dbmdz German BERT models. Retrieved July 21, 2023, from <https://huggingface.co/dbmdz/bert-base-german-uncased>
- De Bruijn, J., de Moel, H., Jongman, B., Wagemaker, J., & Aerts, J. C. (2017). TAGGS: Grouping tweets to improve global geotagging for disaster response. *Natural Hazards and Earth System Sciences Discussions*, 1–22.
- De Greve, L., Martens, G., Lefever, E., Singh, P., & Van Hee, C. (2022). Aspect-based Sentiment Analysis for German: Analyzing “Talk of Literature” Surrounding Literary Prizes on Social Media. *Computational Linguistics in the Netherlands Journal*, *11*, 85–104.
- De Jong, M. (2023). Tracing the downfall of the Nord Stream 2 gas pipeline. *Wiley Interdisciplinary Reviews: Energy and Environment*, e502.
- Dean, B. (2022, January). How Many People Use Twitter in 2022? [New Twitter Stats]. Retrieved March 2, 2023, from <https://backlinko.com/twitter-users#daily-active-users>
- Decker, T., & Menrad, K. (2015). House owners' perceptions and factors influencing their choice of specific heating systems in Germany. *Energy Policy*, *85*, 150–161.
- Dehler-Holland, J., Okoh, M., & Keles, D. (2022). Assessing technology legitimacy with topic models and sentiment analysis—The case of wind power in Germany. *Technological Forecasting and Social Change*, *175*, 121354.
- DeLozier, G., Baldridge, J., & London, L. (2015). Gazetteer-independent toponym resolution using geographic word profiles. *Proceedings of the AAAI Conference on Artificial Intelligence*, *29*(1).
- Der Bundesrat. (2018, May). Bundesrat schlägt Beibehaltung des Wasserzinsmaximums bis 2024 vor. Retrieved December 15, 2023, from <https://www.admin.ch/gov/de/start/dokumentation/medienmitteilungen.msg-id-70763.html>
- Der Bundestag. (2022). Bundestag beschließt AKW-Laufzeitverlängerung bis Mitte April 2023. Retrieved December 10, 2023, from <https://www.bundestag.de/dokumente/textarchiv/2022/kw45-de-atomgesetz-freitag-917474>
- Devine-Wright, P. (2011). Place attachment and public acceptance of renewable energy: A tidal energy case study. *Journal of Environmental Psychology*, *31*(4), 336–343.

- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dickerson, J. P., Kagan, V., & Subrahmanian, V. (2014). Using sentiment to detect bots on twitter: Are humans more opinionated than bots? *2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014)*, 620–627.
- Die Bundesregierung. (2023). Robert Habeck: Bundesminister für Wirtschaft und Klimaschutz. Retrieved December 11, 2023, from <https://www.bundesregierung.de/breg-de/bundesregierung/bundeskabinett/lebenslauf-1974184>
- Dreyer, S. J., Polis, H. J., & Jenkins, L. D. (2017). Changing Tides: Acceptability, support, and perceptions of tidal energy in the United States. *Energy Research & Social Science*, 29, 72–83.
- Dujardin, J., Kahl, A., Kruyt, B., Bartlett, S., & Lehning, M. (2017). Interplay between photovoltaic, wind energy and storage hydropower in a fully renewable Switzerland. *Energy*, 135, 513–525.
- Edry, T., Maani, N., Sykora, M., Elayan, S., Hswen, Y., Wolf, M., Rinaldi, F., Galea, S., & Gruebner, O. (2021). Real-time geospatial surveillance of localized emotional stress responses to COVID-19: a proof of concept analysis. *Health & Place*, 70, 102598.
- El-Kishky, A., Song, Y., Wang, C., Voss, C., & Han, J. (2014). Scalable topical phrase mining from text corpora. *arXiv preprint arXiv:1406.6312*.
- EnDK. (2021). Medienmitteilung vom 13. Dezember 2021 zum Runden Tisch Wasserkraft: Wichtiges Bekenntnis zur Wasserkraft. *Medienmitteilung Konferenz kantonaler Energiedirektoren (ENDK)*. Retrieved December 16, 2023, from https://www.endk.ch/de/dokumentation/aktuelles/medienmitteilung-endk_bpuk_rkgk-wichtiges-bekenntnis-zur-wasserkraft
- ENSI. (2011). Angstbarometer 2011: Die Gefahr einer radioaktiven Verseuchung in der Schweiz ist klein. *Eidgenössisches Nuklearsicherheitsinspektorat*. Retrieved November 2, 2023, from <https://www.ensi.ch/de/2011/11/30/angstbarometer-die-gefahr-einer-radioaktiven-verseuchung-in-der-schweiz-ist-klein/>
- European Commission. (2022a). AEU Taxonomy: Commission welcomes the result of today's vote by the European Parliament on the Complementary Delegated Act. *European Commission Press Release*. Retrieved December 10, 2023, from https://ec.europa.eu/commission/presscorner/detail/en/ip_22_4349
- European Commission. (2022b). REPowerEU: A plan to rapidly reduce dependence on Russian fossil fuels and fast forward the green transition. *European Commission Press Release*. Retrieved April 16, 2023, from https://ec.europa.eu/commission/presscorner/detail/en/IP_22_3131
- Farghali, M., Osman, A. I., Mohamed, I. M., Chen, Z., Chen, L., Ihara, I., Yap, P.-S., & Rooney, D. W. (2023). Strategies to save energy in the context of the energy crisis: a review. *Environmental Chemistry Letters*, 1–37.

- Feizollah, A., Ainin, S., Anuar, N. B., Abdullah, N. A. B., & Hazim, M. (2019). Halal products on Twitter: Data extraction and sentiment analysis using stack of deep learning algorithms. *IEEE Access*, 7, 83354–83362.
- Ferraz Costa, A., Yamaguchi, Y., Juci Machado Traina, A., Traina Jr, C., & Faloutsos, C. (2015). Rsc: Mining and modeling temporal activity in social media. *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 269–278.
- Ferroni, F., & Reichmuth, A. (2017). Die verheerende Bilanz von Solarenergie. *Basler Zeitung online*. Retrieved December 11, 2023, from <https://www.bazonline.ch/die-verheerende-bilanz-von-solarenergie-454156025225>
- Finkel, J. R., Grenager, T., & Manning, C. D. (2005). Incorporating non-local information into information extraction systems by gibbs sampling. *Proceedings of the 43rd annual meeting of the association for computational linguistics (ACL'05)*, 363–370.
- Fuso Nerini, F., Sovacool, B., Hughes, N., Cozzi, L., Cosgrave, E., Howells, M., Tavoni, M., Tomei, J., Zerriffi, H., & Milligan, B. (2019). Connecting climate action with other Sustainable Development Goals. *Nature Sustainability*, 2(8), 674–680.
- Gallwitz, F., & Kreil, M. (2022). Investigating the validity of botometer-based social bot studies. *Multidisciplinary International Symposium on Disinformation in Open Online Media*, 63–78.
- Garske, S. I., Elayan, S., Sykora, M., Edry, T., Grabenhenrich, L. B., Galea, S., Lowe, S. R., & Gruebner, O. (2021). Space-time dependence of emotions on Twitter after a natural disaster. *International journal of environmental research and public health*, 18(10), 5292.
- Gasparotto, J., & Martinello, K. D. B. (2021). Coal as an energy source and its impacts on human health. *Energy Geoscience*, 2(2), 113–120.
- Geels, F. W., & Verhees, B. (2011). Cultural legitimacy and framing struggles in innovation journeys: A cultural-performative perspective and a case study of Dutch nuclear energy (1945–1986). *Technological Forecasting and Social Change*, 78(6), 910–930.
- Gelernter, J., & Balaji, S. (2013). An algorithm for local geoparsing of microtext. *GeoInformatica*, 17, 635–667.
- GFS Zürich. (2012). Angstbarometer 2012: Fukushima-Effekt schon fast wieder verpufft. Retrieved October 18, 2023, from <https://gfs-zh.ch/wp-content/uploads/2015/02/Angstbarometer-2012.pdf>
- Ghahremanlou, L., Sherchan, W., & Thom, J. A. (2015). Geotagging twitter messages in crisis management. *The Computer Journal*, 58(9), 1937–1954.
- Go, A., Huang, L., Bhayani, R., et al. (2009). Twitter sentiment analysis. *Entropy*, 17, 252.
- Götz, U. (2011, February). Proteste gegen Castor Transporte. *Deutschlandfunk*. Retrieved December 9, 2023, from <https://www.deutschlandfunk.de/proteste-gegen-castor-transporte-100.html>
- Grace, R. (2021). Toponym usage in social media in emergencies. *International Journal of Disaster Risk Reduction*, 52, 101923.

- Greer, C. F., & Ferguson, D. A. (2011). Using Twitter for promotion and branding: A content analysis of local television Twitter sites. *Journal of Broadcasting & Electronic Media*, 55(2), 198–214.
- Gritta, M., Pilehvar, M. T., & Collier, N. (2020). A pragmatic guide to geoparsing evaluation: Toponyms, Named Entity Recognition and pragmatics. *Language resources and evaluation*, 54, 683–712.
- Gritta, M., Pilehvar, M. T., Limsopatham, N., & Collier, N. (2018). What's missing in geographical parsing? *Language Resources and Evaluation*, 52, 603–623.
- Grover, C., Tobin, R., Byrne, K., Woollard, M., Reid, J., Dunn, S., & Ball, J. (2010). Use of the Edinburgh geoparser for georeferencing digitized historical collections. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368, 3875–3889.
- Guhr, O., Schumann, A.-K., Bahrmann, F., & Böhme, H. J. (2020). Training a Broad-Coverage German Sentiment Classification Model for Dialog Systems. *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 1627–1632.
- Habdank, M., Rodehutsors, N., & Koch, R. (2017). Relevancy assessment of tweets using supervised learning techniques: Mining emergency related tweets for automated relevancy classification. *2017 4th International conference on information and communication technologies for disaster management (ICT-DM)*, 1–8.
- Hahmann, S., Purves, R., & Burghardt, D. (2014). Twitter location (sometimes) matters: Exploring the relationship between georeferenced tweet content and nearby feature classes. *Journal of Spatial Information Science*, (9), 1–36.
- Hambrick, M. E., Simmons, J. M., Greenhalgh, G. P., & Greenwell, T. C. (2010). Understanding professional athletes' use of Twitter: A content analysis of athlete tweets. *International Journal of Sport Communication*, 3(4), 454–471.
- Han, B., Cook, P., & Baldwin, T. (2014). Text-based twitter user geolocation prediction. *Journal of Artificial Intelligence Research*, 49, 451–500.
- Hasegawa, A., Tanigawa, K., Ohtsuru, A., Yabe, H., Maeda, M., Shigemura, J., Ohira, T., Tominaga, T., Akashi, M., Hirohashi, N., et al. (2015). Health effects of radiation and other health problems in the aftermath of nuclear accidents, with an emphasis on Fukushima. *The Lancet*, 386(9992), 479–488.
- Hecht, B., Hong, L., Suh, B., & Chi, E. H. (2011). Tweets from Justin Bieber's heart: the dynamics of the location field in user profiles. *Proceedings of the SIGCHI conference on human factors in computing systems*, 237–246.
- Hehn, N., & Miosga, M. (2015). Die Zukunft der Windenergie in Bayern nach Einführung der 10 H-Regel. *Informationen zur Raumentwicklung*, 6, 631–644.
- Hermwille, L., & Kiyar, D. (2022). Late and expensive: The political economy of coal phase-out in Germany. In M. Jakob & J. C. Steckel (Eds.), *The Political Economy of Coal* (pp. 21–39). Routledge.

- Hiemstra, D. (2000). A probabilistic justification for using tf x df term weighting in information retrieval. *International Journal on Digital Libraries*, 3, 131–139.
- Hoang, M., Bihorac, O. A., & Rouces, J. (2019). Aspect-based sentiment analysis using bert. *Proceedings of the 22nd nordic conference on computational linguistics*, 187–196.
- Holechek, J. L., Geli, H. M., Sawallah, M. N., & Valdez, R. (2022). A global assessment: can renewable energy replace fossil fuels by 2050? *Sustainability*, 14(8), 4792.
- Holt, R. (1999). A responsible energy future. *Science*, 285(5428), 662–662.
- Hornung, J. (2023). Social identities and deadlocked debates on nuclear energy policy. *Journal of European Public Policy*, 1–25.
- Hu, X., Zhou, Z., Li, H., Hu, Y., Gu, F., Kersten, J., Fan, H., & Klan, F. (2022). Location reference recognition from texts: A survey and comparison. *arXiv preprint arXiv:2207.01683*.
- Hu, Y., & Adams, B. (2020). Harvesting big geospatial data from natural language texts. In M. Werner & Y.-Y. Chiang (Eds.), *Handbook of Big Geospatial Data* (pp. 487–507). Springer.
- Hugging Face. (2020). Transformers Trainer class. Retrieved July 23, 2023, from https://huggingface.co/transformers/v3.0.2/main_classes/trainer.html
- Hugging Face. (2021). Fine-tuning a pretrained model. Retrieved July 28, 2023, from <https://huggingface.co/docs/transformers/training#train-with-pytorch-trainer>
- Hugging Face. (2023). Models. Retrieved July 21, 2023, from <https://huggingface.co/models>
- Huszár, F., Ktena, S. I., O'Brien, C., Belli, L., Schlaikjer, A., & Hardt, M. (2022). Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences*, 119(1), e2025334119.
- Hutchinson, A. (2021). New Study Shows That 25% of Twitter Users Produce 97% of All Tweets. *Social Media Today*. Retrieved May 11, 2023, from <https://www.socialmediatoday.com/news/new-study-shows-that-25-of-twitter-users-produce-97-of-all-tweets/610146/#:~:text=50%25%20of%20all%20tweets%2C%20based,use%20of%20Quote%20Tweets%20here>).
- Ibar-Alonso, R., Quiroga-García, R., & Arenas-Parra, M. (2022). Opinion mining of green energy sentiment: A russia-ukraine conflict analysis. *Mathematics*, 10(14), 2532.
- IEA. (2020). Europe - Total energy supply. Europe data explorer. *International Energy Agency (IEA)*. Retrieved April 7, 2023, from <https://www.iea.org/regions/europe>
- infratest dimap. (2022). Knappe Mehrheit für vorübergehend stärkere Kohleverstromung und längere AKW-Laufzeiten. *ARD Trend Morgenmagazin*. Retrieved July 16, 2023, from <https://www.infratest-dimap.de/umfragen-analysen/bundesweit/umfragen/aktuell/knappe-mehrheit-fuer-voruebergehend-staerkere-kohleverstromung-und-laengere-akw-laufzeiten/>
- IPCC. (2018). *Global Warming of 1.5°C: An IPCC special report on impacts of global warming of 1.5°C above pre-industrial levels in context of strengthening response to climate change, sustainable development, and efforts to eradicate poverty* (V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. R. Shukla, et al., Eds.). Cambridge University Press.

- IPCC. (2023). Summary for Policymakers. In Core Writing Team, H. Lee, & J. Romero (Eds.), *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 1–34).
- IRSN. (2011). Accident du 12 septembre 2011 dans l'installation Centraco de Marcoule. *Institut de Radioprotection et de Sûreté Nucléaire (IRSN)*. Retrieved December 10, 2023, from <https://www.irsn.fr/savoir-comprendre/surete/accident-12-septembre-2011-dans-linstallation-centraco-marcoule>
- Jackoway, A., Samet, H., & Sankaranarayanan, J. (2011). Identification of live news events using Twitter. *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, 25–32.
- Jackson, R. B., Vengosh, A., Carey, J. W., Davies, R. J., Darrah, T. H., O'sullivan, F., & Pétron, G. (2014). The environmental costs and benefits of fracking. *Annual review of Environment and Resources*, 39, 327–362.
- Jayanti, S. (2022). Europe's Energy Crisis Is Going to Get Worse. The World Will Bear the Cost. *Time*. Retrieved April 7, 2023, from <https://time.com/6209272/europes-energy-crisis-getting-worse/>
- Jelodar, H., Wang, Y., Yuan, C., Feng, X., Jiang, X., Li, Y., & Zhao, L. (2019). Latent Dirichlet allocation (LDA) and topic modeling: models, applications, a survey. *Multimedia Tools and Applications*, 78, 15169–15211.
- Jenkins, K., McCauley, D., Heffron, R., Stephan, H., & Rehner, R. (2016). Energy justice: A conceptual review. *Energy Research & Social Science*, 11, 174–182.
- Jeong, S. Y., Kim, J. W., Kim, Y. S., Joo, H. Y., & Moon, J. H. (2021). Sentiment analysis of nuclear energy-related articles and their comments on a portal site in Rep. of Korea in 2010–2019. *Nuclear Engineering and Technology*, 53(3), 1013–1019.
- Jiang, H., Lin, P., & Qiang, M. (2016). Public-opinion sentiment analysis for large hydro projects. *Journal of Construction Engineering and Management*, 142(2), 05015013.
- Jockers, M. L. (2013). *Macroanalysis: Digital methods and literary history*. University of Illinois Press.
- Kaiser, R. (2020). Bäume, die die Welt bedeuten. Der Hambacher Forst als Symbol der deutschen Klimabewegung. *Soziologiemagazin*, 13(2), 11–12.
- Kallbekken, S. (2023). Research on public support for climate policy instruments must broaden its scope. *Nature Climate Change*, 13(3), 206–208.
- Kalt, G., Wiedenhofer, D., Görg, C., & Haberl, H. (2019). Conceptualizing energy services: A review of energy and well-being along the Energy Service Cascade. *Energy Research & Social Science*, 53, 47–58.
- Karami, A., Lundy, M., Webb, F., & Dwivedi, Y. K. (2020). Twitter and research: A systematic literature review through text mining. *IEEE access*, 8, 67698–67717.

- Karimzadeh, M., Huang, W., Banerjee, S., Wallgrün, J. O., Hardisty, F., Pezanowski, S., Mitra, P., & MacEachren, A. M. (2013). GeoTxt: a web API to leverage place references in text. *Proceedings of the 7th workshop on geographic information retrieval*, 72–73.
- Kaufer, R., & Lein, P. (2018). Widerstand im Hambacher Forst: Analyse einer anarchistischen Waldbesetzung. *Forschungsjournal Soziale Bewegungen*, 31(4).
- Keller, T. R., & Klinger, U. (2019). Social bots in election campaigns: Theoretical, empirical, and methodological implications. *Political Communication*, 36(1), 171–189.
- Kersting, J., & Geierhos, M. (2020). Aspect Phrase Extraction in Sentiment Analysis with Deep Learning. *ICAART (1)*, 391–400.
- Khairnar, J., & Kinikar, M. (2013). Machine learning algorithms for opinion mining and sentiment classification. *International Journal of Scientific and Research Publications*, 3(6), 1–6.
- Khan, M. I., Yasmin, T., & Shakoor, A. (2015). Technical overview of compressed natural gas (CNG) as a transportation fuel. *Renewable and Sustainable Energy Reviews*, 51, 785–797.
- Kherwa, P., & Bansal, P. (2019). Topic modeling: a comprehensive review. *EAI Endorsed transactions on scalable information systems*, 7(24), 1–16.
- Kim, S. Y., Ganesan, K., Dickens, P., & Panda, S. (2021). Public sentiment toward solar energy—opinion mining of twitter using a transformer-based language model. *Sustainability*, 13(5), 2673.
- Kim, Y., Kim, M., & Kim, W. (2013). Effect of the Fukushima nuclear disaster on global public acceptance of nuclear energy. *Energy policy*, 61, 822–828.
- Klinglmair, A., Bliem, M. G., & Brouwer, R. (2015). Exploring the public value of increased hydropower use: a choice experiment study for Austria. *Journal of Environmental Economics and Policy*, 4(3), 315–336.
- Knox, E. (2005). Oil combustion and childhood cancers. *Journal of Epidemiology & Community Health*, 59(9), 755–760.
- Komendantova, N., & Neumueller, S. (2020). Discourses about energy transition in Austrian climate and energy model regions: Turning awareness into action. *Energy & Environment*, 31(8), 1473–1497.
- Kosmajac, D., & Keselj, V. (2019). Twitter bot detection using diversity measures. *Proceedings of the 3rd International Conference on Natural Language and Speech Processing*, 1–8.
- Kristiansen, S., Bonfadelli, H., & Kovic, M. (2018). Risk perception of nuclear energy after Fukushima: Stability and change in public opinion in Switzerland. *International Journal of Public Opinion Research*, 30(1), 24–50.
- Krohn, S., & Damborg, S. (1999). On public attitudes towards wind power. *Renewable energy*, 16(1-4), 954–960.
- Lai, E. L., Moyer, D., Yuan, B., Fox, E., Hunter, B., Bertozzi, A. L., & Brantingham, P. J. (2016). Topic time series analysis of microblogs. *IMA Journal of Applied Mathematics*, 81(3), 409–431.

- Lai, T. M., Tran, Q. H., Bui, T., & Kihara, D. (2020). A simple but effective bert model for dialog state tracking on resource-limited systems. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8034–8038.
- Lanz, S., & Morgenthaler, M. (2022, September). *Lebenswerte Zukunft* (Online survey commissioned by the BKW). Sotomo. Retrieved March 2, 2023, from <https://lebensraeume.bkw.ch/news/studie-lebensraeume>
- Lee, B.-S. (2012). A Study of the Correlation between the Spatial Attributes on Twitter. *2012 IEEE 28th International Conference on Data Engineering Workshops*, 337–340.
- Lee, T. M., Markowitz, E. M., Howe, P. D., Ko, C.-Y., & Leiserowitz, A. A. (2015). Predictors of public climate change awareness and risk perception around the world. *Nature climate change*, 5(11), 1014–1020.
- Leetaru, K., Wang, S., Cao, G., Padmanabhan, A., & Shook, E. (2013). Mapping the global Twitter heartbeat: The geography of Twitter. *First Monday*, 18(5–6).
- Lei, Z., Wan, H., Zhang, W., Feng, S., Chen, Z., Zheng, Q., & Luo, M. (2022). BIC: Twitter Bot Detection with Text-Graph Interaction and Semantic Consistency. *arXiv preprint arXiv:2208.08320*.
- Leonard, M. D., Michaelides, E. E., & Michaelides, D. N. (2020). Energy storage needs for the substitution of fossil fuel power plants with renewables. *Renewable Energy*, 145, 951–962.
- Levenshtein, V. I., et al. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet physics doklady*, 10(8), 707–710.
- Li, R., Crowe, J., Leifer, D., Zou, L., & Schoof, J. (2019). Beyond big data: Social media challenges and opportunities for understanding social perception of energy. *Energy Research & Social Science*, 56, 101217.
- Lieberman, M. D., Samet, H., & Sankaranarayanan, J. (2010). Geotagging with local lexicons to build indexes for textually-specified spatial data. *2010 IEEE 26th international conference on data engineering (ICDE 2010)*, 201–212.
- Linville, D. L., McGee, S. E., & Hicks, L. K. (2012). Colleges' and universities' use of Twitter: A content analysis. *Public relations review*, 38(4), 636–638.
- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1), 1–167.
- Liu, P., Koivisto, S., Hiippala, T., van der Lijn, C., Väisänen, T., Nurmi, M., Toivonen, T., Vehkakoski, K., Pyykönen, J., Virmasalo, I., et al. (2022). Extracting locations from sport and exercise-related social media messages using a neural network-based bilingual toponym recognition model. *Journal of Spatial Information Science*, 24, 31–61.
- Liu, X., Yuan, X., Ma, F., & Xia, J. (2023). The increasing risk of energy droughts for hydropower in the Yangtze River basin. *Journal of Hydrology*, 621, 129589.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

- Liu, Z., & Na, J.-C. (2018). Aspect-based sentiment analysis of nuclear energy tweets with attentive deep neural network. *Maturity and Innovation in Digital Libraries: 20th International Conference on Asia-Pacific Digital Libraries, ICADL 2018, Hamilton, New Zealand, November 19-22, 2018, Proceedings 20*, 99–111.
- Lloyd, S., & Nakamura, T. (2022). Public Perceptions of Renewable Energy in the Philippines. *Sustainability*, *14*, 9906.
- Loureiro, M. L., & Alló, M. (2020). Sensing climate change and energy issues: Sentiment and emotion analysis with social media in the UK and Spain. *Energy Policy*, *143*, 111490.
- Luderer, G., Madeddu, S., Merfort, L., Ueckerdt, F., Pehl, M., Pietzcker, R., Rottoli, M., Schreyer, F., Bauer, N., Baumstark, L., et al. (2022). Impact of declining renewable energy costs on electrification in low-emission scenarios. *Nature Energy*, *7*(1), 32–42.
- Lundberg, J., Nordqvist, J., & Laitinen, M. (2019). Towards a language independent Twitter bot detector. *DHN*, 308–319.
- Mac, R., & Hsu, T. (2023). From Twitter to X: Elon Musk begins erasing an icon internet brand. *The New York Times*. Retrieved August 3, 2023, from <https://www.nytimes.com/2023/07/24/technology/twitter-x-elon-musk.html>
- Mahmud, J., Nichols, J., & Drews, C. (2012). Where is this tweet from? inferring home locations of twitter users. *Proceedings of the International AAAI Conference on Web and Social Media*, *6*(1), 511–514.
- Maier, G. (2014). OpenStreetMap, the wikipedia map. *Region*, *1*(1), R3–R10.
- Makice, K. (2009). *Twitter API: Up and running: Learn how to build applications with the Twitter API*. O'Reilly Media, Inc.
- Mäntylä, M. V., Graziotin, D., & Kuutila, M. (2018). The evolution of sentiment analysis – A review of research topics, venues, and top cited papers. *Computer Science Review*, *27*, 16–32.
- Mardani, A., Jusoh, A., Zavadskas, E. K., Cavallaro, F., & Khalifah, Z. (2015). Sustainable and renewable energy: An overview of the application of multiple criteria decision making techniques and approaches. *Sustainability*, *7*(10), 13947–13984.
- Marine-Roig, E. (2022). Content analysis of online travel reviews. In Z. Xiang, M. Fuchs, U. Gretzel, & W. Höpken (Eds.), *Handbook of e-Tourism* (pp. 1–26). Springer.
- Mariotti, S., Köhnen, E., Scheler, F., Sveinbjörnsson, K., Zimmermann, L., Piot, M., Yang, F., Li, B., Warby, J., Musiienko, A., et al. (2023). Interface engineering for high-performance, triple-halide perovskite–silicon tandem solar cells. *Science*, *381*(6653), 63–69.
- Markard, J., Isoaho, K., & Widdel, L. (2023). Discourses around decline: Comparing the debates on coal phase-out in the UK, Germany and Finland. *Technologies in Decline: Socio-Technical Approaches to Discontinuation and Destabilisation*, 119–144.
- Markard, J., Rinscheid, A., & Widdel, L. (2021). Analyzing transitions through the lens of discourse networks: Coal phase-out in Germany. *Environmental Innovation and Societal Transitions*, *40*, 315–331.

- Martin, M., & Islar, M. (2021). The ‘end of the world’ vs. the ‘end of the month’: understanding social resistance to sustainability transition agendas, a lesson from the Yellow Vests in France. *Sustainability Science*, *16*, 601–614.
- Maslow, A. H. (2000). A theory of human motivation. *Classics in Management Thought-Edward Elgar Publishing*, *1*, 450.
- Maynard, D. G., & Greenwood, M. A. (2014). Who cares about sarcastic tweets? investigating the impact of sarcasm on sentiment analysis. *Language Resources and Evaluation Conference (LREC) 2014 proceedings*.
- McAdam, D., & Boudet, H. (2012). *Putting social movements in their place: Explaining opposition to energy projects in the United States, 2000–2005*. Cambridge University Press.
- McCurley, K. S. (2001). Geospatial mapping and navigation of the web. *Proceedings of the 10th international conference on World Wide Web*, 221–229.
- Mellon, J., & Prosser, C. (2017). Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users. *Research & Politics*, *4*(3), 1–9.
- Mez, L. (2012). Nuclear energy—Any solution for sustainability and climate protection? *Energy Policy*, *48*, 56–63.
- Mihm, A. (2017). Klimaschutz: Solarstrom ist überteuert und ineffizient. *Frankfurter Allgemeine Zeitung*. Retrieved December 11, 2023, from <https://www.faz.net/aktuell/wirtschaft/hohe-co2-vermeidungskosten-fuer-photovoltaik-15365573.html>
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Möhl, S. (2010). 120.000 Demonstranten – Anti-Atomkraft-Bewegung lebt. *Welt*. Retrieved December 9, 2023, from <https://www.welt.de/aktuell/article7325832/120-000-Demonstranten-Anti-Atomkraft-Bewegung-lebt.html>
- Mohtasham, J. (2015). Renewable energies. *Energy Procedia*, *74*, 1289–1297.
- Moriarty, P., & Honnery, D. (2012). What is the global potential for renewable energy? *Renewable and Sustainable Energy Reviews*, *16*(1), 244–252.
- Morino, Y., Ohara, T., & Nishizawa, M. (2011). Atmospheric behavior, deposition, and budget of radioactive materials from the Fukushima Daiichi nuclear power plant in March 2011. *Geophysical research letters*, *38*(7), 1–7.
- Morstatter, F., Pfeffer, J., Liu, H., & Carley, K. (2013). Is the sample good enough? comparing data from twitter’s streaming api with twitter’s firehose. *Proceedings of the international AAAI conference on web and social media*, *7*(1), 400–408.
- Müller, S., Backhaus, N., & Buchecker, M. (2020). Mapping meaningful places: A tool for participatory siting of wind turbines in Switzerland? *Energy Research & Social Science*, *69*, 101573.
- Müller-Hansen, F., Lee, Y. T., Callaghan, M., Jankin, S., & Minx, J. C. (2022). The German coal debate on Twitter: Reactions to a corporate policy process. *Energy Policy*, *169*, 113178.

- Müller-Hansen, F., Repke, T., Baum, C. M., Brutschin, E., Callaghan, M. W., Debnath, R., Lamb, W. F., Low, S., Lück, S., Roberts, C., et al. (2023). Attention, sentiments and emotions towards emerging climate technologies on Twitter. *Global Environmental Change*, *83*, 102765.
- Munawer, M. E. (2018). Human health and environmental impacts of coal combustion and post-combustion wastes. *Journal of Sustainable Mining*, *17*(2), 87–96.
- Munro, R. (2011). Subword and spatiotemporal models for identifying actionable information in Haitian Kreyol. *Proceedings of the fifteenth conference on computational natural language learning*, 68–77.
- National Geographic Education. (2023). Fossil Fuels: What is a fossil fuel and what is being done to make fossil fuels more environmentally friendly? *National Geographic*. Retrieved August 24, 2023, from <https://education.nationalgeographic.org/resource/fossil-fuels/>
- Nauman, F., & Herschel, M. (2022). *An introduction to duplicate detection*. Springer Nature.
- Negara, E. S., Triadi, D., & Andryani, R. (2019). Topic modelling twitter data with latent dirichlet allocation method. *2019 International Conference on Electrical Engineering and Computer Science (ICECOS)*, 386–390.
- Nicoletti, G., Arcuri, N., Nicoletti, G., & Bruno, R. (2015). A technical and environmental comparison between hydrogen and some fossil fuels. *Energy Conversion and Management*, *89*, 205–213.
- Noblet, C. L., Teisl, M. F., Evans, K., Anderson, M. W., McCoy, S., & Cervone, E. (2015). Public preferences for investments in renewable energy production and energy efficiency. *Energy Policy*, *87*, 177–186.
- Nominatim. (2023a). Importances. *Nominatim Manual*. Retrieved November 23, 2023, from <https://nominatim.org/release-docs/latest/customize/Importance/>
- Nominatim. (2023b). Place Ranking in Nominatim. *Nominatim Manual*. Retrieved November 23, 2023, from <https://nominatim.org/release-docs/latest/customize/Ranking/#search-rank>
- Novotny, J. (2019). Twitter bot detection & categorization-a comparative study of machine learning methods.
- Nuortimo, K., & Härkönen, J. (2018). Opinion mining approach to study media-image of energy production. Implications to public acceptance and market deployment. *Renewable and Sustainable Energy Reviews*, *96*, 210–217.
- Nuortimo, K., Härkönen, J., & Karvonen, E. (2018). Exploring the global media image of solar power. *Renewable and Sustainable Energy Reviews*, *81*, 2806–2811.
- Ochsenbein, T. (2023). Schweizer wollen AKW zurück. *Blick*. Retrieved November 2, 2023, from <https://www.blick.ch/politik/blick-umfrage-zur-energiepolitik-zeigt-schweizer-wollen-akws-zurueck-id18391702.html>
- Okeku, K. E. (2018). *Stability Control and Reliable Performance of Wind Turbines*. InTech.
- Olson-Hazboun, S. K., Krannich, R. S., & Robertson, P. G. (2016). Public views on renewable energy in the Rocky Mountain region of the United States: Distinct attitudes, exposure, and other key predictors of wind energy. *Energy Research & Social Science*, *21*, 167–179.

- ORF.at. (2022a, June). Kohlekraftwerk wird reaktiviert: Gewessler kündigt weitere Schritte an. *ORF News*. Retrieved December 11, 2023, from <https://orf.at/stories/3271993/>
- ORF.at. (2022b, August). Wassermangel in Norwegen: Nächstes Problem für Europas Stromnetz. *ORF News*. Retrieved December 16, 2023, from <https://orf.at/stories/3280222/>
- Osman, A. I., Chen, L., Yang, M., Msigwa, G., Farghali, M., Fawzy, S., Rooney, D. W., & Yap, P.-S. (2023). Cost, environmental impact, and resilience of renewable energy under a changing climate: a review. *Environmental Chemistry Letters*, *21*(2), 741–764.
- Owusu, P. A., & Asumadu-Sarkodie, S. (2016). A review of renewable energy sources, sustainability issues and climate change mitigation. *Cogent Engineering*, *3*(1), 1167990.
- Pak, A., Paroubek, P., et al. (2010). Twitter as a corpus for sentiment analysis and opinion mining. *LREc*, *10*(2010), 1320–1326.
- Pang, B., Lee, L., et al. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in information retrieval*, *2*(1–2), 1–135.
- Paradesi, S. M. (2011). Geotagging tweets using their content. *Twenty Fourth International FLAIRS Conference*, 355–356.
- Purves, R. S., Koblet, O., & Adams, B. (2022). *Unlocking Environmental Narratives: Towards Understanding Human Environment Interactions through Computational Text Analysis*. Ubiquity Press.
- Qazi, A., Hussain, F., Rahim, N. A., Hardaker, G., Alghazzawi, D., Shaban, K., & Haruna, K. (2019). Towards sustainable energy: a systematic review of renewable energy sources, technologies, and public opinions. *IEEE access*, *7*, 63837–63851.
- Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A Python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.
- Rahimnejad, M., Ghoreyshi, A., Najafpour, G., Younesi, H., & Shakeri, M. (2012). A novel microbial fuel cell stack for continuous production of clean energy. *International journal of hydrogen energy*, *37*(7), 5992–6000.
- Rahimnejad, M., Adhami, A., Darvari, S., Zirepour, A., & Oh, S.-E. (2015). Microbial fuel cell as new technology for bioelectricity generation: A review. *Alexandria Engineering Journal*, *54*(3), 745–756.
- Raiwa, M., Kneip, N., Walther, C., & Wendt, K. (2023). Dem Kernbrennstoff auf der Spur: Nukleare Forensik an kleinsten Umweltpartikeln. *Physik in unserer Zeit*, *54*(1), 24–29.
- Rand, J., & Hoen, B. (2017). Thirty years of North American wind energy acceptance research: What have we learned? *Energy research & social science*, *29*, 135–148.
- Rapp, H.-J. (2019). 1.000 Jahre Reinhardswald—ein Jubiläum zur rechten Zeit. *Jahrb. Naturschutz Hess*, *18*, 43–46.
- Rapp, H.-J. (2021). Windkraft um jeden Preis - Aber bitte nicht im Reinhardswald. *Jahrb. Naturschutz Hess*, *20*, 95–101.
- Řehůřek, R., & Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora. *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50.

- Reinhardt, W. (2009). How people are using Twitter during conferences. *Proceeding of 5th EduMedia Conference*.
- Remus, R., Quasthoff, U., & Heyer, G. (2010). SentiWS-A Publicly Available German-language Resource for Sentiment Analysis. *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC '10)*.
- Rivest, R. (1992). *The MD5 message-digest algorithm* (tech. rep.).
- Roberts, M. E., Stewart, B. M., & Airoidi, E. M. (2016). A model of text for experimentation in the social sciences. *Journal of the American Statistical Association*, 111(515), 988–1003.
- Roesslein, J. (2018, March). tweepy Documentation, Release 3.6.0. Retrieved January 12, 2023, from <https://buildmedia.readthedocs.org/media/pdf/tweepy/v3.6.0/tweepy.pdf>
- Roggenkamp, M. M., de Graaf, K. J., & Fleming, R. C. (2021). *Energy law, climate change and the environment*. Edward Elgar Publishing.
- Rothman, D., & Gulli, A. (2022). *Transformers for Natural Language Processing: Build, train, and fine-tune deep neural network architectures for NLP with Python, PyTorch, TensorFlow, BERT, and GPT-3*. Packt Publishing Ltd.
- Ruhnau, O., Stiewe, C., Muessel, J., & Hirth, L. (2023). Natural gas savings in Germany during the 2022 energy crisis. *Nature Energy*, 1–8.
- Ryoo, K., & Moon, S. (2014). Inferring twitter user locations with 10 km accuracy. *Proceedings of the 23rd International Conference on World Wide Web*, 643–648.
- Rzepka, D. (2022). Ein Grüner reaktiviert die Kohle: Habecks "bittere" Kehrtwende [only available via web.archive.org]. *ZDF Nachrichten*. Retrieved December 11, 2023, from <https://www.zdf.de/nachrichten/politik/habeck-gas-kohle-bdi-100.html>
- Sagcan, M., & Karagoz, P. (2015). Toponym recognition in social media for estimating the location of events. *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, 33–39.
- Salazar, J., Liang, D., Nguyen, T. Q., & Kirchoff, K. (2019). Masked language model scoring. *arXiv preprint arXiv:1910.14659*.
- Salman, S., & Liu, X. (2019). Overfitting mechanism and avoidance in deep neural networks. *arXiv preprint arXiv:1901.06566*.
- Sander, H. (2016). *Die Klimagerechtigkeitsbewegung in Deutschland. Entwicklungen und Perspektiven* (Study commissioned by the Rosa-Luxemburg-Stiftung).
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Scharpf, F. (1999). *Governing in Europe: Effective and democratic?* Oxford University Press.
- Scherhauer, P., Höltinger, S., Salak, B., Schauppenlehner, T., & Schmidt, J. (2017). Patterns of acceptance and non-acceptance within energy landscapes: A case study on wind energy expansion in Austria. *Energy Policy*, 109, 863–870.
- Schmidt, T., Fehle, J., Weissenbacher, M., Richter, J., Gottschalk, P., & Wolff, C. (2022). Sentiment Analysis on Twitter for the Major German Parties during the 2021 German Federal Election.

- Proceedings of the 18th Conference on Natural Language Processing (KONVENS 2022)*, 74–87.
- Schmitt, M., Steinheber, S., Schreiber, K., & Roth, B. (2018). Joint aspect and polarity classification for aspect-based sentiment analysis with end-to-end neural networks. *arXiv preprint arXiv:1808.09238*.
- Schmitt, X., Kubler, S., Robert, J., Papadakis, M., & LeTraon, Y. (2019). A replicable comparison study of NER software: StanfordNLP, NLTK, OpenNLP, SpaCy, Gate. *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, 338–343.
- Schopf, T., Klimek, S., & Matthes, F. (2022). Patternrank: leveraging pretrained language models and part of speech for unsupervised keyphrase extraction. *arXiv preprint arXiv:2210.05245*.
- Schulz, A., Hadjakos, A., Paulheim, H., Nachtwey, J., & Mühlhäuser, M. (2013). A multi-indicator approach for geolocation of tweets. *Proceedings of the International AAAI Conference on Web and Social Media*, 7(1), 573–582.
- Schweizerische Eidgenossenschaft. (2017, May). Volksabstimmung vom 21. Mai 2017. Erläuterung des Bundesrates - Energiegesetz (EnG). Retrieved December 15, 2023, from https://www.admin.ch/dam/gov/de/Dokumentation/Abstimmungen/Mai17/LowRes_Volksabstimmung_Erlaeuterung_21.05.2017_DE.pdf.download.pdf/Erl%C3%A4uterungen%20des%20Bundesrates%20-%20Volksabstimmung%20vom%2021.%20Mai%202017.pdf
- Sedhai, S., & Sun, A. (2017). Semi-supervised spam detection in Twitter stream. *IEEE Transactions on Computational Social Systems*, 5(1), 169–175.
- Segreto, M., Principe, L., Desormeaux, A., Torre, M., Tomassetti, L., Tratzi, P., Paolini, V., & Petracchini, F. (2020). Trends in social acceptance of renewable energy across Europe – A literature review. *International Journal of Environmental Research and Public Health*, 17(24), 9161.
- Semin, R. A. B. (2008). A technical review of compressed natural gas as an alternative fuel for internal combustion engines. *Am. J. Eng. Appl. Sci*, 1(4), 302–311.
- Sharma, S., & Gupta, V. (2022). Role of twitter user profile features in retweet prediction for big data streams. *Multimedia Tools and Applications*, 81(19), 27309–27338.
- Shi, W., & Demberg, V. (2019). Next sentence prediction helps implicit discourse relation classification within and across domains. *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, 5790–5796.
- Shipman, J. W. (2013). Tkinter 8.5 reference: a GUI for Python. *New Mexico Tech Computer Center*, 54, 356–359.
- Siebel, A. (2022, December). Energiekrise aus Verbrauchersicht [Summary of the survey]. *EY Deutschland*. Retrieved March 2, 2023, from https://www.ey.com/de_de/energy-resources/umfrage-zur-energiekrise-2022-aus-verbrauchersicht

- Singhal, S., Agarwal, S., Arora, S., Sharma, P., & Singhal, N. (2017). Upgrading techniques for transformation of biogas to bio-CNG: a review. *International Journal of Energy Research*, *41*(12), 1657–1669.
- Slovic, P. (1987). Perception of risk. *Science*, *236*(4799), 280–285.
- Sokołowski, J., Lewandowski, P., & Frankowski, J. (2023). *How to prevent another Yellow Vest Movement? Assessing public preference for a carbon tax with a discrete choice experiment* (IBS working paper).
- spaCy.io. (n.d.). spaCy: Fact and Figures. Retrieved March 2, 2023, from <https://spacy.io/usage/facts-figures#benchmarks>
- Spiegel.de. (2011, January). E.onAKW: Experten rätseln über Mini-Riss im Reaktor-Innenbereich. Retrieved December 9, 2023, from <https://www.spiegel.de/wissenschaft/technik/e-on-akw-experten-raetseln-ueber-mini-riss-im-reaktor-innenbereich-a-739714.html>
- Spiegel.de. (2019, March). Modellrechnung des DLR: Windkraftanlagen töten im Sommer täglich Milliarden Insekten. Retrieved December 11, 2023, from <https://www.spiegel.de/wissenschaft/natur/windkraftanlagen-toeten-im-sommer-taeglich-milliarden-insekten-a-1259462.html>
- Staddon, P. L., & Depledge, M. H. (2015). Fracking cannot be reconciled with climate change mitigation policies. *Environmental Science and Technology*, *49*(14), 8269–8270.
- Statistik Austria. (2021). Energieeinsatz der Haushalte. *Statistik Austria*. Retrieved April 9, 2023, from <https://www.statistik.at/statistiken/energie-und-umwelt/energie/energieeinsatz-der-haushalte>
- Stede, J., & May, N. (2019). Strikte Mindestabstände bremsen den Ausbau der Windenergie. *DIW Wochenbericht*, *86*(48), 895–903.
- Steffen, B., & Patt, A. (2022). A historical turning point? Early evidence on how the Russia-Ukraine war changes public support for clean energy policies. *Energy Research & Social Science*, *91*, 102758.
- Stokel-Walker, C. (2023). Twitter's \$42,000-per-Month API Prices Out Nearly Everyone. *Wired*. Retrieved March 10, 2023, from <https://www.wired.co.uk/article/twitter-data-api-prices-out-nearly-everyone>
- Süddeutsche Zeitung. (2018). Naturschutz - Aßlar: Kritik an milliardenfachem Tod von Fischen durch Wasserkraft. *Süddeutsche Zeitung*. Retrieved December 16, 2023, from <https://www.sueddeutsche.de/wissen/naturschutz-asslar-kritik-an-milliardenfachem-tod-von-fischen-durch-wasserkraft-dpa.urn-newsml-dpa-com-20090101-180415-99-898844>
- Sufi, F. K., & Khalil, I. (2022). Automated disaster monitoring from social media posts using AI-based location intelligence and sentiment analysis. *IEEE Transactions on Computational Social Systems*.
- Sun, C., Qiu, X., Xu, Y., & Huang, X. (2019). How to fine-tune bert for text classification? *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*, 194–206.

- Sütterlin, B., & Siegrist, M. (2017). Public acceptance of renewable energy technologies from an abstract versus concrete perspective and the positive imagery of solar power. *Energy Policy*, *106*, 356–366.
- Tagesspiegel.de. (2022, December). Erneut Weltrekord: Solarzelle aus Deutschland erreicht höchsten Wirkungsgrad. Retrieved December 13, 2023, from <https://www.tagesspiegel.de/wissen/wieder-weltrekord-solarzelle-aus-deutschland-erreicht-hochsten-wirkungsgrad-9058582.html>
- Tang, D., Qin, B., & Liu, T. (2015). Deep learning for sentiment analysis: successful approaches and future challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *5*(6), 292–303.
- Tao, K., Abel, F., Hauff, C., Houben, G.-J., & Gadiraju, U. (2013). Groundhog day: near-duplicate detection on twitter. *Proceedings of the 22nd international conference on World Wide Web*, 1273–1284.
- Tatham, M., & Peters, Y. (2023). Fueling opposition? Yellow vests, urban elites, and fuel taxation. *Journal of European Public Policy*, *30*(3), 574–598.
- Teisl, M. F., McCoy, S., Marrinan, S., Noblet, C. L., Johnson, T., Wibberly, M., Roper, R., & Klein, S. (2015). Will offshore energy face “Fair winds and following seas”? Understanding the factors influencing offshore wind acceptance. *Estuaries and coasts*, *38*, 279–286.
- Tollefson, J., & Gibney, E. (2022). Nuclear-fusion lab achieves ‘ignition’: what does it mean? *Nature News*. Retrieved December 13, 2023, from <https://www.nature.com/articles/d41586-022-04440-7>
- Tornes, A. (2021). Product News - Enabling the future of academic research with the Twitter API. *Twitter Developers*. Retrieved March 2, 2023, from <https://developer.twitter.com/en/blog/product-news/2021/enabling-the-future-of-academic-research-with-the-twitter-api>
- Trieb, F., Gerz, T., & Geiger, M. (2018). Modellanalyse liefert Hinweise auf Verluste von Fluginsekten in Windparks. *Energiewirtschaftliche Tagesfragen*, *68*(11), 51–55.
- Tripathy, J., Dash, R., Pattanayak, B. K., & Mohanty, B. (2021). Automated phrase mining using POST: The best approach. *2021 1st Odisha international conference on electrical power engineering, communication and computing technology (ODICON)*, 1–6.
- Turconi, R., Boldrin, A., & Astrup, T. (2013). Life cycle assessment (LCA) of electricity generation technologies: Overview, comparability and limitations. *Renewable and sustainable energy reviews*, *28*, 555–565.
- Turkenburg, W. C., Faaij, A., et al. (2000). Renewable energy technologies. In UNDP, UNDESA, & WEC (Eds.), *World energy assessment: energy and the challenge of sustainability* (pp. 219–272). UNDP.
- Twitter. (2023). Tweet location FAQ. *Twitter*. Retrieved February 24, 2023, from <https://help.twitter.com/en/safety-and-security/tweet-location-settings>

- Twitter Developer Platform. (2023a). *Search Tweets*. Retrieved January 11, 2023, from <https://developer.twitter.com/en/docs/twitter-api/tweets/search/api-reference/get-tweets-search-all>
- Twitter Developer Platform. (2023b). Twitter API. Retrieved January 11, 2023, from <https://developer.twitter.com/en/docs/twitter-api>
- Twitter (X) Developers. (2023, February). Twitter API updates information [Accessed: 2024-01-02]. <https://twitter.com/XDevelopers/status/1621026986784337922>
- UNECE. (2022). *REN 21: Renewable Energy Status Report 2022*. United Nations Economic Commission for Europe. Retrieved November 18, 2023, from https://www.ren21.net/wp-content/uploads/2019/05/REN21_UNECE2022_FullReport.pdf
- United Nations. (2015). Paris Agreement. *UNFCCC*. Retrieved May 13, 2023, from https://unfccc.int/sites/default/files/english_paris_agreement.pdf
- United Nations. (2022). Renewable energy – powering a safer future. *United Nations Climate Action*. Retrieved July 20, 2023, from <https://www.un.org/en/climatechange/raising-ambition/renewable-energy>
- Usuga-Cadavid, J. P., Lamouri, S., Grabot, B., & Fortin, A. (2022). Using deep learning to value free-form text data for predictive maintenance. *International Journal of Production Research*, *60*(14), 4548–4575.
- Vågerö, O., Bråte, A., Wittemann, A., Robinson, J. Y., Sirotko-Sibirskaya, N., & Zeyringer, M. (2023). Machine Learning of Public Sentiments toward Wind Energy in Norway. *arXiv preprint arXiv:2304.02388*.
- Vaillancourt, K., Labriet, M., Loulou, R., & Waaub, J.-P. (2008). The role of nuclear energy in long-term climate scenarios: An analysis with the World-TIMES model. *Energy Policy*, *36*(7), 2296–2307.
- Varol, O., Ferrara, E., Davis, C., Menczer, F., & Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. *Proceedings of the international AAAI conference on web and social media*, *11*(1), 280–289.
- Vaske, J. J., Don Carlos, A. W., Manfredo, M. J., & Teel, T. L. (2023). Evaluating alternative survey methodologies in human dimensions of wildlife research. *Human Dimensions of Wildlife*, *28*(4), 320–334.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
- Venus, T. E., Hinzmann, M., & Gerdes, H. (2022). Public Acceptance of Hydropower. *Novel Developments for Sustainable Hydropower*, 29–40.
- Vidadili, N., Suleymanov, E., Bulut, C., & Mahmudlu, C. (2017). Transition to renewable energy and sustainable energy development in Azerbaijan. *Renewable and Sustainable Energy Reviews*, *80*, 1153–1161.

- Viswanathan, S., Damodaran, N., Simon, A., George, A., Anand Kumar, M., & Soman, K. (2019). Detection of duplicates in Quora and Twitter corpus. *Advances in Big Data and Cloud Computing: Proceedings of ICBDDCC18*, 519–528.
- Voss, J. (2022). Kohleausstieg: Wie sich das Rheinische Revier neu erfinden will. *National Geographic*. Retrieved December 11, 2023, from <https://www.nationalgeographic.de/umwelt/2022/02/kohleausstieg-wie-sich-das-rheinische-revier-neu-erfinden-will>
- Vrana, V., Kydros, D., Kotzaivazolou, I., & Pechlivanaki, I. (2023). EU Citizens' Twitter Discussions of the 2022–23 Energy Crisis: A Content and Sentiment Analysis on the Verge of a Daunting Winter. *Sustainability*, 15(2), 1322.
- Vringer, K., & Carabain, C. L. (2020). Measuring the legitimacy of energy transition policy in the Netherlands. *Energy policy*, 138, 111229.
- Vuichard, P., Stauch, A., & Dällenbach, N. (2019). Individual or collective? Community investment, local taxes, and the social acceptance of wind energy in Switzerland. *Energy Research & Social Science*, 58, 101275.
- Wagner, B., Hauer, C., Schoder, A., & Habersack, H. (2015). A review of hydropower in Austria: Past, present and future development. *Renewable and Sustainable Energy Reviews*, 50, 304–314.
- Wang, J., & Hu, Y. (2019). Enhancing spatial and textual analysis with EUPEG: An extensible and unified platform for evaluating geoparsers. *Transactions in GIS*, 23(6), 1393–1419.
- Wang, J., Hu, Y., & Joseph, K. (2020). NeuroTPR: A neuro-net toponym recognition model for extracting locations from social media messages. *Transactions in GIS*, 24(3), 719–735.
- Weissenbacher, D., Tahsin, T., Beard, R., Figaro, M., Rivera, R., Scotch, M., & Gonzalez, G. (2015). Knowledge-driven geospatial location resolution for phylogeographic models of virus migration. *Bioinformatics*, 31(12), i348–i356.
- Wheeler, D. A. (1988). Atmospheric dispersal and deposition of radioactive material from Chernobyl. *Atmospheric Environment (1967)*, 22(5), 853–863.
- Wiebe, J., Bruce, R., & O'Hara, T. P. (1999). Development and use of a gold-standard data set for subjectivity classifications. *Proceedings of the 37th annual meeting of the Association for Computational Linguistics*, 246–253.
- Wiertz, T., Kuhn, L., & Mattissek, A. (2023). A turn to geopolitics: Shifts in the German energy transition discourse in light of Russia's war against Ukraine. *Energy Research & Social Science*, 98, 103036.
- Wojatzki, M., Ruppert, E., Holschneider, S., Zesch, T., & Biemann, C. (2017). Germeval 2017: Shared task on aspect-based sentiment in social media customer feedback. *Proceedings of the GermEval*, 1–12.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. (2020). Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, 38–45.

- Wolske, K. S., Stern, P. C., & Dietz, T. (2017). Explaining interest in adopting residential solar photovoltaic systems in the United States: Toward an integration of behavioral theories. *Energy research & social science*, 25, 134–151.
- Wu, J., Teng, E., & Cao, Z. (2022). Twitter Bot Detection Through Unsupervised Machine Learning. *2022 IEEE International Conference on Big Data (Big Data)*, 5833–5839.
- Wüstenhagen, R., Wolsink, M., & Bürer, M. J. (2007). Social acceptance of renewable energy innovation: An introduction to the concept. *Energy policy*, 35(5), 2683–2691.
- Xu, H., Tang, T., Zhang, B., & Liu, Y. (2022). Automatic sentiment analysis of public opinion on nuclear energy. *Kerntechnik*, 87(2), 167–175.
- Xue, J., Chen, J., Chen, C., Zheng, C., Li, S., & Zhu, T. (2020). Public discourse and sentiment during the COVID 19 pandemic: Using Latent Dirichlet Allocation for topic modeling on Twitter. *PloS one*, 15(9), e0239441.
- Yacouby, R., & Axman, D. (2020). Probabilistic extension of precision, recall, and f1 score for more thorough evaluation of classification models. *Proceedings of the first workshop on evaluation and comparison of NLP systems*, 79–91.
- Yan, J. (2006). Bot, Cyborg and Automated Turing Test: (Or “Putting the Humanoid in the Protocol”). *International Workshop on Security Protocols*, 190–197.
- Yang, M., Liang, Y., Zhao, W., Xu, W., Zhu, J., & Qu, Q. (2018). Task-oriented keyphrase extraction from social media. *Multimedia Tools and Applications*, 77, 3171–3187.
- Yang, S., & Zhang, H. (2018). Text mining of Twitter data using a latent Dirichlet allocation topic model and sentiment analysis. *International Journal of Computer and Information Engineering*, 12(7), 525–529.
- Yin, W., & Fan, H. (2023). Perception changes and the attribution of the impact of Lancang-Mekong hydropower dams in the media of riparian countries from 1971 to 2020. *Journal of Hydrology: Regional Studies*, 45, 101302.
- Ying, X. (2019). An overview of overfitting and its solutions. *Journal of physics: Conference series*, 1168, 022022.
- Yoosuf, S., & Yang, Y. (2019). Fine-grained propaganda detection with fine-tuned BERT. *Proceedings of the second workshop on natural language processing for internet freedom: censorship, disinformation, and propaganda*, 87–91.
- Zachmann, G., Sgaravatti, G., & McWilliams, B. (2022). European natural gas imports. *Bruegel Dataset*, 20.
- Zhang, B., Niu, N., Li, H., & Wang, Z. (2023). Assessing the efforts of coal phaseout for carbon neutrality in China. *Applied Energy*, 352, 121924.
- Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4), e1253.
- Zhang, W., Li, X., Deng, Y., Bing, L., & Lam, W. (2022). A survey on aspect-based sentiment analysis: Tasks, methods, and challenges. *IEEE Transactions on Knowledge and Data Engineering*, 35(11), 11019–11038.

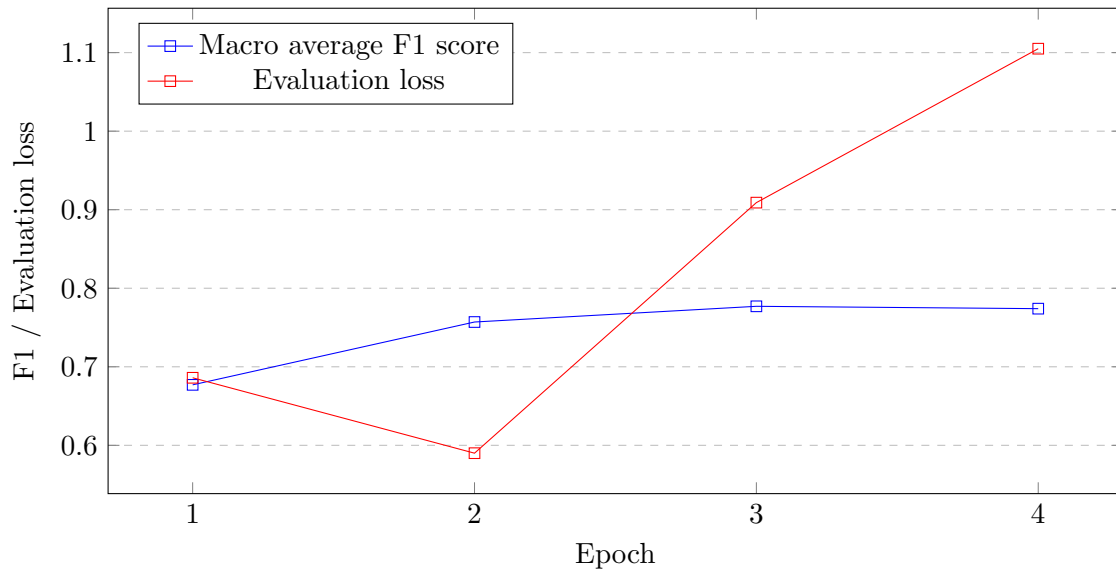
- Zhang, X. (2023). Diffusion dynamics and digital movement: The emergence and proliferation of the German-speaking# FridaysForFuture network on twitter. *Social Movement Studies*, 1–22.
- Zhao, J., & Luo, Y. (2021). A framework to address cognitive biases of climate change. *Neuron*, 109(22), 3548–3551.
- Zheng, X., Han, J., & Sun, A. (2018). A survey of location prediction on twitter. *IEEE Transactions on Knowledge and Data Engineering*, 30(9), 1652–1671.
- Zhou, B., Zou, L., Hu, Y., & Qiang, Y. (2023). TopoBERT: Plug and Play Toponym Recognition Module Harnessing Fine-tuned BERT. *arXiv preprint arXiv:2301.13631*.

9 Appendix

Appendix A: Model performance after each epoch during the fine-tuning process

The figure below shows the evaluation F1 score and the evaluation loss after each epoch during the final fine-tuning phase. Due to a missing parameter in the HuggingFace Trainer class, checkpoint 4 was chosen for the final aspect-based sentiment model. Compared to checkpoint 2, checkpoint 4 was found to predict the sentiment of most tweets with a high confidence. Hence, this led to a larger evaluation loss as the disparity between the true label and the predicted label of misclassified tweets maximized compared to the other epochs. Checkpoint 3 would have been the optimal choice due to its largest F1 and a lower evaluation loss.

Changes of F1 and evaluation loss over epochs



Appendix B: Temporal shift of the discourse

During microreading, it was observed that people referred to the health and environmental dimensions more often in the first half of the study period while the climate and energy supply dimensions became dominant in the second half. This qualitative observation was quantitatively found in the data for nuclear power. The following figure displays this discourse shift by plotting the monthly frequencies of characteristic words (normalized by the total unique words per month).

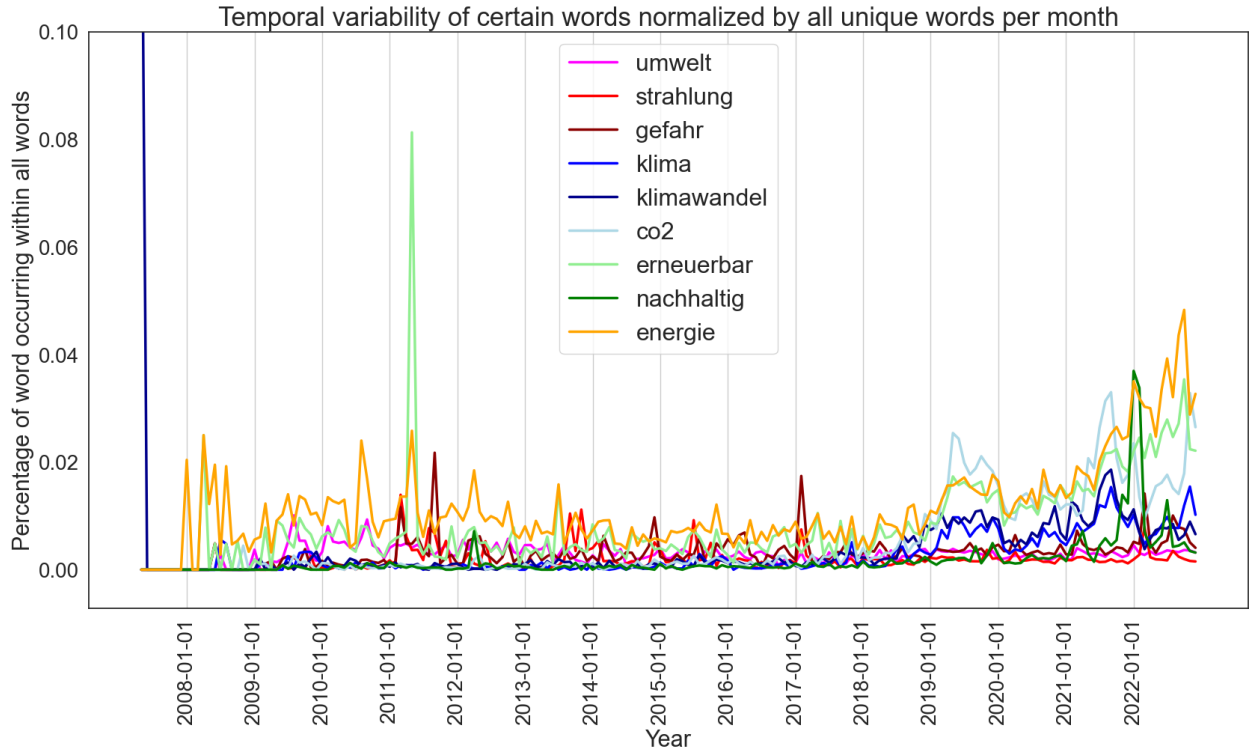


Figure 43: Shift of prioritized dimensions observed in data for nuclear power. Starting in 2018, the climate and energy supply topics replaced previously dominant environmental and health arguments.

Personal declaration

I hereby declare that the submitted thesis is the result of my own, independent work. All external sources are explicitly acknowledged in the thesis.



Marco Stutz

30.01.2024