

# Report on the 1st Workshop on Argumentation Knowledge Graphs (ArgKG 2021) at AKBC 2021

Tirthankar Ghosal

ÚFAL, MFF, Charles University, CZ  
*ghosal@ufal.mff.cuni.cz*

Khalid Al-Khatib

University of Groningen, NL  
*Khalid.alkhatib@rug.nl*

Yufang Hou

IBM Research, Ireland  
*yhou@ie.ibm.com*

Anita de Waard

Elsevier, USA  
*a.dewaard@elsevier.com*

Dayne Freitag

SRI International, USA  
*freitag@ai.sri.com*

## Abstract

The first workshop on Argumentation Knowledge Graphs (ArgKG) was held virtually at the Automated Knowledge Base Construction (AKBC 2021) conference on October 7, 2021. ArgKG @ AKBC 2021 brought together the Computational Argumentation and Knowledge Graphs communities, aiming to promote cross-pollination of ideas and encourage discussions and collaborations between the two communities. This paper describes the workshop and compiles several of its findings and insights.

**Date:** 7 October, 2021.

**Website:** <https://argkg21.argmining.org>.

## 1 Introduction

Knowledge Graphs (KGs) have demonstrated remarkable effectiveness on various knowledge-intensive computational tasks in a range of domains, including journalism and science [Wang et al., 2019; Masoud et al., 2021]. Most work on the automatic construction of KGs has focused narrowly on extracting factual assertions from the source material (e.g., [Mondal et al., 2021; Bosselut et al., 2019]), ignoring an important dimension—argumentation. Many texts that communicate factual knowledge couch that knowledge in argument, in language that asserts, interprets, and reasons. Arguments are essential for forming stances on controversial topics, arriving at successful decisions in the face of uncertainty, and establishing consensus on problems requiring interpretation. The active research community of *computational argumentation* has investigated computational models for mining (e.g., [Fromm et al., 2021]), assessing (e.g., [Toledo et al., 2019]), and generating (e.g., [Alshomary et al., 2021]) textual argumentation in several domains [Lawrence and Reed, 2019; Al Khatib et al., 2021a]. Still, many important questions lack concrete answers, including

---

the best approaches to formalizing argumentation as a graph, deploying such graphs at scale, and integrating argumentative knowledge with factual knowledge typically found in KGs.

This inaugural workshop on Argumentative Knowledge Graphs (ArgKG) explored the automatic construction of KGs that encode argumentative knowledge and structures and incorporate argumentation into factual knowledge graphs, drawing together researchers focusing on NLP, automatic KG construction, and computational analysis of argumentation. The workshop established a basis for constructing argumentation knowledge graphs and highlighted noteworthy topics and research directions.

## 2 Submissions

As this was the first ArgKG workshop, we invited new papers and extended abstracts. The submissions were non-archival, and the authors had the flexibility to submit a final version of their paper to a different venue. Thanks to our program committee, each submission to ArgKG received at least two reviews to help the authors further refine their work for the next venue.

**Full Papers:** Papers were non-archival, with the flexibility to submit in any venue in addition to the ArgKG workshop. Accepted papers are listed on the workshop schedule, and available in OpenReview<sup>1</sup>. We invited papers of all types, including mature works, late-breaking results, work-in-progress, and position papers. Articles under review at other venues were also welcomed. Papers were restricted to 10 single-column pages, excluding references. Appendices were to be put after references and submitted in one PDF document. All submissions and supplementary materials had to be anonymized.

**Extended Abstracts:** We invited submissions of extended abstracts (2 pages maximum) related to the research topic. which could include previously published results, late-breaking results, and work in progress. Abstracts were lightly reviewed to ensure that the topic was within the scope of the workshop. The abstracts were non-archival and participants were free to submit their work for publication elsewhere.

We highlighted the following *topics of interest*:

- Argumentation knowledge graph (ArgKG) design, construction, and evaluation.
- Applications of ArgKG to different tasks such as question answering, machine-reading and reasoning, legal case analysis, search and retrieval, argument quality assessment, and argument generation.
- Systems and toolkits for building, maintaining, and using ArgKGs.
- Augmentation of traditional general-purpose KGs with argumentation knowledge.
- Relation between argumentation knowledge and other kinds of knowledge such as common sense.
- Visualization and exploratory analysis of ArgKG.
- Alternative (e.g., deep learning) representations of argumentation knowledge.

In addition to the regular call for new submissions, we invited some very relevant papers after careful scrutiny. We invited the corresponding authors to present their already published papers in ArgKG. We discuss those in Section 5.

---

<sup>1</sup><https://openreview.net/group?id=AKBC.ws/2021/Workshop/ArgKG>

---

## 3 Format of the Workshop

The ArgKG Workshop acted as a community-building exercise for communities interested in the intersection of Argumentation and Knowledge Graphs. We conducted the six + hour-long workshop virtually over Zoom with the breaks and social interactions over Gather Town<sup>2</sup>. We had three keynote talks (each of 30 minutes), one panel session (90 minutes), five invited papers (each of 20 minutes), and four presentations (10 minutes each) in ArgKG. The workshop schedule is here<sup>3</sup>. We witnessed enthusiastic participation, and at one point in time, around 50 + people joined the session.

## 4 Keynote Speakers

ArgKG featured three keynote speakers who had worked on computational argumentation spanning their careers.

### 4.1 [Keynote Talk 1] Reconstructing Arguments and Knowledge Claims – A Hermeneutical Perspective

**Abstract:** Argument analysis clarifies argumentative texts by reconstructing arguments in standardized ways (e.g., premise-conclusion structures, inference graphs, debate maps). Similarly, factual information in a corpus may be analyzed and represented as knowledge graphs. Argument analysis and knowledge reconstruction can be understood as text-interpretation tasks. From a methodological point of view, both tasks are closely intertwined: Argument reconstruction presupposes the explication of knowledge claims; and the reconstruction of knowledge claims requires the assessment of justificatory (i.e., argumentative) relations. This talk explores the methodological relation between argument mining and knowledge graph construction from the perspective of the humanities. It closes with an outlook on using pre-trained neural language models for these interpretative tasks.

**Speaker Bio:** Gregor Betz<sup>4</sup> is a professor of philosophy of science at the Karlsruhe Institute of Technology and coordinating KIT's DebateLab. He studies and applies normative models of reasoning and argumentation, with a focus on multi-agent argumentation. More specifically, Gregor has developed a theory for representing and evaluating controversial argumentation, contributed to the computational study of argumentative debate dynamics, and shown that common argumentative practices (e.g., critique, plurality, consensus) are truth-conducive. His theoretical investigations are the background for his more practical work on the analysis and evaluation of philosophical, scientific, and political debates. Gregor and his colleagues have developed software tools for argument analysis, especially Argunet and Argdown, and use these tools extensively in research, teaching, and scientific policy advice.

---

<sup>2</sup><https://gather.town/>

<sup>3</sup><https://argkg21.argmining.org/schedule.html>

<sup>4</sup>[https://www.philosophie.kit.edu/mitarbeiter\\_betz\\_gregor.php](https://www.philosophie.kit.edu/mitarbeiter_betz_gregor.php)

---

## 4.2 [Keynote Talk 2] Conclusion Synthesis: A Tale of Two Towers

**Abstract:** In this talk, I focus on the task of conclusion synthesis, which consists in (automatically) generating a conclusion given a set of premises. I contrast different approaches to conclusion synthesis ranging from purely knowledge-based through to purely language model based. The knowledge-based approach generates conclusions from a knowledge base by using a template-based approach. The benefit of this method is that it generates conclusions that can be interacted with and scrutinized by a human expert that can challenge the arguments and perform sensitivity analysis to understand how premises systematically relate to the conclusions. The drawback is the need to have formalized knowledge in the application domain, medical decision-making in our example case. The purely language model based approach uses a generative model to infer conclusions from (textual) premises and does not require formalized knowledge and can be applied to any domain. However, the process for generating conclusions is not transparent in the sense that it is unclear how a system derives conclusions from premises. The argumentation process lacks accountability and does not support critical scrutiny as well as sensitivity analysis by a human counterpart. We draw some conclusions from our observations regarding the potential for bringing the discussed methods together.

**Speaker Bio:** Philipp Cimiano<sup>5</sup> is a full professor for semantic computing at Bielefeld University. He graduated in computer science from Stuttgart University and obtained his Ph.D. and Habilitation from the Karlsruhe Institute of Technology. His area of expertise is in the field of knowledge-based systems, knowledge representation and management, as well as natural language processing. He has done seminal work and published a number of monographs on knowledge acquisition and ontology learning from text as well as ontology-driven interpretation of natural language. He is currently coordinating the priority program RATIO “Robust Argumentation Machines” funded by the German Science Foundation.

## 4.3 [Keynote Talk 3] Modeling Knowledge in Computational Argumentation

**Abstract:** Our society is overloaded with information, and people are getting more passive and dependent on information delivered by technologies. In order to build computational models that support human decision-making in a truthful way, such models should be able to handle knowledge. In this talk, I’m going to focus on two problems in computational argumentation—argument mining and counter evidence retrieval—and try to answer what knowledge is useful to solve these problems and how to incorporate that knowledge into computational models.

**Speaker Bio:** Yohan Jo<sup>6</sup> worked on natural language processing for various aspects of dialogue in domains like argumentation, education, and clinical notes. His interests lie in modeling human reasoning in conversation, and his Ph.D. thesis focuses on computational analysis and generation of argumentation. Yohan recently joined Amazon for the Alexa AI team as an applied scientist.

---

<sup>5</sup><http://www.sc.cit-ec.uni-bielefeld.de/team/philipp-cimiano/>

<sup>6</sup><https://yohanjo.github.io>

---

## 5 Invited Papers

We shortlisted a set of very relevant papers that would aptly fit the theme of ArgKG and stimulate discussions in the workshop. We then reached out to the authors. The following are the papers whose authors agreed to present and participate in the ensuing discussions.

- **Employing Argumentation Knowledge Graphs for Neural Argument Generation** by [Al Khatib et al. \[2021b\]](#) published in ACL-IJCNLP 2021. Here the authors build several knowledge graphs and utilize their compositions to be encoded in a set of arguments. The arguments and encoded knowledge are used to fine-tune a general text generation model, driving the model to deliver good-quality arguments. Authors Khalid Al-Khatib and Yufang Hou, who are incidentally one of the organizers of ArgKG, were the presenters.
- **Argument Invention from First Principles** by [Bilu et al. \[2019\]](#) published in ACL 2019. In this work, the authors explicitly defined a taxonomy of principled recurring arguments and, given a controversial topic, automatically identified which of these arguments were relevant to the topic. The authors claim that this is for the first time an approach to argument invention is formalized and made explicit in the context of NLP. Co-author Ariel Gera from IBM, Israel, presented the paper.
- **Syntopical Graphs for Computational Argumentation Tasks** by [Barrow et al. \[2021\]](#) published in ACL-IJCNLP 2021. In this work, the authors introduce the *syntopical graph*, a data structure for linking claims within a collection, inspired by syntopical reading. This reading process emphasizes comparing and contrasting viewpoints to improve topic understanding. The authors define a syntopical graph as a typed multi-graph where nodes represent claims and edges represent different possible pairwise relationships, such as entailment, paraphrase, or support. Author Joe Barrow from the University of Maryland, who interned at Adobe Research while carrying out this work, was the presenter.
- **Covid-on-the-Web: Knowledge Graph and Services to Advance COVID-19 Research** by [Michel et al. \[2020\]](#) published in ISWC 2020. In this paper, the authors discuss their Covid-on-the-Web project that aims to allow biomedical researchers to access, query, and make sense of COVID-19 related literature. Their project combines and extends tools to process, analyze and enrich the COVID-19 Open Research Dataset (CORD-19) [[Wang et al., 2020](#)]. The resource comprises two main knowledge graphs describing (1) named entities mentioned in the CORD-19 corpus and linked to DBpedia, Wikidata, and other BioPortal vocabularies, and (2) arguments extracted using ACTA [[Mayer et al., 2019](#)], a tool automating the extraction and visualization of argumentative graphs, meant to help clinicians analyze clinical trials and make decisions. Santiago Marro and Aline Menin from I3S Lab, INRIA, CNRS, were the presenters.
- **Use of Claim Graphing and Argumentation Schemes in Biomedical Literature: A Manual Approach to Analysis** by [Moser and Mercer \[2020\]](#) published in the 7th ArgMining workshop at COLING 2020. In this work, the authors report on an investigation of the large-scale argumentation structure found when examining five biochemistry journal publications. Their investigation suggests that argumentation schemes originally designed for genetic research articles may transfer to experimental biomedical literature in general. The authors present two novel models: Claim Graph and Model of Informational Hierarchy (MIH), to better understand the large-scale argumentation structure of a com-

---

plete biochemistry paper. Co-author Robert Mercer from the University of Western Ontario presented their work.

## 6 Panel Discussion

We had a stimulating panel discussion on the topic **Computational Argumentation and Knowledge Graphs: Recipe for miracles?** with five panel members. The objective of the panel was to discuss some critical aspects for modeling computational argumentation in the light of knowledge graphs. The panel session was moderated by Anita de Waard. Our panel members were:

- **Davide Ceolin**<sup>7</sup> obtained his Ph.D. at the Vrije University Amsterdam. He is currently a tenure-track researcher in the Human-Centered Data Analytics group at Centrum Wiskunde & Informatica (CWI) in Amsterdam and a guest researcher for the User-Centric Data Science group at the Vrije Universiteit Amsterdam. His research focuses on the transparent assessment of multidimensional online information quality and has been published in diverse venues like WebSci, EKAW, IJDQ, IP&M. His research encompasses the use of argumentation theory, logical methods, and crowdsourcing to better and more extensively assess the quality of online information. He is the leader of the starting NL eScience Center eTec Eye of the Beholder project, where he will investigate how to make AI pipelines for information quality assessment transparent.
- **Nancy Green**<sup>8</sup> is Associate Professor Emerita of the Department of Computer Science, University of North Carolina Greensboro in the U.S. Her current research interests are in logic-based computational models of natural argument, rhetoric, scientific discovery, and AI ethics.
- **Iryna Gurevych**<sup>9</sup> (PhD 2003, U. Duisburg-Essen, Germany) is professor of Computer Science and the director of the Ubiquitous Knowledge Processing (UKP) Lab at the Technical University (TU) of Darmstadt in Germany. She joined TU Darmstadt in 2005 (tenured as full professor in 2009). Her main research interests are in machine learning for large-scale language understanding, text semantics, and argument mining. Iryna's work has received numerous awards. Examples are: ACL fellow 2020, and the first Hessian LOEWE Distinguished Chair (2,5 mil. Euro) in 2021. Currently, Iryna is the president of SIGDAT, and she is the co-director of the ELLIS NLP program. She has been PC co-chair of ACL 2018 and has been elected to be the future president (2023) of the International Association for Computational Linguistics (ACL).
- **Simone Teufel**<sup>10</sup> is Professor of Language and Information at the Department of Computer Science and Technology at Cambridge University. After achieving her first degree in Computer Science from the University of Stuttgart, she did her Ph.D. research at the Cognitive Science in Edinburgh University, on the topic of scientific rhetoric detection and

---

<sup>7</sup><https://www.cwi.nl/people/davide-ceolin>

<sup>8</sup><https://compsci.uncg.edu/faculty/green/>

<sup>9</sup>[https://www.informatik.tu-darmstadt.de/ukp/ukp\\_home/head\\_ukp/index.en.jsp](https://www.informatik.tu-darmstadt.de/ukp/ukp_home/head_ukp/index.en.jsp)

<sup>10</sup><https://www.cl.cam.ac.uk/~sht25/>

---

summarization. She spent two years as a postdoc at Columbia University working on medical information extraction. Since 2001, she has been a permanent member of staff at the Department of Computer Science and Technology (formerly Computer Laboratory). She is also a frequent collaborator at the Tokyo Institute of Technology. Her research interests include argumentation mining, reasoning, discourse processing, and all aspects of human evaluation of NLP tasks.

- **Yuxiao Ye**<sup>11</sup> is a Ph.D. candidate in the Department of Computer Science and Technology at the University of Cambridge, supervised by Prof. Simone Teufel. Yuxiao has two years of industrial research experience at Alibaba, working on knowledge-based question answering. He obtained his MSc in Speech and Language Processing at the University of Edinburgh. He entered a Math degree at Tongji University at age 13, before reorienting himself to Foreign Language and Literatures at Tsinghua University and graduating in 2016. His current research is on multi-party informal argumentation.

The panelists started by discussing how their work related to knowledge graphs and then collectively started to define what an argumentation knowledge graph is. Collectively, they discussed four different layers in argumentation knowledge graphs:

- a basic layer which formed by concepts, linked by arcs
- a next layer, namely that of predicates
- a next layer involved the discourse structure
- the final layer represents argument quality, strength, and direction

All agreed that there are no good methodologies to uniformly move from text to predicates, and tools are necessary to generate such Knowledge Graphs.

A further agreement was on annotations for a task; instead, we need to find an efficient way to gather large-scale datasets for a large number of tasks in different domains of argumentation. Apart from learning from text alone, it was considered essential to explore learning from the environment since argumentation is essentially about communication and persuasion.

A key is the ability to model the belief system and background of the arguers to understand better how logically unsound arguments are proposed and accepted during discussions. The importance of modeling values in argumentation was stressed. It was pointed out that often participants disagree with each others' values instead of the factuality of the various statements. A great goal would be to help people to understand each other better by showing the inference process of each side which considers values and other dimensions of argumentation.

## 7 Accepted Papers and Abstracts

We accepted the following papers/extended abstracts via Open Review<sup>12</sup>, which was the submission platform for AKBC 2021<sup>13</sup>.

---

<sup>11</sup><https://yeyuxiao.github.io/>

<sup>12</sup><https://openreview.net/group?id=AKBC.ws/2021/Workshop/ArgKG>

<sup>13</sup><https://www.akbc.ws/2021/>

- 
- **Improving Quality Assessment of Online Reviews Using Formal Argumentation Theory and Knowledge Graphs**<sup>14</sup> In this extended abstract, the author introduce a previously developed framework for the assessment of the quality of product reviews by using formal argumentation theory and machine learning. Author Davide Ceolin, who was also one of our panelists, presented this work.
  - **Annotating Implicit Reasoning in Arguments with Causal Links**<sup>15</sup> In this full paper, the authors propose a semi-structured template to represent argumentation knowledge that explicates the implicit reasoning in arguments via causality. They create a novel two-phase annotation process with simplified guidelines and show how to collect and filter high-quality implicit reasonings via crowdsourcing. Author Keshav Singh presented their paper.
  - **My research relevant to the Argument Knowledge Graph workshop**<sup>16</sup> Nancy Green, who was one of our panelists and a veteran researcher on computational argumentation, presented her several research projects over the years: arguments represented in propositional logic generated from knowledge bases, GenIE project, scientific argumentation in genetics research articles, argumentation schemes by inductive logic programming, etc.
  - **ArguMeet: An Argument Diagramming Schema for Meeting Conversations**<sup>17</sup> Here in this work, the authors propose an argumentation diagramming schema which extracts argumentative portions from meeting conversations to create a visual representation that provides a high-level view of the argumentation that has taken place during the meeting resulting in a set of possible action points. Author Saprativa Bhattacharjee presented their work.

As mentioned, the accepted submissions were non-archival.

## 8 Discussion and Future Directions

Our workshop was hosted by AKBC, a forum fostering research into the automated construction of “conventional” knowledge graphs, the focus of which is often to capture ontological and factual knowledge expressed in text corpora. A recurring theme in our discussion was the significance of that work for argument mining. For some participants, work on knowledge graphs offers a useful structuring metaphor; argumentation is intrinsically relational, and a logical near-term focus is the assimilation of individual argumentative relations into a graph as a way to summarize the arguments found in a document or, more interestingly, a set of documents belong to some technical focus. A separate question is whether factual knowledge graphs provide value for argumentation analysis. The clear consensus is that they can, in principle, supply background knowledge and argumentative context, information critical to a correct rendering of a text’s arguments, but the noise and incompleteness typical of automatically constructed knowledge graphs poses a continuing challenge.

Much of the current work on knowledge graphs focus on using large-scale language models (LSLMs), such as BERT, as a potential substitute or supplement for explicitly constructed

---

<sup>14</sup><https://openreview.net/forum?id=7T8xeZYGCKA>

<sup>15</sup><https://openreview.net/forum?id=GFa0tnhAe-W>

<sup>16</sup><https://tinyurl.com/argkg2021-nancy-green>

<sup>17</sup><https://openreview.net/forum?id=JcIVRT5TozL>

---

graphs. The ArgKG community is actively investigating this option, but the place of LSLMs in the argumentation analysis remains to be determined. Some workshop participants maintained that because argumentation is tied closely to reasoning, these opaque models can never provide the basis of a complete solution. However, the methodology LSLMs embody, with their attention to multiple modes of communication in the performance of multiple tasks, presumably *is* the wave of the future. Background knowledge is ultimately critical to computational treatments of argumentation, which are therefore subject to the classic knowledge acquisition bottleneck. Several participants cited the ease with which children acquire much of this knowledge through mechanisms that are presumably more consonant with the training of LSLMs.

Argument mining addresses a number of use cases, and it borders on a number of adjacent problems that could be scoped in or out of the discipline. Argumentation analysis often seeks to assess reliability and persuasion, but these concepts involve phenomena that are not strictly logical and are challenging to model. How one judges the ultimate quality of an argument depends partly on the author's identity (or perceived tribal affiliation), one's own tribal affiliation, the success with which the author invokes certain frames or evokes shared values, etc. These effects are most visible in popular communication, but even scientific discourse is subject to similar dynamics. For some workshop participants, these phenomena must be sequestered to understand an argument's logical structure. Others expressed enthusiasm for modeling these phenomena to pursue a more complete account of concepts such as argument reliability. However, it was noted that work on computational treatments of the relevant factors, such as frames or ethos, is still in its infancy. Given the enthusiastic participation and the community interest, we look forward to the next iteration of ArgKG in 2022, with more stimulating sessions on the synergy of computational argumentation with knowledge graphs.

## 9 Organizers

The ArgKG workshop was organized by:

- **Tirthankar Ghosal**<sup>18</sup> is a researcher at the Institute of Formal and Applied Linguistics, Charles University Prague, Czech Republic. His main research interests are NLP/ML for Scientific Discourse Processing and Peer Reviews, Text/Dialogue Summarization, Argumentation Mining.
- **Khalid Al-Khatib** is an Assistant Professor in Natural Language Processing at the University of Groningen. His main research interests are argumentation mining and generation, discourse analysis, and bias analysis.
- **Yufang Hou** is a research scientist at IBM Research Ireland. Her main research interests include discourse analysis, argument mining, IE, and summarization from the scientific literature.
- **Anita de Waard** is VP of Research Collaborations, where her work focuses on working with academic and industry partners on projects pertaining to progressing modes and frameworks for scholarly communication. De Waard has worked on discourse analysis of scientific narratives, with an emphasis on finding key epistemic components in biological text.

---

<sup>18</sup><https://member.acm.org/~tghosal/>

- 
- **Dayne Freitag** is a Program Director in the Artificial Intelligence Center at SRI International. His research interests include multimodal information extraction from the scientific literature.

## Acknowledgements

We thank the organizers of AKBC 2021, especially the workshop chairs, Waleed Ammar and Eunsoo Choi, for their support to host ArgKG, providing with the necessary infrastructure, publicity, complimentary registrations for the organizers and our invited speakers, panelists. We also thank our program committee for helping us carve an exciting inaugural ArgKG program. Our program committee members were:

- Gully Burns, Chan Zuckerberg Initiative, US
- Philipp Cimiano, Bielefeld University, Germany
- Sourish Dasgupta, RAXter Inc, India
- Jennifer D’Souza, TIB Hannover, Germany
- Paul Groth, University of Amsterdam, Netherlands
- Anne Lauscher, University of Mannheim, Germany
- Jodi Schneider, University of Illinois at Urbana Champaign, US
- Muskaan Singh, Charles University, Czech Republic
- George Tsatsaronis, Elsevier, Germany
- Serena Villata, CNRS, Italy
- Henning Wachsmuth, Paderborn University, Germany
- Ellie Young, Knowledge Graph Conference

## References

- Khalid Al Khatib, Tirthankar Ghosal, Yufang Hou, Anita de Waard, and Dayne Freitag. Argument mining for scholarly document processing: Taking stock and looking ahead. In *Proceedings of the Second Workshop on Scholarly Document Processing*, pages 56–65, Online, June 2021a. Association for Computational Linguistics. doi: 10.18653/v1/2021.sdp-1.7. URL <https://aclanthology.org/2021.sdp-1.7>.
- Khalid Al Khatib, Lukas Trautner, Henning Wachsmuth, Yufang Hou, and Benno Stein. Employing argumentation knowledge graphs for neural argument generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4744–4754, Online, August 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.366. URL <https://aclanthology.org/2021.acl-long.366>.
- Milad Alshomary, Wei-Fan Chen, Timon Gurcke, and Henning Wachsmuth. Belief-based generation of argumentative claims. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 224–233, Online, April

- 
2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.17. URL <https://aclanthology.org/2021.eacl-main.17>.
- Joe Barrow, Rajiv Jain, Nedim Lipka, Franck Dernoncourt, Vlad Morariu, Varun Manjunatha, Douglas Oard, Philip Resnik, and Henning Wachsmuth. Syntopical graphs for computational argumentation tasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1583–1595, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.126. URL <https://aclanthology.org/2021.acl-long.126>.
- Yonatan Bilu, Ariel Gera, Daniel Hershcovich, Benjamin Sznajder, Dan Lahav, Guy Moshkovich, Anael Malet, Assaf Gavron, and Noam Slonim. Argument invention from first principles. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1013–1026, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1097. URL <https://www.aclweb.org/anthology/P19-1097>.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. COMET: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1470. URL <https://aclanthology.org/P19-1470>.
- Michael Fromm, Evgeniy Faerman, Max Berrendorf, Siddharth Bhargava, Ruoxia Qi, Yao Zhang, Lukas Dennert, Sophia Selle, Yang Mao, and Thomas Seidl. Argument mining driven analysis of peer-reviews. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):4758–4766, May 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16607>.
- John Lawrence and Chris Reed. Argument mining: A survey. *Comput. Linguistics*, 45(4):765–818, 2019. doi: 10.1162/coli\_a\_00364. URL [https://doi.org/10.1162/coli\\_a\\_00364](https://doi.org/10.1162/coli_a_00364).
- Maraim Masoud, Bianca Pereira, John McCrae, and Paul Buitelaar. Automatic Construction of Knowledge Graphs from Text and Structured Data: A Preliminary Literature Review. In Dagmar Gromann, Gilles Sérasset, Thierry Declerck, John P. McCrae, Jorge Gracia, Julia Bosque-Gil, Fernando Bobillo, and Barbara Heinisch, editors, *3rd Conference on Language, Data and Knowledge (LDK 2021)*, volume 93 of *Open Access Series in Informatics (OASISs)*, pages 19:1–19:9, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-199-3. doi: 10.4230/OASISs.LDK.2021.19. URL <https://drops.dagstuhl.de/opus/volltexte/2021/14555>.
- Tobias Mayer, Elena Cabrio, and Serena Villata. Acta a tool for argumentative clinical trial analysis. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6551–6553. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/953. URL <https://doi.org/10.24963/ijcai.2019/953>.

- 
- Franck Michel, Fabien Gandon, Valentin Ah-Kane, Anna Bobasheva, Elena Cabrio, Olivier Corby, Raphaël Gazzotti, Alain Giboin, Santiago Marro, Tobias Mayer, Mathieu Simon, Serena Villata, and Marco Winckler. Covid-on-the-web: Knowledge graph and services to advance covid-19 research. In Jeff Z. Pan, Valentina Tamma, Claudia d’Amato, Krzysztof Janowicz, Bo Fu, Axel Polleres, Oshani Seneviratne, and Lalana Kagal, editors, *The Semantic Web – ISWC 2020*, pages 294–310, Cham, 2020. Springer International Publishing. ISBN 978-3-030-62466-8. URL [https://link.springer.com/chapter/10.1007/978-3-030-62466-8\\_19](https://link.springer.com/chapter/10.1007/978-3-030-62466-8_19).
- Ishani Mondal, Yufang Hou, and Charles Jochim. End-to-end construction of NLP knowledge graph. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1885–1895, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.165. URL <https://aclanthology.org/2021.findings-acl.165>.
- Eli Moser and Robert E. Mercer. Use of claim graphing and argumentation schemes in biomedical literature: A manual approach to analysis. In *Proceedings of the 7th Workshop on Argument Mining*, pages 88–99, Online, December 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.argmining-1.10>.
- Assaf Toledo, Shai Gretz, Edo Cohen-Karlik, Roni Friedman, Elad Venezian, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. Automatic argument quality assessment - new datasets and methods. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5625–5635, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1564. URL <https://aclanthology.org/D19-1564>.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, Kathryn Funk, Rodney Kinney, Ziyang Liu, William Merrill, Paul Mooney, Dewey A. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex D. Wade, Kuansan Wang, Chris Wilhelm, Boya Xie, Douglas Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. CORD-19: the covid-19 open research dataset. *CoRR*, abs/2004.10706, 2020. URL <https://arxiv.org/abs/2004.10706>.
- Peilu Wang, Hao Jiang, Jingfang Xu, and Qi Zhang. Knowledge Graph Construction and Applications for Web Search and Beyond. *Data Intelligence*, 1(4):333–349, 11 2019. ISSN 2641-435X. doi: 10.1162/dint\_a\_00019. URL [https://doi.org/10.1162/dint\\_a\\_00019](https://doi.org/10.1162/dint_a_00019).