



Reconstructing Arguments from Noisy Text

Introduction to the RANT project

Natalie Dykes¹ · Stefan Evert¹ · Merlin Göttlinger² · Philipp Heinrich¹  · Lutz Schröder²

Received: 14 February 2020 / Accepted: 29 May 2020 / Published online: 1 July 2020
© The Author(s) 2020

Abstract

Social media are of paramount importance to public discourse. RANT aims to contribute methods and formalisms for extracting, representing, and processing arguments from noisy text found in social media discussions, using a large corpus of pre-referendum Brexit tweets as a running case study. We identify recurring linguistic argumentation patterns in a corpus-linguistic analysis and formulate corresponding corpus queries to extract arguments automatically. Given the huge amount of social media data available, our approach aims at high precision at the possible price of low recall. Argumentation patterns are directly associated with logical patterns in a dedicated formalism and accordingly, individual arguments are directly parsed as logical formulae. The logical formalism for argument representation features a broad range of modalities capturing real-life modes of expression. We cast this formalism as a family of instance logics in the generic framework of coalgebraic logic and complement it by a flexible framework to represent relationships between arguments; including standard relations like attack and support but also relations extracted from metadata. Some relations are inferred from the logical content of individual arguments. We are in the process of developing suitable generalizations of various extension semantics for argumentation frameworks combined with corresponding algorithmic methods to allow for the automated retrieval of large-scale argumentative positions.

The work is supported by the DFG under Project 377333057
“Reconstructing Arguments from Noisy Text (RANT)”.

Natalie Dykes
natalie.mary.dykes@fau.de

Stefan Evert
stefan.evert@fau.de

Merlin Göttlinger
merlin.goettlinger@fau.de

✉ Philipp Heinrich
philipp.heinrich@fau.de

Lutz Schröder
lutz.schroeder@fau.de

¹ Computational Corpus Linguistics, FAU Erlangen-Nürnberg, Erlangen, Germany

² Theoretical Computer Science, FAU Erlangen-Nürnberg, Erlangen, Germany

1 Introduction

Public discourse is becoming increasingly dominated by social media, creating a need for robust analysis of arguments raised in computer-mediated communication. The objective of RANT is to explore the possibility of conducting such an analysis in a formally grounded manner despite the high level of both syntactic and semantic variability inherent in the exchanged arguments, based on a large corpus of Twitter messages on the Brexit referendum. In view of the sheer size of the corpus, we follow a high-precision/low-recall approach where we fine-tune corpus queries to extract arguments matching a set of logical patterns (manually developed by inspection of the corpus) and embed the harvested arguments into an expressive logical/argumentation-theoretic framework.

Peculiarities of social media data occur both on a linguistic and on an argumentation-theoretical level.

Linguistic challenges result from the largely unmoderated environment of Twitter, imposing little structure or stylistic guidelines upon users. Tweets often feature multimodality in the form of links, videos or pictures and non-

standard language is prevalent. Linguistic phenomena like abbreviations, colloquial style and typographic errors challenge both standard NLP pipelines and traditional frameworks for extracting argumentation. This leads to poor performance on tasks such as detecting argumentative utterances or splitting claims and premises [9, 13].

Argumentation theoretic challenges mostly result from the informal nature of conversations on social media. Arguments in day to day conversations tend to feature a high degree of implicitness because users normally assume common knowledge on the side of the listeners. It is possible for whole premises, conclusions, and any intermediate steps of arguments to remain implicit [5]. This effect is amplified by the fast pace of social media communication and especially Twitter's limit of 140 characters¹. Such incomplete arguments are referred to as *enthymemes* in the literature [28] and making the missing information explicit is difficult even for human annotators [4].

In this setting of incomplete defeasible arguments, persuasion through rhetoric strategies such as selection, arrangement, or phrasing of argumentative units is highly relevant [27]. Meta argumentation like *ad hominem* arguments and accusations of fallacies (especially *red herring*) is ubiquitous. In the particular case of implicit premises, the distinction between fallacies and possibly valid enthymemes becomes blurred – especially for automated analysis.

Formal Argumentation is concerned with representing and reasoning about argumentation in a machine understandable format. We refer to the *Handbook of Formal Argumentation* [3] for a general introduction, but a brief outline of the relevant aspects will be given in the following.

In RANT, formalisms for representing arguments extracted from our data are particularly important. As social media are a rich source of relevant metadata, formalisms based on the argument interchange format (AIF) like the social AIF (S-AIF) [16], AIF plus (AIF⁺) [21], and inference anchoring theory (IAT) [6] are particularly interesting because they can directly represent dialogue and speakers.

In the AIF and related formalisms, argumentation is represented as a graph with different types of nodes and corresponding meanings. Edges are also modelled as nodes in order to avoid the need for hyper-edges. The AIF was designed to represent monologues and later extended to AIFP where locutions are related through dialogue moves [18, 21] to incorporate dialogical argumentation. The extension to AIFP also brought support for representing IAT, i.e. to incorporate illocutionary force into argument maps for re-

lating locutions and illocutions² as well as inferences [20]. Furthermore, the SAIF incorporates the authors of locutions into the model for enhanced dialogue analysis and statistics.

We plan to utilise a logic reasoner as part of our argument map construction pipeline for automatically drawing some inferences from illocutions alone. For this aim we incorporate ideas from structured argumentation – such as labelling nodes with logic formulae. Besides other advantages, this will allow us to construct argument maps from a “soup of utterances” on Twitter, where extracting argumentative information from the dialog context is impossible.

Due to the peculiarities of argumentation on social media and especially on Twitter, existing argumentation formalisms require substantial adaptation and extension to be applicable in this context. For instance, most formalisms have no support for meta-argumentation like *ad hominem* arguments or accusations of fallacies. Especially the former phenomenon is very common on social media, although itself a well-known fallacy, and hence merits consideration in the choice of formalism. Moreover, only very few formalisms have a convenient representation of enthymemes.

2 Data and corpus queries

2.1 Data

Our dataset is based on roughly 23 million tweet IDs containing the string “Brexit” collected between May 5 and August 24, 2016.³ We downloaded all available tweets via the official Twitter API (which amounts to roughly 20 million tweets) and processed the dataset as follows.

Firstly, we only consider original tweets (no retweets), which amount to approximately half of all successfully downloaded tweets; the distribution of originals is shown in




Fig. 1 Number of original English tweets in our data base by day. The dashed bar indicates the day of the referendum (23rd June 2016); only tweets posted before the referendum are included in our corpus

¹ The size limit has recently been extended to 280 characters, but our corpus data pre-date this change.

² In speech act theory, the *locution* designates what is actually said, while the *illocution* is the aspect of the speech act as an action directed at causing a desired effect. In our case the desired effect would be to convince the other speakers of some point.

³ see <https://zenodo.org/record/263584/>.

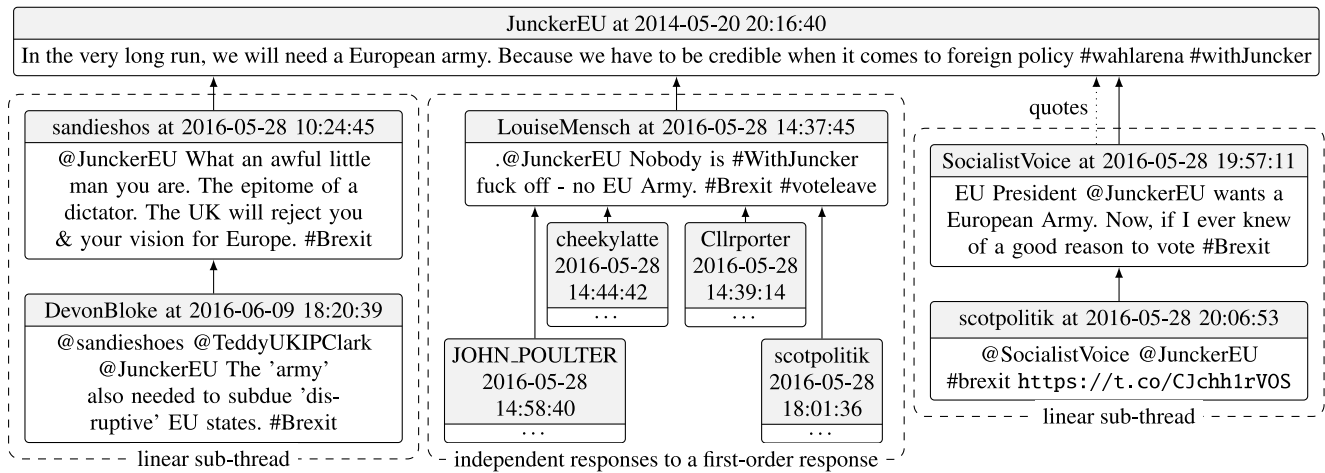


Fig. 2 Example reply thread with several linear sub-threads and independent responses

Fig. 1. Secondly, we only keep tweets identified as *English* by Twitter, which leaves more than 7 million of the original tweets; the most frequent other languages were Spanish (655,799), French (314,609), German (200,066), Italian (190,530), and *undefined* (160,756). Thirdly, for the development phase, we restricted the dataset to tweets posted before the referendum on June 23, 2016, aiming at a relatively consistent dataset and expecting substantial differences between arguments presented before and after the referendum. Our final implementation will later be applied to the full dataset, enabling the comparison of argumentation patterns over time.

In addition, we resolved reply-threads by retrieving all available tweets for which there is a reply in our dataset, but excluding non-English tweets. In this way, we can access dialogues between users, which are more likely to contain arguments. One example of such a reply thread is shown in Fig. 2. Note that our final database therefore contains tweets sent before May 5, 2016 and tweets that do not contain the search string “Brexit”. There are 215,744 reply-threads involving 688,905 tweets in our data set; 85% of them (183,188) could be resolved to their root.

Last but not least, we excluded near-duplicates (most of them likely generated by social bots – see [25] for details on our deduplication heuristics) as long as they are not part of a reply-thread. This way, we only consider genuine original content. Our final corpus consists of approximately 2.4 million tweets.

2.2 Linguistic Annotation

We used off-the-shelf software tools for tokenization and coarse POS tagging [19]⁴ and a custom lemmatizer based on

⁴ This POS tagger uses a set of 25 single-character tags specifically designed for tweets.

work by Minnen et al. [17]. POS taggers categorize words according to their parts of speech, e.g. *verb* or *noun*. Lemmatizers group together inflected forms of the same stem (e.g. *takes*, *took*, *taken* are all mapped to the lemma *take*). We additionally ran a tool for phrase chunking and named entity recognition (NER) [22, 23] – which also tags tokens with around 50 fine-grained POS tags following Penn Treebank style – combining the different tokenization layers in a post-processing step. Having a linguistically enriched data basis is an important prerequisite for formulating precise queries, cf. Sec. 2.3.

The annotated corpus has a total size of 32 million tokens. Unsurprisingly, the most frequent words forms are *Brexit* and the corresponding hashtag *#Brexit*, which together make up around 4% of all tokens, followed by function words such as determiners, punctuation marks, and prepositions. The most frequent content lemmas (after *brexit* and the auxiliary *be*) are *vote* and *eu* (both about 0.8% relative frequency).⁵ The coarse-grained POS system tagged approximately one third of all tokens as verb or noun, followed by proper nouns, prepositions, punctuation marks, determiners, adjectives, hashtags, URLs, pronouns, and adverbs. The NER system detected around 10 million noun phrases, 4 million verb phrases, 3 million prepositional phrases, and 2 million named entities. The annotation of a typical tweet is shown in Table 1; this tweet contains a match of the query presented in Sec. 4.⁶

⁵ Note that all lemmas have been converted to lower case.

⁶ Automatic linguistic annotation is prone to making errors, of course. In the example in Table 1, *Remain* has incorrectly been tagged as *V* and hence been included in the verb phrase. Such errors have to be taken into account when formulating queries.

Table 1 Linguistic annotation of a tweet sent on June 17, 2016 (ID: 743726246568173568), in vertical format. Tokenized text can be read from top to bottom in the first column. The following two columns represent the fine-grained and the coarse-grained POS tags, respectively. Lemmas are displayed in the last column. Additionally, phrase chunking and named entity/event recognition is displayed on the right-hand side; the last verb phrase was recognised as an event *vote*. The query match is indicated by a vertical bar on the left-hand side, anchor points and respective regions next to it, cf. Sec. 4

<i>token</i>	<i>POS</i>	<i>lemma</i>	
#Brexit	HT	#	#brexit
.	.	,	.
@0 As	IN	P	as
a	DT	D	a
@1 scientist	NN	N	scientist
.	.	.	.
@2 I	PRP	O	i
know	VBP	V	know
the	DT	D	the
immense	JJ	A	immense
value	NN	N	value
of	IN	P	of
@3 collaboration	NN	N	collaboration
.	.	.	.
So	RB	P	so
I	PRP	O	i
will	MD	V	will
vote	VB	V	vote
Remain	VB	V	remain
https[...]	UH	U	https://t.co/5iasjoJBMw

2.3 Query Architecture

Our corpus queries serve to extract argumentation from the data (cf. Sec. 4). Having queries as the central element of extraction allows us to combine lexical and grammatical patterns with word lists. At the same time, the formulation of explicit queries incorporating a fixed linguistic structure allows us to handle the noisy data prevalent in social media, as the queries can capture typical phenomena on the level of syntax, vocabulary and phraseology.

Our query architecture builds on the IMS Corpus Workbench (CWB) [10], a system designed for enabling complex linguistic searches on large corpora. The query language is based on regular expressions and allows for the incorporation of various levels of annotation. All grammatical information added to the corpus during preprocessing (cf. Sec. 2.2) can be accessed for each individual word or region – for instance, `[pos="N"]` will retrieve any word identified as a noun by the POS tagger, while `[lemma="have"]` finds all forms of *have* (*have*, *has*, *had*, *having*). Similarly, phrase chunks like `<np>...</np>` specify a sequence tagged as a noun phrase. These elements can be freely combined: `<np> [pos="N"] + </np>` matches a noun phrase consisting only of one or more nouns. For initial query development,

we used the web-based concordancing front-end CQPweb [15], allowing us to browse query results, view and sort context displays and perform statistical analyses.

To support the particular needs of RANT, we implemented our own python-wrapper around CWB⁷ and developed a bespoke web-application to manage word lists and queries, and to display results in a way tailored to the needs of argument extraction. In comparison to CQPweb, our app places its central focus on enabling the management of multiple query patterns rather than on individual queries and their statistical properties. This is achieved in particular by supporting complex macros and allowing the user to build and semi-automatically expand word lists.

3 Logical Models

3.1 Argument Representation

As mentioned in Sec. 1, argument representation formalisms need to be substantially amended to meet the challenges of social media argumentation. This led us to develop the trichotomic argumentation framework (T-AF) [14].

Part of a TAF graph (as depicted in Fig. 3) is essentially (up to the graphical presentation) an IAT/AIFP representation of a conversation, featuring hyper-edges labelled with inference schemes between nodes labelled with formulae in modal logic. We left out a further part containing the speakers and their relations, as this is currently not in the focus of our argument extraction pipeline.

These deviations mitigate the following problems posed by the challenges of social media argumentation mining and representation:

Enthymemes are incorporated by representing all relevant illocutions as logic formulae and having inferences between those when it is clear from context even though it might not be a full argument. This relies on the assumption that active participants of the discussion either understood the inference or would have asked for clarification, thereby leveraging the human inference capabilities lacking in an argument mining pipeline.

Argument schemes are annotated where recoverable as identifying such schemes can be beneficial when trying to recover the further structure of arguments [11].

Meta-arguments can be represented via the flexible inference edges that can attack entities (ad hominem arguments) and other relations (relevance attacks) besides illocutions.

⁷ <https://pypi.org/project/cwb-ccc/>.

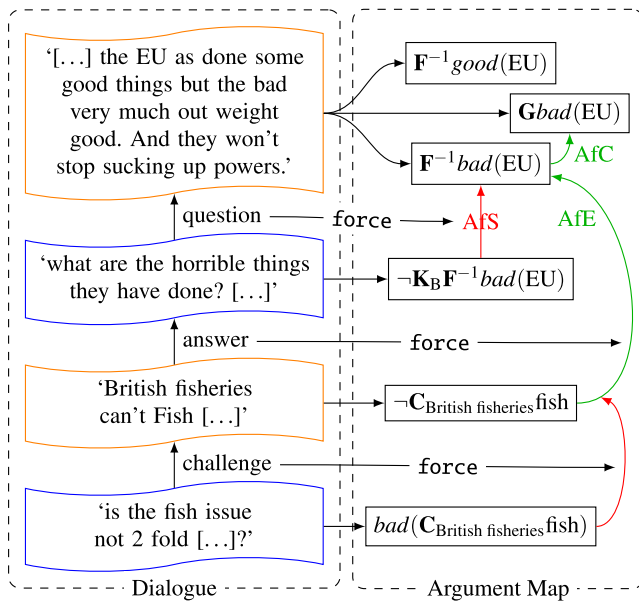


Fig. 3 Exemplary Twitter conversation on the Brexit referendum. *AfC* Argument from Commitment, *AfE* Argument from Evidence, *AfS* Argument from Source. Arguments are represented as modal formulae in suitable modal logics, in the present example featuring modalities F^{-1} at some point in the past, G always in the future, K_a agent a knows that, and C_a agent a can

Uncertainty is incorporated at multiple locations in our formalism. At the *locution* interpretation edges, we incorporate the case where locutions could be interpreted in different ways by linking to multiple alternative illocutions. At the *illocutions*, we can use similarity-based reasoning to reason about inferences in the presence of bad spelling, abbreviations, or different references to the same entities. Alternatively one could use suitable modalities for directly incorporating uncertainty in the formulae.

3.2 Coalgebraic Logic

As apparent in Fig. 3, the representation of real-world arguments involves a wide variety of modal operators. Recall that modal logics (including most current description logics [2]) are traditionally equipped with a relational semantics. This setup does fit some of the modal operators involved, such as the temporal operators featuring in Fig. 3. On the other hand, the semantics of modalities of knowledge and ability often lives outside the relational world, and involves, e.g., neighbourhood systems or game structures (e.g. [1, 26]). Moreover, relational semantics certainly does not suffice for uncertainty, vagueness, or defaults, which instead require probabilistic, fuzzy, or preferential structure in the semantics.

Coalgebraic logic has emerged as a common framework for logics featuring such extended modalities [7]. It is based on casting semantic models as *coalgebras* $X \rightarrow FX$ for

a given functor F in the paradigm of *universal coalgebra*, with the basic example of such a functor being the powerset construction $F = \mathcal{P}$ on sets, whose coalgebras reproduce the base case of relational systems. Further standard examples of coalgebras include probabilistic systems, preferential structures, and game structures. Coalgebraic logic thus supports a wide range of non-standard modalities such as *probably/with probability at least p*, *if – then normally*, or *X can enforce*. Modularity results in coalgebraic logic [24] allow for unrestricted combination of such modalities in the sense of a fusion of modal logics, in formulae such as $\text{ParliamentaryDemocracy} \Rightarrow [\text{Parliament}] \text{NewExecutive}$ ‘in a parliamentary democracy, parliament can normally force a change of executive’ (unless exceptional situations occur such as an irregular suspension of parliament). The above-mentioned *logical patterns* are thus essentially formulae in coalgebraic logic with placeholders for subformulae. Reasoning support is provided, either off-the-shelf or by easy implementation of further instance logics, via the generic *Coalgebraic Ontology Logic* (COOL) reasoner [12]. An important desideratum for further research is to provide support for *similarity reasoning* (e.g. [8]) at this level of generality, to ameliorate problems caused by deviations in vocabulary and phrasing.

4 Argument Extraction

Our current inventory consists of 25 formal patterns representing logical constituents in everyday argumentation, such as *position to know* arguments:

$$(?0 : \text{entity}) \Rightarrow \downarrow x. K_x (?1 : \text{entity})$$

Roughly speaking, the above example pattern states that ‘?0 are in a position to know ?1’, often phrased in the form ‘As a ?0, I know ?1’ (for brevity, we elide the implied claim ‘I am a ?0’.) Formally, the \Rightarrow connective is a default conditional ‘if – then normally’; the standard \downarrow -binder of hybrid logic binds the name x to the current individual (‘I’); and K_x is an epistemic modality, read ‘ x knows that’. The latter form of epistemic modality poses new challenges for logical reasoning, being indexed by a variable individual x . The example match (Table 1) to the query described next thus yields, slightly abbreviated, the complete formula $\text{scientist} \Rightarrow \downarrow x. K_x \text{value_of_collaboration}$. Because our logical formulae are more abstract than their linguistic representations, we have more queries than patterns, with a current set of 67 corpus queries. Queries are designed to capture as many instances of a given pattern in a particular linguistic context as possible, while maximising precision.

One of the queries associated with the pattern above, formulated in the syntax described in Sec. 2.3, is

```
[lemma = "as|being"]
@0:[::]
<np>
  []*
  [lemma = $common_people | lemma = $origin |
  lemma = $profession | lemma = $supporter]
  []*
</np>
@1:[::]
", "? [pos_ark = "O|N|Z"]
<vp>
  [pos_ner != "MD|TO"]*
  [lemma = $verbs_know]
  []*
</vp>
@2:[::]
( /region[vp] | /region[np] | /be_ap[] |
 /region[pp] | /region[advp] | /region[adjp] )+
@3:[::]
```

For each query, automatically annotated elements on the grammatical level are combined with custom-designed reusable macros and word lists. In the example, the entity claiming expertise is designated by a noun phrase (<np>) containing an element from one of four word lists describing persons. For instance, \$profession includes entries like *scientist*, *historian*, *economist* while \$common_people are generic person terms (*person*, *dude*, *gal*). pos_ner references the fine-grained part-of-speech annotation, allowing us to filter out modal verbs and infinitive markers in particular (MD, TO) while the tags assigned by pos_ark capture coarser categories (N = noun, Z = name, O = pronoun, ...). The macro /be_ap[] captures a form of *be* followed by an adjective phrase (*am sure*, *is certain*). Elements beginning with an @ followed by a number and a colon are target markers, which allow us to extract individual words or ranges (i.e. the words between two target markers), corresponding to a particular slot in the logical formula. In our example, entity ?0 is expressed by the region from @0: to @1: while ?1 is expressed by the region from @2: to @3:.

Evaluation In a preliminary evaluation study, we have manually checked the matches of a different query (associated with a concept similarity pattern) in order to assess our deduplication algorithm and to determine initial precision results. Automatic near-duplicate detection identified 33 duplicates among 99 matches; there were no false positives, but 12 false negatives. The large number of duplicate tweets skews the precision results to some extent: In the raw corpus, 17 out of the 99 matches were false positives (83% precision); in the manually deduplicated corpus, pre-

cision decreases to 72% (15 false positives out of 54 unique matches).

5 Conclusions

RANT follows a corpus-linguistic approach to extract argument patterns from a large dataset and transform matches and associated argumentation structure into a combined logical and argumentation-theoretic formalism whose development constitutes part of the project work. Ongoing work aims to extend the logical and structural expressiveness of the framework and to increase coverage and precision of the queries. An additional issue to be addressed is developing suitable evaluation measures; in particular, measuring recall remains challenging.

Funding Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alur R, Henzinger T, Kupferman O (2002) Alternating-time temporal logic. *J ACM* 49:672–713
2. Baader F, Calvanese D, McGuinness DL, Nardi D, Patel-Schneider PF (eds) (2003) *The description logic handbook*. Cambridge University Press, Cambridge
3. Baroni P, Gabbay D, Giacomin M, van der Torre L (eds) (2018) *Handbook of formal argumentation*. College Publications, London
4. Boltužić F, Šnajder J (2016) Fill the gap! Analyzing implicit premises between claims from online debates. *Argum Min* 3:124–133. <https://doi.org/10.18653/v1/W16-2815>
5. Bosc, T., Cabrio, E., Villata, S.: Tweeties squabbling: Positive and negative results in applying argument mining on social media. In: *Computational Models of Argument, COMMA 2016* (2016)
6. Budzynska, K., Reed, C.: Speech acts of argumentation: Inference anchors and peripheral cues in dialogue. In: *Computational Models of Natural Argument, CMNA 2011* (2011)
7. Cirstea C, Kurz A, Pattinson D, Schröder L, Venema Y (2011) Modal logics are coalgebraic. *Comput J* 54:31–41
8. Dubois D, Esteva F, Garcia P, Godo L, Prade H (1995) Similarity-based consequence relations. In: *Symbolic and quantitative approaches to reasoning and uncertainty ECSQARU 1995*. LNCS, vol 946. Berlin, Heidelberg, New York, Springer, pp 171–180
9. Dusmanu M, Cabrio E, Villata S (2017) Argument mining on Twitter: arguments, facts and sources. In: *Empirical Methods in Natural Language Processing EMNLP 2017*. ACL, Stroudsburg, pp 2317–2322

10. Evert S, Hardie A (2011) Twenty-first century corpus workbench: Updating a query architecture for the new millennium. *Corpus linguistics, CL* 2011. University of Birmingham, Birmingham
11. Feng VW, Hirst G (2011) Classifying arguments by scheme. In: Annual Meeting of the Association for Computational Linguistics ACL 2011. ACL, Stroudsburg, pp 987–996
12. Gorin D, Pattinson D, Schröder L, Widmann F, Wißmann T (2014) COOL – a generic reasoner for coalgebraic hybrid logics (system description). In: Automated Reasoning IJCAR 2014. LNCS, vol 8562. Springer, Cham, pp 396–402
13. Goudas T, Louizos C, Petasis G, Karkaletsis V (2014) Argument extraction from news, blogs, and social media. In: Artificial intelligence: methods and applications SETN 2014. Springer, Cham, pp 287–299
14. Göttlinger M, Schröder L (2018) Trichotomic Argumentation Representation. arXiv. 1812.06745v1. <https://arxiv.org/abs/1812.06745>
15. Hardie A (2012) Cqpweb – combining power, flexibility and usability in a corpus analysis tool. *Int J Corpus Ling* 17(3):380–409
16. Lawrence J, Snaith M, Konat B, Budzynska K, Reed C (2017) Debating technology for dialogical argument: sensemaking, engagement, and analytics. *ACM Trans Internet Technol* 17(3):1–23
17. Minnen G, Carroll J, Pearce D (2001) Applied morphological processing of English. *Nat Lang Eng* 7(3):207–223
18. Modgil S, McGinnis J (2008) Towards characterising argumentation based dialogue in the argument interchange format. In: Argumentation in multi-agent systems ArgMAS 2007. Springer, Berlin, Heidelberg, pp 80–93
19. Owoputi, O., O'Connor, B., Dyer, C., Gimpel, K., Schneider, N., Smith, N.: Improved part-of-speech tagging for online conversational text with word clusters. In: Human Language Technologies, HLT-NAACL 2013, pp. 380–390 (2013)
20. Reed, C., Wells, S., Budzynska, K., Devereux, J.: Building arguments with argumentation: the role of illocutionary force in computational models of argument. In: Computational Models of Argument, COMMA 2010, pp. 415–426 (2010)
21. Reed, C., Wells, S., Devereux, J., Rowe, G.: AIF+: Dialogue in the argument interchange format. In: Computational Models of Argument, COMMA 2008, pp. 311–323 (2008)
22. Ritter A, Clark S, Mausam, Etzioni O (2011) Named entity recognition in tweets: an experimental study. In: Empirical Methods in Natural Language Processing EMNLP 2011. ACL, Stroudsburg, pp 1524–1534
23. Ritter A, Mausam, Etzioni O, Clark S (2012) Open domain event extraction from Twitter. In: Knowledge Discovery and Data Mining KDD 2012. ACM, Stroudsburg, pp 1104–1112
24. Schröder L, Pattinson D (2011) Modular algorithms for heterogeneous modal logics via multi-sorted coalgebra. *Math Struct Comput Sci* 21(2):235–266
25. Schäfer F, Evert S, Heinrich P (2017) Japan's 2014 general election: Political bots, right-wing internet activism and PM Abe Shinzo's hidden nationalist agenda. *Big Data* 5(4):294–309
26. Vardi M (1989) On the complexity of epistemic reasoning. In: Logic in Computer Science LICS 1989. IEEE, pp 243–252
27. Wachsmuth H, Stede M, Baff RE, Khatib KA, Skeppstedt M, Stein B (2018) Argumentation synthesis following rhetorical strategies. In: Computational Linguistics COLING 2018. ACL, Stroudsburg, pp 3753–3765
28. Walton D, Reed C, Macagno F (2008) Argumentation schemes. Cambridge University Press, Cambridge