

The Meant, the Said, and the Understood: Conversational Argument Search and Cognitive Biases

Johannes Kiesel*
Bauhaus-Universität Weimar
Weimar, Germany
johannes.kiesel@uni-weimar.de

Henning Wachsmuth
Paderborn University
Paderborn, Germany
henningw@upb.de

Damiano Spina*
RMIT University
Melbourne, Australia
damiano.spina@rmit.edu.au

Benno Stein
Bauhaus-Universität Weimar
Weimar, Germany
benno.stein@uni-weimar.de

ABSTRACT

Many questions of public interest do not have a single answer but come with a set of choices, each of which with its pros and cons. An “objective” information system can help explore the underlying argument space, and, if equipped with a conversational interface, it can create the experience of lively discussions resembling those from our daily lives. However, users will (subconsciously) extend the provided information by assumptions that adhere to their cognitive biases. In this regard, note that biases do not arise only from the underlying data or the employed algorithms, but also from the way the information is presented—especially in audio-only channels. Our paper brings attention to bias-related challenges of conversational interfaces for argument search systems. We identify research questions that address these challenges, and we propose ideas and methods to tackle them.

CCS CONCEPTS

• **Human-centered computing** → **Interaction paradigms**; • **Information systems** → **Users and interactive retrieval**.

KEYWORDS

argument search, cognitive bias, conversational search, intelligent assistants, selection bias

ACM Reference Format:

Johannes Kiesel, Damiano Spina, Henning Wachsmuth, and Benno Stein. 2021. The Meant, the Said, and the Understood: Conversational Argument Search and Cognitive Biases. In *3rd Conference on Conversational User Interfaces (CUI '21)*, July 27–29, 2021, Bilbao (online), Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3469595.3469615>

1 INTRODUCTION

Conversational user interfaces of intelligent assistants, such as Amazon Alexa, Google Assistant, and Apple Siri, have the potential to address complex information needs even for controversial

*Johannes Kiesel and Damiano Spina contributed equally to the paper.

CUI '21, July 27–29, 2021, Bilbao (online), Spain



This work is licensed under a Creative Commons Attribution International 4.0 License.
ACM ISBN 978-1-4503-8998-3/21/07.
<https://doi.org/10.1145/3469595.3469615>

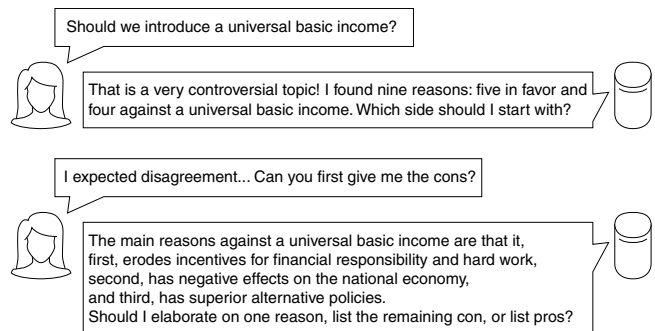


Figure 1: Start of a conversational argument search session with a fictitious intelligent assistant (as imagined in prior work [16]) on the topic of a universal basic income.

topics, but are at the moment mostly limited to answer with facts expressed in a few words. However, intelligent assistants that answer queries on controversial topics could be a solid asset to fight misinformation, disinformation, and societal polarization by exposing the background and reasoning behind different views [13, 18]. This scenario is called conversational argument search [16, 28].

Unfortunately, there is a large discrepancy between the amount of information an intelligent assistant can convey and the exploration a controversial topic demands. At first glance, the fictitious dialog in Figure 1 seems to provide an adequate overview of the topic discussed. However, as illustrated in Figure 2,¹ the topic’s structure is way more complex than suggested by the dialog.

Additionally, it is critical to understand that the way intelligent assistants expose information has a significant impact on how users perceive the information [5, 7]. Already the decision which view of a controversial topic to present an argument for first has biasing effects an assistant should compensate for. Otherwise, uncompensated subconscious influences bear a high risk, as illustrated by a proposal for an upcoming EU Regulation for Artificial Intelligence: “The prohibitions covers practices that have a significant potential to manipulate persons through subliminal techniques beyond their consciousness [...] in order to materially distort their behavior in a manner that is likely to cause them or another person psychological or physical harm” [11].

¹Source: <https://www.kialo.com/should-there-be-a-universal-basic-income-ubi-1634>




Figure 2: Kialo’s discussion topology for “Should There be a Universal Basic Income (UBI)?” Each pro (green) and con (red) is represented as a ring segment. Pros and cons of these are on the next outer ring, and so on. Hovering over a segment shows the corresponding premise and its “impact,” as voted by the Kialo community, in the form of one bar.

Motivation. As conversational and voice interfaces are becoming more habitual, we assume that conversational argument search will become an attractive feature of intelligent assistants [16]. Unfortunately, our current understanding of conversational argument search, as well as of providing complex answers via intelligent assistants in general, is insufficient. However, we argue that if one could design a conversational interface that makes the complexities of the topics and of its decisions transparent to the user and that compensates for the user’s cognitive biases, one could create an effective tool to help users comprehend the complexities that many of today’s questions of public interest have. To this end, this paper aims to identify open challenges that the research community needs to address to better understand how to design conversational argument search systems, particularly to make them bias-aware.

2 BACKGROUND

The paper at hand discusses the implications of combining argument search [33] and conversational search [4] concerning biases and—by extension—fairness, accountability, confidentiality, and transparency (FACT) [20]. We argue that these issues cannot be addressed through algorithm design and data collection alone, but require a careful interface design, as well.

2.1 Conversational Argument Search

Conversational argument search systems provide their users with arguments in favor of and against user-specified statements by means of a conversation. The first argument search systems appeared recently [31, 33], presenting both pros and cons instead of a direct answer [22], but still resembling the classical search engine interface. To provide for an interaction closer to human discussions, related research suggests employing either voice search [16] or avatars [35]. Even beyond, IBM’s Project Debater [27] is the first system able to engage in a formal debate with humans. However,

whereas participants in a debate aim to convince their audience, the users of an argument search system aim to inform themselves.²

A major challenge of conversational argument search systems is summarizing the pros and cons for a quick overview [16] (recall Figure 2). One approach is to visualize the distribution of arguments according to certain aspects, such as the argument’s topics [1]. Other possible aspects could be the argument’s moral sentiment [17], the facets examined in comparative arguments [26], or the types of evidence given [23]. Other approaches include clustering and summarizing the arguments by a handful of key points [6] and providing an overview of an argumentative discourse via a structure-based comparison [15]. How to embed such overviews into a conversation, especially a voice-based one, has barely been investigated yet [16].

People searching for meaning [24] might be strongly attracted to argument search, particularly to a conversational interface that allows for vague questions, implying responsibility of the system for fairness and transparency [2]. It is still unclear how to select arguments for a fair treatment [20] of views, groups, and users. In many cases, absolute fairness will be impossible, especially with a narrow channel such as voice, but also due to biases in the data (e.g., social biases [30]). Transparency [2] is essential in such a situation, that is, the argument selection should be explained and gaps should be indicated. However, research on how to increase the transparency of conversational argument search systems is still lacking.

2.2 Bias in Conversational Information Access

Cognitive biases, which are systematic patterns that may make users deviate from rational decision making, have been analyzed at great length in psychology (cf. Kahneman [14]). Among these, *anchor bias* refers to relying too strongly on the first piece of information about a topic encountered. On the other hand, *availability bias* refers to overestimating the importance of information that comes to one’s mind more easily.³ Clearly, these and other biases have implications on how to present arguments in a fair manner in conversational interfaces.

Prior research on the effects of cognitive biases—negative as well as positive ones—exists for both information access and conversational systems, but little for systems at their intersection. For information access, Azzopardi [5] provides an overview by surveying different studies that analyzed cognitive biases across different domains and search process steps. Different cognitive biases influence people’s information-seeking and retrieval behaviors in decision-making processes. For instance, Novin and Meyers [19] observed four cognitive biases of students researching controversial science topics with a screen-based search engine: priming, anchoring, framing, and availability. As exemplary research for conversational systems, Santhanam et al. [25] observed anchor bias in judgments for dialog systems evaluations.

Different paths exist for alleviating this issue, but these are barely explored so far. Conceptually, Gerritse et al. [12] distinguish three strategies for conversational systems to deal with biases they detect

²One reason for argument search could be to prepare for a debate, though.

³Given the cognitive load and memory load of *listening* to information in an audio channel, the anchor bias effect may also be observed in the opposite direction: users may get anchored to the latest information they listened to. Likewise, the effect of availability bias may also be more prominent, as the higher cognitive load of listening causes earlier pieces of information to be less prevalent in memory.




Figure 3: Overview of the proposed research questions. Read arrows as “How can the <entity 1> <verb> the <entity 2>?”

in the results or in a user’s questions: ignore, tell the user, or provide options to counterweigh/diversify. In some cases, though, it may be sufficient to change the presentation of results, for example, by utilizing the intelligent assistants’ voice, speech style, and linguistic attributes [9] or by modifying conversational style (e.g., high consideration vs. high involvement [32]). The effect of such changes in the presentation on the different cognitive biases is still unclear.

3 CHALLENGES AND OPPORTUNITIES

The challenges toward bias-aware conversational argument search are related, on the one hand, to users and their cognitive bias, and, on the other hand, to intelligent assistants and their argumentative selection bias [30]. Figure 3 illustrates the different interactions among these entities relevant to conversational user interfaces. From these, we derive five main research questions along with a (non-exhaustive) list of possible action steps listed in Table 1. We detail the steps below, and we discuss the opportunities they entail.

RQ1. How can the user and the assistant understand each other? Cooperative conversations require the conversants to establish a common ground for the conversation, known as grounding [8], also in argument search. One part of grounding for the assistant is either to provide personalization options or to adapt automatically to each user’s knowledge and beliefs. The other part for the assistant is to detect and react to the user’s mental model of itself. For both parts, novel experimental methodologies and a shift towards longitudinal studies are needed to capture both short- and long-term effects of users’ cognitive biases on their behavior. Approaches for all mentioned issues likely require gathering personal data, making it necessary to develop privacy-aware interaction guidelines [3].

RQ2. How can the assistant explain its argumentative selection bias? Systems should be accountable for informing the user about their different argumentative selection biases. This accountability involves characterizing biases in a transparent and intelligible way. For instance, the assistant could employ measures on media bias from journalism and media communication [10, 29] to provide the user with a sense of the political views covered—and missing—in its response. Further investigation is also needed on how to communicate such measurements most effectively: up-front, with each piece of information, as a summary, or yet in a different manner?

Table 1: Research questions (RQs) derived from the arrows in Figure 3 along with proposed action steps to address them.

RQ	Action Step
1. How can the user and the assistant understand each other?	<ul style="list-style-type: none"> - Investigate on short/long-term effects and mental models - Develop privacy-aware interaction guidelines
2. How can the assistant explain its argumentative selection bias?	<ul style="list-style-type: none"> - Identify intuitively understandable bias categories - Investigate how to make bias explicit
3. How can the user control the assistant’s argumentative selection bias?	<ul style="list-style-type: none"> - Identify cue phrases that specify argumentative selection biases - Investigate on personas for different argumentative selection biases
4. How can the assistant compensate for the user’s cognitive biases?	<ul style="list-style-type: none"> - Investigate strategies to encourage users to explore - Identify conversation styles that least provoke cognitive biases
5. How can the assistant help the user identify their cognitive biases?	<ul style="list-style-type: none"> - Identify hints at the application of cognitive biases - Identify strategies to explain the users their cognitive biases

RQ3. How can the user control the assistant’s argumentative selection bias? Once users understand the argumentative selection bias, the ability to control the bias would provide them with another means for steering the conversation. But what are intuitive interaction methods for such control? One avenue for research are cue phrases to ask for argumentative selection biases explicitly, such as “what would a right-wing politician say?” Another one is to develop different argumentative personas, such as the right-wing politician, and to allow the user to choose whom to talk to. At its extreme, this approach would allow for a conversation with several different personas at once where the user would be the moderator.

RQ4. How can the assistant compensate for the user’s cognitive biases? Conversely, the assistant needs to factor in the user’s cognitive biases when providing arguments. For instance, one could investigate how to compensate for biases that lead to premature assumptions by encouraging users to explore the arguments both more broadly (e.g., different views) and more deeply (e.g., counter-arguments [34]). More subtle, conversational styles (e.g., whether to barge in) and communication strategies likely affect cognitive biases, though it is still unclear how and to what extent. Measurements and experimental instrumentation from fields such as cognitive science and social science may inform the design of corresponding experiments. In particular, such research is relevant for the auditing of intelligent assistants in terms of not being detrimental in decision-making processes (cf. the regulation proposal mentioned in Section 1).

RQ5. How can the assistant help the user identify their cognitive biases? Another way to minimize the adverse effects of cognitive biases is to help the user identifying them. Such help requires the assistant to be able to detect the application of biases within the conversation. Similarly to RQ2, the assistant needs to create awareness of these biases in an intuitive and intelligible way. On a positive note, users might be particularly open to being educated on biases during argument search. Moreover, even generic reminders could make users more attentive [21]. As with RQ1, longitudinal

studies seem necessary to better understand the long-term effects of such communication strategies.

4 CONCLUSION

In this thought-provoking paper, we have argued that more efforts are needed for understanding how to design and develop effective conversational argument search systems. In particular, more attention needs to be paid to the implications that conversational interfaces have in argumentative contexts. This is because (i) intelligent assistants can barely retrieve information without a selection bias and, (ii) they may create or reinforce undesirable cognitive biases on the user's side. We have identified relevant research questions and action steps to address current challenges in designing bias-aware conversational argument search systems.

Many of the action steps that we propose will probably provide insights beyond conversational argument search. Cognitive biases will affect the reception of complex information via a narrow (audio-only) communication channel in most scenarios. Audio, in turn, may be the only channel available for certain individuals, such as people from low-literacy communities or the visually impaired. For example, how would a conversation sound like in order to verify the latest rumor about the efficacy of a COVID-19 vaccine effectively? As for this and similar tasks, research on bias-aware conversational argument search can shed light on critical recurring aspects of conversational interfaces for information-seeking.

ACKNOWLEDGMENTS

Damiano Spina is the recipient of an Australian Research Council DECRA Research Fellowship (DE200100064).

We would like to thank the anonymous reviewers for their valuable feedback. We would also like to thank the participants of the Future Conversations workshop at CHIIR 2021 [28] for the valuable discussions, which helped consolidate the ideas in this paper.

REFERENCES

- [1] Yamen Ajjour, Henning Wachsmuth, Dora Kiesel, Patrick Riehmman, Fan Fan, Giuliano Castiglia, Rosemary Adejoh, Bernd Fröhlich, and Benno Stein. 2018. Visualization of the Topic Space of Argument Search Results in args.me. In *2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018) - System Demonstrations*. ACL, 60–65. <http://aclweb.org/anthology/D18-2011>
- [2] James Allan, Jaime Arguello, Leif Azzopardi, Peter Bailey, Tim Baldwin, Krisztian Balog, Hannah Bast, Nick Belkin, Klaus Berberich, Bodo von Billerbeck, Jamie Callan, Rob Capra, Mark Carman, Ben Carterette, Charles L. A. Clarke, Kevyn Collins-Thompson, Nick Craswell, W. Bruce Croft, J. Shane Culpepper, Jeff Dalton, Gianluca Demartini, Fernando Diaz, Laura Dietz, Susan Dumais, Carsten Eickhoff, Nicola Ferro, Norbert Fuhr, Shlomo Geva, Claudia Hauff, David Hawking, Hideo Joho, Gareth Jones, Jaap Kamps, Noriko Kando, Diane Kelly, Jaewon Kim, Julia Kiseleva, Yiqun Liu, Xiaolu Lu, Stefano Mizzaro, Alistair Moffat, Jian-Yun Nie, Alexandra Olteanu, Iadh Ounis, Filip Radlinski, Maarten de Rijke, Mark Sanderson, Falk Scholer, Laurianne Sitbon, Mark Smucker, Ian Soboroff, Damiano Spina, Torsten Suel, James Thom, Paul Thomas, Andrew Trotman, Ellen Voorhees, Arjen P. de Vries, Emine Yilmaz, and Guido Zuccon. 2018. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum* 52, 1 (Aug. 2018), 34–90. <https://doi.org/10.1145/3274784.3274788>
- [3] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Founrey, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, 1–13. <https://doi.org/10.1145/3290605.3300233>
- [4] Avishek Anand, Lawrence Cavedon, Hideo Joho, Mark Sanderson, and Benno Stein. 2020. Conversational Search (Dagstuhl Seminar 19461). *Dagstuhl Reports* 9, 11 (2020), 34–83. <https://doi.org/10.4230/DagRep.9.11.34>
- [5] Leif Azzopardi. 2021. Cognitive Biases in Search: A Review and Reflection of Cognitive Biases in Information Retrieval. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval (CHIIR '21)*. ACM, 27–37. <https://doi.org/10.1145/3406522.3446023>
- [6] Roy Bar-Haim, Lilach Eden, Roni Friedman, Yoav Kantor, Dan Lahav, and Noam Slonim. 2020. From Arguments to Key Points: Towards Automatic Argument Summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. ACL, 4029–4039. <https://doi.org/10.18653/v1/2020.acl-main.371>
- [7] Carlos Castillo. 2019. Fairness and Transparency in Ranking. *SIGIR Forum* 52, 2 (Jan. 2019), 64–71. <https://doi.org/10.1145/3308774.3308783>
- [8] Herbert H. Clark and Susan E. Brennan. 1991. Grounding in Communication. In *Perspectives on Socially Shared Cognition*. American Psychological Association, 127–149.
- [9] Leigh Clark. 2016. *Exploring Vague Language Use and Voice Variation in Human-Agent Interaction*. Ph.D. Dissertation. University of Nottingham.
- [10] Robert M. Entman. 2007. Framing Bias: Media in the Distribution of Power. *Journal of communication* 57, 1 (2007), 163–173.
- [11] European Commission. 2021. Proposal for a Regulation laying down harmonised rules on artificial intelligence. <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence>. Accessed: 25-05-2021.
- [12] Emma J. Gerritse, Faegheh Hasibi, and Arjen P. de Vries. 2020. Bias in Conversational Search: The Double-Edged Sword of the Personalized Knowledge Graph. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval (ICTIR '20)*. ACM, 133–136. <https://doi.org/10.1145/3409256.3409834>
- [13] Andrew Guess, Brendan Nyhan, and Jason Reifler. 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. *European Research Council* 9, 3 (2018), 4.
- [14] Daniel Kahneman. 2012. *Thinking, fast and slow*. Penguin, London.
- [15] Dora Kiesel, Patrick Riehmman, Henning Wachsmuth, Benno Stein, and Bernd Fröhlich. 2021. Visual Analysis of Argumentation in Essays. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 1139–1148. <https://doi.org/10.1109/TVCG.2020.3030425>
- [16] Johannes Kiesel, Kevin Lang, Henning Wachsmuth, Eva Hornecker, and Benno Stein. 2020. Investigating Expectations for Voice-Based and Conversational Argument Search on the Web. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval (CHIIR '20)*. ACM, 53–62. <https://doi.org/10.1145/3343413.3377978>
- [17] Jonathan Kobbe, Ines Rehbein, Ioana Hulpuş, and Heiner Stuckenschmidt. 2020. Exploring Morality in Argumentation. In *Proceedings of the 7th Workshop on Argument Mining*. ACL, 30–40. <https://www.aclweb.org/anthology/2020.argmining-1.4>
- [18] Devi Mallal and Damiano Spina. 2020. How will verified fact-check content be identifiable on conversational devices? <https://virtual.globalfact7.com/agenda/how-will-verified-fact-check-content-be-identifiable-on-conversational-devices>. Online workshop at the GlobalFact7 Virtual Conference. Accessed: 25-05-2021.
- [19] Alamir Novin and Eric Meyers. 2017. Making Sense of Conflicting Science Information: Exploring Bias in the Search Engine Result Page. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval (CHIIR '17)*. ACM, 175–184. <https://doi.org/10.1145/3020165.3020185>
- [20] Alexandra Olteanu, Jean Garcia-Gathright, Maarten de Rijke, Michael D. Ekstrand, Adam Roegiest, Aldo Lipani, Alex Beutel, Alexandra Olteanu, Ana Lucic, Ana-Andreea Stoica, Anubrata Das, Asia Biega, Bart Voorn, Claudia Hauff, Damiano Spina, David Lewis, Douglas W. Oard, Emine Yilmaz, Faegheh Hasibi, Gabriella Kazai, Graham McDonald, Hinda Hamed, Iadh Ounis, Ilse van der Linden, Jean Garcia-Gathright, Joris Baan, Kamuela N. Lau, Krisztian Balog, Maarten de Rijke, Mahmoud Sayed, Maria Panteli, Mark Sanderson, Matthew Lease, Michael D. Ekstrand, Preethi Lahoti, and Toshihiro Kamishima. 2021. FACTS-IR: Fairness, Accountability, Confidentiality, Transparency, and Safety in Information Retrieval. *SIGIR Forum* 53, 2 (March 2021), 20–43. <https://doi.org/10.1145/3458553.3458556>
- [21] Gordon Pennycook, Ziv Epstein, Mohsen Mosleh, Antonio A. Arechar, Dean Eckles, and David G. Rand. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature* 592 (Mar 2021), 590–595. <https://doi.org/10.1038/s41586-021-03344-2>
- [22] Martin Pothast, Matthias Hagen, and Benno Stein. 2021. The Dilemma of the Direct Answer. *SIGIR Forum* 54, 1, Article 14 (Feb. 2021), 12 pages. <https://doi.org/10.1145/3451964.3451978>
- [23] Rutu Rinott, Lena Dankin, Carlos Alzate Perez, Mitesh M. Khapra, Ehud Aharoni, and Noam Slonim. 2015. Show Me Your Evidence – an Automatic Method for Context Dependent Evidence Detection. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. ACL, Lisbon, Portugal, 440–450. <https://doi.org/10.18653/v1/D15-1050>
- [24] Ian Ruthven. 2019. Making Meaning: A Focus for Information Interactions Research. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval (CHIIR '19)*. ACM, 163–171. <https://doi.org/10.1145/3295750.3298938>
- [25] Sashank Santhanam, Alireza Karduni, and Samira Shaikh. 2020. Studying the Effects of Cognitive Biases in Evaluation of Conversational Agents. In *Proceedings*

- of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). ACM, 1–13. <https://doi.org/10.1145/3313831.3376318>
- [26] Matthias Schildwächter, Alexander Bondarenko, Julian Zenker, Matthias Hagen, Chris Biemann, and Alexander Panchenko. 2019. Answering Comparative Questions: Better than Ten-Blue-Links?. In *2019 Conference on Human Information Interaction and Retrieval (CHIIR 2019)*. ACM, 361–365. <https://doi.org/10.1145/3295750.3298916>
- [27] Noam Slonim, Yonatan Bilu, Carlos Alzate, Roy Bar-Haim, Ben Bogin, Francesca Bonin, Leshem Choshen, Edo Cohen-Karlik, Lena Dankin, Lilach Edelstein, Liat Ein-Dor, Roni Friedman-Melamed, Assaf Gavron, Ariel Gera, Martin Gleize, Shai Gretz, Dan Gutfreund, Alon Halfon, Daniel Hershcovich, Ron Hoory, Yufang Hou, Shay Hummel, Michal Jacovi, Charles Jochim, Yoav Kantor, Yoav Katz, David Konopnicki, Zvi Kons, Lili Kotlerman, Dalia Krieger, Dan Lahav, Tamar Lavee, Ran Levy, Naftali Liberman, Yosi Mass, Amir Menczel, Shachar Mirkin, Guy Moshkovich, Shila Ofek-Koifman, Matan Orbach, Ella Rabinovich, Ruty Rinott, Slava Shechtman, Dafna Sheinwald, Eyal Shnarch, Ilya Shnayderman, Aya Soffer, Artem Spector, Benjamin Sznajder, Assaf Toledo, Orith Toledo-Ronen, Elad Venezian, and Ranit Aharonov. 2021. An Autonomous Debating System. *Nature* 591, 7850 (01 Mar 2021), 379–384. <https://doi.org/10.1038/s41586-021-03215-w>
- [28] Damiano Spina, Johanne R. Trippas, Paul Thomas, Hideo Joho, Katriina Byström, Leigh Clark, Nick Craswell, Mary Czerwinski, David Elswiler, Alexander Frummet, Souvick Ghosh, Johannes Kiesel, Irene Lopatovska, Daniel McDuff, Selina Meyer, Ahmed Mourad, Owoicho Paul, Sachin Pathiyan Cherumanal, Daniel Russell, and Laurianne Sitbon. 2021. Report on the Future Conversations Workshop at CHIIR 2021. *SIGIR Forum* 55, 1 (2021), 22 pages.
- [29] Timo Spinde, Felix Hamborg, Karsten Donnay, Angelica Becerra, and Bela Gipp. 2020. Enabling News Consumers to View and Understand Biased News Coverage: A Study on the Perception and Visualization of Media Bias. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020 (JCDL '20)*. ACM, 389–392. <https://doi.org/10.1145/3383583.3398619>
- [30] Maximilian Spliethöver and Henning Wachsmuth. 2020. Argument from Old Man's View: Assessing Social Bias in Argumentation. In *Proceedings of the 7th Workshop on Argument Mining*. ACL, 76–87. <https://www.aclweb.org/anthology/2020.argmining-1.9>
- [31] Christian Stab, Johannes Daxenberger, Chris Stahlhut, Tristan Miller, Benjamin Schiller, Christopher Tauchmann, Steffen Eger, and Iryna Gurevych. 2018. ArgumenText: Searching for Arguments in Heterogeneous Sources. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*. ACL, 21–25. <https://doi.org/10.18653/v1/N18-5005>
- [32] Paul Thomas, Daniel McDuff, Mary Czerwinski, and Nick Craswell. 2020. Expressions of Style in Information Seeking Conversation with an Agent. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*. ACM, 1171–1180. <https://doi.org/10.1145/3397271.3401127>
- [33] Henning Wachsmuth, Martin Potthast, Khalid Al-Khatib, Yamen Ajjour, Jana Puschmann, Jiani Qu, Jonas Dorsch, Viorel Morari, Janek Bevendorff, and Benno Stein. 2017. Building an Argument Search Engine for the Web. In *Proceedings of the 4th Workshop on Argument Mining*. ACL, 49–59. <https://doi.org/10.18653/v1/W17-5106>
- [34] Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 2018. Retrieval of the Best Counterargument without Prior Topic Knowledge. In *56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, Iryna Gurevych and Yusuke Miyao (Eds.). ACL, 241–251. <http://aclweb.org/anthology/P18-1023>
- [35] Klaus Weber, Niklas Rach, Wolfgang Minker, and Elisabeth André. 2020. How to Win Arguments. *Datenbank-Spektrum* 20, 2 (2020), 161–169. <https://doi.org/10.1007/s13222-020-00345-9>