





ReCAP - Information Retrieval and Case-Based Reasoning for Robust Deliberation and Synthesis of Arguments in the Political Discourse*

Ralph Bergmann^(✉) , Ralf Schenkel^(✉) , Lorik Dumani^(✉) , and Stefan Ollinger^(✉) 

University of Trier, Behringstrasse 13, 54296 Trier, Germany

Abstract. The ReCAP project is a recently started project within the DFG priority programm robust argumentation machines (RATIO). It follows the vision of future argumentation machines that support researchers, journalistic writers, as well as human decision makers to obtain a comprehensive overview of current arguments and opinions related to a certain topic, as well as to develop personal, well-founded opinions justified by convincing arguments. Unlike existing search engines, which primarily operate on the textual level, such argumentation machines will reason on the knowledge level formed by arguments and argumentation structures. The focus of ReCAP is on novel contributions to and confluence of methods from information retrieval and knowledge representation and reasoning, in particular case-based reasoning. The aim is to develop methods that are able to capture arguments in a robust and scalable manner, in particular representing, contextualizing, and aggregating arguments and making them available to a user. Together with experts from the political domain real-world scenarios and use cases are worked out. A corpus of semantically annotated argumentations is being created from relevant text sources and will be made available to the argumentation research community.

Keywords: Argumentation · Information Retrieval · Case-Based Reasoning.

1 Introduction

Argumentation is a core aspect of everyday human life, e.g. in medicine, law and politics. On the one hand humans search for arguments to make good own decisions, and on the other hand they search or form own arguments to persuade others. However, in times of information age, big data, and fake news it is almost impossible to manually find all valid and relevant arguments for a certain topic.

An argument consists of a claim or standpoint supported or opposed by reasons or premises [19]. Argument components are usually expressed in natural

* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 375342983 - 2018-2020.

language. Together, the arguments form a graph or *argumentation structure*. *Deliberation* finds and weighs all arguments supporting or opposing some question or topic based on the available knowledge, e.g. by assessing their strength or factual correctness, to enable informed decision making, e.g. for a political action. *Synthesis* tries to generate new arguments for an upcoming topic based on transferring an existing relevant argument to the new topic and adapting it to the new environment.

This paper gives an overview over the recently started project ReCAP which is part of the DFG priority programm *robust argumentation machines* (RATIO)¹. It follows the vision of future *argumentation machines* that support political researchers and journalistic writers in deliberation and synthesis. Unlike existing search engines, which primarily operate on the textual level, such argumentation machines will reason on a knowledge level formed by (argumentative) propositions and argumentation structures. We propose a general architecture for an argumentation machine with focus on novel contributions to and confluence of methods from Information Retrieval (IR) and Knowledge Representation and Reasoning (RI), in particular Case-Based Reasoning. The argumentation machine works closely with argumentation structures in natural language and in order to achieve argumentative reasoning, it abstracts further away from the text by notions of similarity, extraction of facts, validation, clustering, generalization and adaptation of arguments, thereby offering some form of argument competency. Together with experts from the political domain we develop real-world scenarios which feedback into the development of the system. A corpus of high-quality argumentation structures for closely related topics is developed.

Section 2 describes the deliberation and synthesis use cases which drive the research in the project. Then, sect. 3 gives an overall overview of the project before sect. 4 presents the proposed methods that are in the focus of research. Section 5 presents first results of a workshop with our cooperation partners from the political domain. Finally, sect. 6 summarizes related work and sect. 7 concludes the paper.

2 Use Cases on Deliberation and Synthesis

We now illustrate the ReCAP vision by two exemplary use cases from the domain of political journalism. In the *deliberation use case*, we consider a journalist who attempts to write a survey article about a topic that is currently debated (or was debated in the past). Typical examples for such topics from the recent past could be the Brexit, accepting refugees in Europe, or a specific countermeasure against the subprime mortgage crisis. For her survey, she wants to collect arguments in favour of and against the topic, and she decides to search this information on the Web. As of now, she would have to manually collect large amounts of relevant documents (such as news articles, forum entries, blog posts, etc.) using a search engine, then manually extract arguments from the documents, and cluster similar arguments; maybe she also wants to rank the arguments such that

¹ <http://www.spp-ratio.de>

she can focus on the most important ones. The methods and tools that we will develop in the ReCAP project will provide strong support in this case, such that the journalist will be able to focus on journalistic aspects of the problem. The journalist needs to provide a textual description of the topic. The ReCAP argumentation machine will then automatically find documents with argumentative content where the topic is discussed and extract the (possibly complex) argumentation structure. It will then cluster similar arguments and argumentation structures, thus allowing for a concise overview of the discussion. Additional modules will help assessing the strength and the validity of arguments. As an output of the search, the system will present an aggregated view of arguments or even argumentation structures pro and con the topic, weighted by popularity, truthfulness or persuasiveness, with the option to drill-down and look at the (textual) sources of each argument. For a deeper analysis, our system can show, for each argument, how this argument is in turn supported or attacked by other arguments.

In the *synthesis use case*, the journalist attempts to forecast possible future discussions that may emerge about a topic that is just about to become important, based on similar discussions in the past. As a typical example, one could ask which of the arguments pro and con the Brexit would still apply for an exit of the Netherlands from the European Union. Instead of finding documents discussing this topic and extracting arguments from them, the journalist now needs to locate documents on similar topics (in our example, topics on the Brexit) and examine which arguments used there still apply, which need to be modified, and which cannot be applied in the new scenario. Again, the goal of the ReCAP project is to develop methods and tools that will support the journalist with such tasks. Given a topic, we will first determine similar topics based on argument similarity measures; in the example, this could be Brexit or Grexit. The ReCAP argumentation machine then finds documents discussing these related topics, extracts argumentation structures from them, and aggregates them. For each argument used for such a related topic, the system will then estimate if it is still valid in this context, if it needs to be adapted or replaced by an analogous one, or if it does not apply in this context at all and must be removed. As a result, a new argumentation is synthesised by reuse and analogical transfer of existing argument structures that particularly addresses the potential exit of the Netherlands from the EU.

3 Project Overview

The overall project vision is reflected in a preliminary view of the argumentation machine’s architecture, depicted in Fig. 1. This figure serves as an overview of the various basic research questions addressed in the ReCAP project as well as of their interrelationships. The bottom part of this layered architecture shows the textual level of the argumentation machine, addressing argument mining and corpus construction from existing textual sources, leading to semantically annotated argumentation graphs, reflecting the document’s content on the knowledge

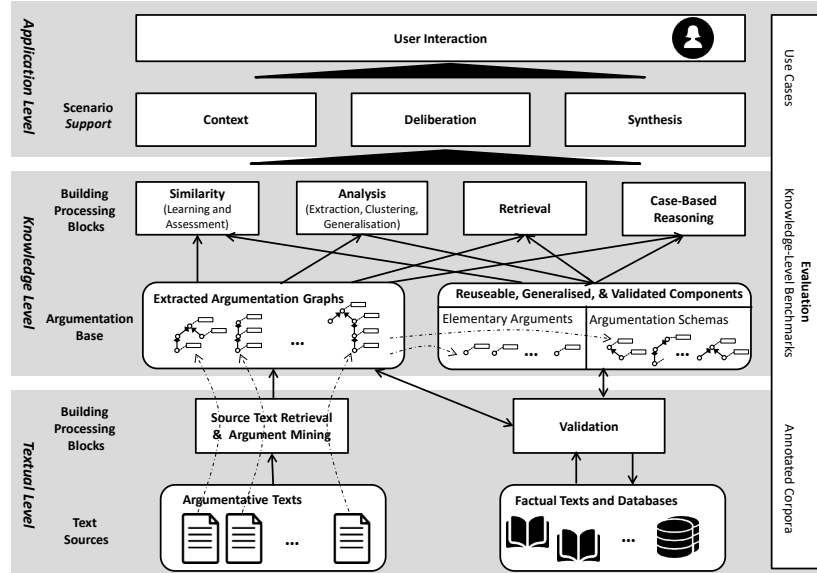


Fig. 1. Layered Architecture of the Envisioned Argumentation Machine

level. The validation of factual statements in arguments based on related text and databases leads to further enhancements of their representation on the knowledge level by including assessments of validity and strength as well as by links to the related textual evidences. The knowledge-level reasoning is positioned on top of the textual level. The extracted, specific argumentation graphs need to be analysed such that their major common constituents (elementary arguments, supporting and attacking relations, argumentation structures) are identified. Together with the specific argumentation graphs they will form the argumentation base for further knowledge processing. We will investigate new similarity measures for comparing arguments and argument structures, supported by machine learning methods for textual similarity. Such a computational notion of similarity is also the core for argumentation graph analysis. In addition argument (structure) retrieval, as required for finding suitable arguments, requires a notion of similarity to measure relevance of arguments. Finally, Case-Based Reasoning aims at supporting synthesis by designing new argument structures by adapting the best existing similar structures from the argumentation base by analogical transfer. The upper level of the architecture encompasses the specific application-oriented components to support deliberation and synthesis as required for the targeted use cases. Their implementation will make use of the knowledge processing building blocks. The context module aims at capturing, analysing, and representing the specific user's context, i.e., the specific issue under consideration as well as specific beliefs and constraints of the user.

We will focus our research in the first three years primarily on the knowledge-level. On the textual level, we address the collection of relevant evidences for validation. Research on argument mining will be deferred. Thus, the transformation of available text sources into semantically annotated argumentation graphs is performed manually. For this purpose, we identify relevant German language text sources from the political domain and manually annotate the argumentation structures that occur. Thereby a research corpus of high-quality is constructed (see sect. 5) which supports the evaluation of the developed methods.

4 Methods

We now give a brief overview of the methods being researched within the overall architecture of the argumentation machine.

4.1 Representation of Argumentation Structures

We will develop a model for representing argumentation structures and their components. An argument is a relationship between a claim and several premises, by which the claim is either attacked or supported. An argumentation structure consists of arguments, forming a graph. Claims and premises have a textual representation in natural language. Our model is essentially extending AIF [6] by support for storing specific meta information on claims, premises and their relationships (e.g. correctness and strength measures, who used the argument, when was it used, etc.) and for explicitly linking arguments and documents. For each argument, we will include provenance information in the form of its exact position in the source document, the annotator, and possibly any further meta information from the automated annotation process. Additionally, we will model argumentation schemes [28] and their instantiation in a concrete argumentation structure.

Our methods often work on the textual representation of premises and claims. A pre-processing step improves the understanding of a proposition in isolation by adding further annotations, e.g. by POS tagging, mapping to an ontology and resolution of anaphora. A model of user context represents the attitude towards certain arguments, documents, experts, and sources, similar to the context definition by Brewka & Eiter [8], thus allowing for a personalised argument ranking, retrieval, and adaptation. The context for a user group generalises this to groups of users (e.g. right-wing populists), allowing for a targeted analysis of the argumentation within a certain population.

4.2 Quantitative Properties of Arguments

The validity or factual correctness of a proposition estimates the degree of truth of the proposition. Validity measures will exploit the textual representation of a premise in two different ways. In a first line of work, we will develop methods

that evaluate the validity of a proposition by connecting it to a factual knowledge base, extending on existing work for RDF facts [11,17]. We will apply and modify pattern-based information extraction methods that extract facts to check from the proposition. If the fact is found in the knowledge base, this is a strong indication that the fact is correct. We will use the YAGO knowledge base [21] which has the advantage to include both multilingual information and temporal information on facts. As pure fact-checking against a knowledge base will only work for a subset of all facts used in argumentations, we plan to investigate an alternative approach that attempts to find the (possibly rephrased) proposition in a large text corpus or even on the Web and estimates its correctness based on frequency and/or authority of sources; the latter may depend on the context. We will build on existing methods such as those proposed by Leong and Cucerzan [16], which focus on reformulating a factual statement, not an argumentative proposition, and operate only on Wikipedia, exploiting specific properties like citations and inter-article linkage. In a next step, we will consider for all premises linked to a specific claim their strength or importance for the discussion, potentially relative to a given context. The goal here is to develop a ranking function. Popularity may be a reasonable start, but will usually not be enough since the strongest argument may not be given often, whereas a weak argument that is known to many people may be given frequently.

4.3 Similarity and Analysis of Argumentation Structures

Similarity is a core concept relevant when reasoning with argumentation structures extracted from text. Similarity measures will be considered for different purposes, in particular for similarity-based retrieval of graphs (see sect. 4.4) for deliberation as well as for CBR (see sect. 4.5) to support the synthesis of new argumentation graphs. Following the local-global principle established in CBR [23] the measure will be decomposed in local similarity measures assessing the similarity of all information available in the representation. Those measures will include combined structural and textual similarity measures [25,27], the semantic closeness of the arguments and relation types based on related ontological information [22] as well as the numerical similarity of certain attributes, such as validity and strength estimations. The global similarity is computed using an optimization process [4] which creates an alignment between the nodes and edges of the two argumentation graphs based on the local similarities. We will analyse the computational complexity of this optimisation problem and develop heuristic methods for finding good approximate solutions under acceptable time constraints. The developed measures will be evaluated in the context of their application and purpose based on systematically constructed ranking experiments with real users.

Building upon this research on similarity, we will further consider additional methods that support the analysis and decomposition of the argumentation graphs. We will develop clustering algorithms for arguments and argumentation graphs based on hierarchical clustering [14], such as divisive clustering using

k-medoid splitting, making use of the developed similarity measures. Generalisation can be performed without changing the structure of the graph by replacing one or several arguments by more general arguments [18]. The previously performed clustering provides the cluster label as a general argument for a cluster member. In case of hierarchical clusters, generalisation can be achieved on several levels of generality. Finally, we will identify frequently occurring elementary arguments, argumentation graphs and subgraphs are identified, extracted, generalised (if possible), and stored in a separate part of the argument base of reusable argumentation components.

4.4 Retrieval of Argumentation Structures

When exploring the arguments for some topic, a user may want to identify other topics where an argument was used, similar argumentation structures in different topics, or where a specific partial instantiation of an argumentation scheme is used. We will identify typical information needs based on discussions with experts. To enable such complex queries, we will define a query language based on a graph query language like SPARQL [20] that allows specification of strict constraints (e.g. a minimal strength), vague constraints (e.g. textual content of arguments) and constraints on the graph structure (e.g. relations, higher-order relations and scheme instantiations [28]). Since important information may be available only in the documents from which arguments were extracted/mined, but not in the extracted arguments, the language will also allow to specify vague and strict constraints on source documents. The language will also include means to refer to a predefined context. Scoring and ranking methods for the results will combine graph-based similarity measures (sect. 4.3) with content-based scores for conditions on documents and argument properties, building on our earlier work [10] for knowledge bases. We plan to apply standard learning-to-rank techniques that combine a large number of query-independent and query-dependent features. For training these models, we will develop an annotated corpus of structured queries together with relevant results, and we will develop tools for constructing these relevance assessments based on crowdsourcing. We will also develop an easy-to-use query interface with the option to explore detailed solutions, including faceted browsing.

4.5 Case-Based Reasoning with Argumentation Structures

The work on similarity and retrieval is extended towards a comprehensive CBR approach for the synthesis of argumentation graphs by reuse and adaptation of argumentation graphs and propositions from the argumentation base. This research builds upon previous work on CBR for legal argumentation ([7,24]) as well as on the previous work on process-oriented CBR ([5]). We will research on new adaptation methods that iteratively transform a retrieved argumentation graph towards an adapted argumentation graph that is better suited to the query than the original graph. In particular we will transfer the main concepts successfully developed for the compositional and operator-based adaptation of

workflow graphs [5] towards argumentation graphs. This includes methods for learning the required adaptation knowledge from the argumentation base. For learning adaptation operators, pairs of similar argumentation graphs from the argumentation base will be searched and compared (e.g. a mapping between the two graphs is constructed). The identified differences will be analysed and turned into a formal operator description that is able to bridge this difference between the two graphs. While the general principle underlying this learning approach is established in CBR [9], its application to argumentation graphs is absolutely novel. The adaptation process itself will then apply the learned adaptation knowledge iteratively on the retrieved argumentation graph, leading to a local search process. This search process (which can be implemented, for example, as a hill-climbing or a stochastic local search) aims at optimising two criteria in parallel: a) the similarity of the adapted argumentation graph to the current query (how well does the synthesised argumentation match the claim) and b) the validity and strength of it.

5 Workshop with cooperation partners

We organized a workshop with expert groups in the fields of journalistic writing (led by Dr. Damian Trilling, University of Amsterdam) and political research (led by Dr. Lasse Cronqvist, Trier University). The goal of this workshop was on the one hand to develop a comprehensive understanding of the problems and workflows of each occupational group in order to sketch visionary tools, and on the other hand to elaborate concrete use cases for deliberation and synthesis. For the latter, topics and sources of argumentative text had to be identified.

5.1 Operational workflows and potential tools

Nowadays, *journalism* can be divided into three main branches: *classic journalism*, *investigative journalism* and *online journalism*. Classic journalism describes a ground-in routine with superficial investigation and known sources where journalists are short of time and need to publish multiple articles per day; thus there is less need for argumentation tools. A journalist in investigative journalism deals with a subject rigorously. She concerns with phenomena in society, e.g. right-wing extremism, and needs to search forums, blogs, etc. There is a lot of data that needs to be extracted and there supporting tools could be helpful. In online journalism articles of print media are enriched e.g. with info graphics, headlines, etc. This work is mainly done by other journalists which are neither always familiar with the themes nor have much time to address the themes in depth. Thus supporting tools would be useful for these journalists, e.g. a deliberation tool that generates an overview of all arguments for or against a given topic.

A simplified example of such an overview is depicted in Fig. 2, showing arguments for and against merging the school forms *Hauptschule* and *Realschule* in Rhineland-Palantine, grouped by argument type. The Dutch company *Argumentenfabriek* created similar argument maps (mostly in Dutch, some in English) in a manual process [1]. A potential deliberation tool could generate such

argument maps automatically for arguments in German. It could also visualize the temporal dynamics by allowing to focus on arguments used in specific time frame.

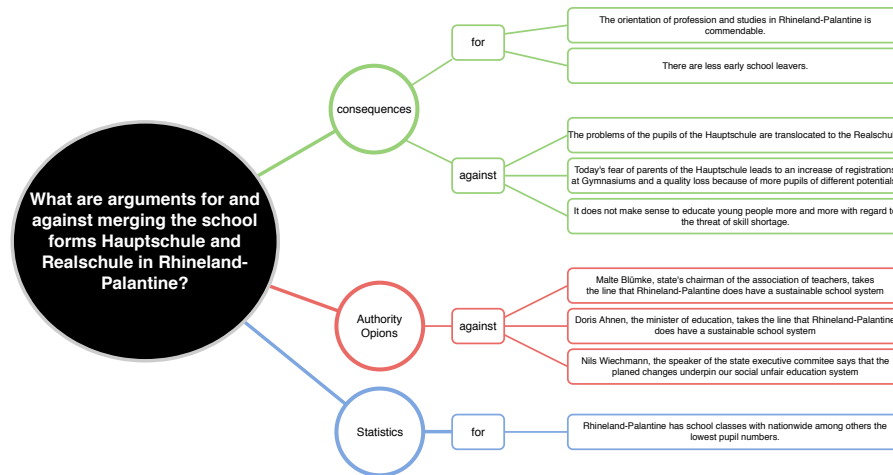


Fig. 2. Vision of an overview of arguments in an argument map

Contrary to journalists there does not exist a ground-in routine for *political researchers*. The way of working of a political researcher is sometimes very individual and depends on particular cases. Frequently, they work with annotating software like MaxQDA. They annotate 20 to 30 texts on a certain topic and then analyze passages of a certain argument type, focusing rather in the content of the text than the argument structure. The workshop with the political researchers confirmed that the tools envisioned with the journalist expert group would also be useful for these users. Both expert groups agreed that OVA [2] is a suitable tool to annotate arguments in texts.

5.2 Topics and Sources for Use Cases

The topics for use cases should be relevant to society and should not be too specific in order to find enough argumentative texts. To ensure a high-quality annotation with argument structures, typical texts discussing the topic should neither be too complicated for a non-expert nor be too simplistic, and the topic needs to have a manageable size in terms of the different arguments used. In addition to that, the meaning of typical concepts used should not change significantly over time. Since the focus of the project is on arguments in German, topics should be limited to Germany. Especially concerning the synthesis of arguments, information needs to be transferable between instances of topics.

We committed to the topic of *education policy* because it meets all these requirements. A core component of education policy is that it is under the responsibility of federal states. Arguments are clearly limited to states and therefore it is possible to transfer arguments between states. In order to have a diverse starting point, we decided to begin with three federal states (Rhineland-Palatinate, Hamburg, and Bavaria). Within this overall theme, a number of specific detailed topics were identified, e.g. the discussion of G8 vs. G9, the question if children should spend full days at school, or the debate if Hauptschule and Realschule should be merged.

For every topic, textual sources of arguments had to be identified and the argument structures within them manually annotated. Since the quality of arguments stands and falls with these sources we had to choose these with care. A good source are the protocols of plenary debates since they include different viewpoints, they are of medium complexity, and large volumes of them are freely available. Another central source are news articles, which also include different viewpoints, but are usually neither always open nor free to use. We will also consider information provided by political parties and lobby groups because these are usually rich of arguments.

6 Related Work

Today, many corpora with annotated argumentation already exist. Important sources for such corpora are aifdb.org (from University of Dundee) which provides argumentation structures in the Argument Interchange Format [6], the IBM Debating Technologies datasets based on Wikipedia articles, corpora provided by the UKP Lab in Darmstadt, and corpora built by the Applied CL Discourse Lab in Potsdam. However, none of these corpora is directly tied to specific usage scenarios and queries for deliberation and synthesis of arguments.

The validation of factual statements is an important problem not only in argumentation, but in many other fields. ClaimBuster [15] automatically finds, but does not check, facts in a debate that are worth checking. TruthTeller of the Washington Post (which is now offline) matched statements made by politicians in a speech to a database of pre-checked statements. PolitiFact (<http://www.politifact.com/>) performs thorough manual fact-checking for selected facts and rates the accuracy of statements. A number of rather limited methods for automatic validation of factual statements exist. Some approaches find, given an RDF triple, Web documents supporting this fact, converting it to various textual representations; recent examples include Defacto [11] and Multi-Verifier [29]. Validation of factual statements represented in textual form has seen much less work. Existing approaches essentially rely on paraphrasing a statement and finding it in a reference collection like Wikipedia [16]. Textual entailment [3] can be seen as a way of finding other statements that entail the statement in question. Habernal and Gurevych [13] estimate the convincingness of arguments using neural networks.

Stab et al. [26] present ArgumenText, an argument retrieval system capable of retrieving topic-relevant sentential arguments from a large collection of diverse Web texts for any given controversial topic. The system first retrieves relevant documents, then it identifies arguments and classifies them as “pro” or “con”, and presents them ranked by relevance in a web interface. Gutfreund et al. [12] introduce a system that automatically generates arguments supporting and contesting a given point of view about a controversial topic.

7 Conclusion

In this paper, we gave an overview of the main approach and the involved methods of the recently started project ReCAP that aims at developing an argumentation machine and related applications in the domain of politics. This project will be linked with the other projects in the priority programme RATIO. In particular, we aim at incorporating research results from projects which focus on argument mining during later phases of our roadmap for research. Thereby we aim at completing our vision of a full argumentation machine working without manual preparation of textual resources.

References

1. Example of an overview map of arguments from Argumentenfabriek. <https://www.argumentenfabriek.nl/media/2276/10032-tno-co2-engels.pdf>. Last accessed 8 June 2018.
2. OVA tool developed by ARG-tech. <http://ova.arg-tech.org/>. Last accessed 8 June 2018.
3. I. Androutsopoulos and P. Malakasiotis. A survey of paraphrasing and textual entailment methods. *J. Artif. Intell. Res. (JAIR)*, 38:135–187, 2010.
4. R. Bergmann and Y. Gil. Similarity assessment and efficient retrieval of semantic workflows. *Information Systems*, 40:115–127, 2014.
5. Ralph Bergmann and Gilbert Müller. Similarity-based Retrieval and Automatic Adaptation of Semantic Workflows. In *Synergies Between Knowledge Engineering and Software Engineering*, number 626 in Advances in Intelligent Systems and Computing, pages 31–54. Springer, 2017.
6. F. Bex, J. Lawrence, M. Snaith, and C. Reed. Implementing the argument web. *Commun. ACM*, 56(10):66–73, 2013.
7. L.Karl Branting. A reduction-graph model of precedent in legal analysis. *Artificial Intelligence*, 150(1):59 – 95, 2003.
8. G. Brewka and T. Eiter. Argumentation context systems: A framework for abstract group argumentation. In *Proc. 10th Int. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pages 44–57, 2009.
9. S. Craw, N. Wiratunga, and R. C Rowe. Learning adaptation knowledge to improve case-based reasoning. *Artificial Intelligence*, 170(16):1175–1192, 2006.
10. S. Elbassuoni, M. Ramanath, R. Schenkel, M. Sydow, and G. Weikum. Language-model-based ranking for queries on RDF-graphs. In *Proc. 18th ACM Conf. on Information and Knowledge Management*, pages 977–986, 2009.

11. D. Gerber, D. Esteves, J. Lehmann, L. Bühmann, R. Usbeck, A. Ngonga Ngomo, and R. Speck. Defacto - temporal and multilingual deep fact validation. *J. Web Sem.*, 35:85–101, 2015.
12. D. Gutfreund, Y. Katz, and N. Slonim. Automatic arguments construction — from search engine to research engine, 2016.
13. I. Habernal and I. Gurevych. Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM. In *ACL*, 2016.
14. J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 2006.
15. N. Hassan, B. Adair, J. Hamilton, C. Li, M. Tremayne, J. Yang, and C. Yu. The quest to automate fact-checking. In *Proc. of the 2015 Computation+Journalism Symposium*, 2015.
16. C. W. Leong and S. Cucerzan. Supporting factual statements with evidence from the web. In *CIKM*, pages 1153–1162, 2012.
17. Steffen Metzger, Shady Elbassuoni, Katja Hose, and Ralf Schenkel. S3K: Seeking statement-supporting top-k witnesses. In *Proc. 20th ACM Conf. on Information and Knowledge Management (CIKM)*, Glasgow, UK, 2011.
18. G. Müller and R. Bergmann. Generalization of Workflows in Process-Oriented Case-Based Reasoning. In *Proc. of the 28th FLAIRS Conference*, pages 391–396, Hollywood (Florida), USA, 2015. AAAI Press.
19. A. Peldszus and M. Stede. From argument diagrams to argumentation mining in texts: A survey. *Int. J. Cogn. Inform. Nat. Intell.*, 7(1):1–31, January 2013.
20. J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. *ACM Trans. Database Syst.*, 34(3), 2009.
21. T. Rebele et al. YAGO: A multilingual knowledge base from Wikipedia, Wordnet, and Geonames. In *Proc. 15th Int. Semantic Web Conference (ISWC)*, 2016.
22. P. Resnik. Using information content to evaluate semantic similarity in a taxonomy. In *Proceedings of the 14th Int. Joint Conf. on Artificial Intelligence*, pages 448–453. Morgan Kaufmann, 1995.
23. M. M. Richter and R. O. Weber. *Case-Based Reasoning - A Textbook*. Springer, 2013.
24. Edwina L. Rissland, Kevin D. Ashley, and Karl Branting. Case-based reasoning and law. *The Knowledge Engineering Review*, 20(3):293–298, 2005.
25. R. Schenkel, A. Theobald, and G. Weikum. Semantic similarity search on semistructured data with the XXL search engine. *Information Retrieval*, 8(4):521–545, December 2005.
26. C. Stab, J. Daxenberger, C. Stahlhut, T. Miller, B. Schiller, C. Tauchmann, S. Eger, and I. Gurevych. Argumenttext: Searching for arguments in heterogeneous sources. In *Proceedings of the 2018 Conference of the NAACL*, pages 21–25, 2018.
27. M. Theobald, H. Bast, D. Majumdar, R. Schenkel, and G. Weikum. TopX: Efficient and versatile top-k query processing for semistructured data. *VLDB Journal*, 17(1):81–115, January 2008.
28. D. N. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, 2009.
29. T. Wang, Q. Zhu, and S. Wang. Multi-verifier: A novel method for fact statement verification. *World Wide Web*, 18(5):1463–1480, 2015.