

DOI:10.1145/2500891

**Improve online public discourse
by connecting opinions across blogs,
editorials, and social media.**

**BY FLORIS BEX, JOHN LAWRENCE,
MARK SNAITH, AND CHRIS REED**

Implementing the Argument Web

ARGUMENT AND DEBATE are cornerstones of civilized society and intellectual life. As online activity usurps many traditional forms of communication, we would hope to see these processes alive and well on the Web as well. But we do not. Too many mechanisms for online interaction hamper and discourage debate, facilitating poor-quality argument and fuzzy thinking. Needed are new tools, systems, and standards engineered into the heart of the Web to encourage debate, facilitate good argument, and promote a new online critical literacy. This is the Argument Web vision, involving a Web platform combining linked argument data with software tools that make online debate intuitive for its participants, including mediators, students, academics, broadcasters, and bloggers.


New opinions are constantly being presented on websites, blogs, news sites, and discussion forums, challenged and evaluated by a diverse worldwide user group. An important problem is the semantic

structure of argumentative viewpoints; for example, whether one person agrees with another or whether a particular statement conflicts with a claim is not captured. A further problem is there is no representation of how arguments across the Web relate to one another and contribute to the overall picture. Despite the numerous vocal communities on the Web, they remain relatively isolated because opinions are not connected.

Needed is an infrastructure allowing for interconnected arguments to be posted anywhere on the Web through a comprehensive underlying ontology of argument. This is the Argument Web,^{13,14} a URI-addressable structure of linked argument data making it possible to follow a line of argument (on a particular topic or by a particular person) across disparate forums, comments, editorials, and multimedia resources. A number of bespoke tools have been developed as part of the Argument Web implementation. Various annotation and analysis tools have been developed for academics and trained discourse analysts. Moreover, while argument analysis may be a specialized skill, most people conduct an argument, give reasons, conclude, and give grounds for disagreement every day; it is this intuitive skill the Argument Web aims to support explicitly. Familiar interfaces (such as

» key insights

- **The Web's focus on popularity (such as reflected through "like" buttons and number of followers) instead of rationality (such as through "agree" or "disagree" buttons and number of people subscribing to an opinion) discourages online discourse.**
- **The Web's numerous vocal communities (such as bloggers, Reddit, and Twitter) are rather isolated, with opinions expressed on one website not directly connected to opinions expressed on other websites.**
- **The Argument Web includes new tools, systems, and standards for linking argument data, allowing opinions to be connected across the Web through semantically meaningful links.**



blogging and instant messaging) have been adapted, allowing users to navigate opinions and express agreement or disagreement.

In 2007, some of the basic ideas behind the Argument Web were first expressed by Rahwan et al.¹⁴ This article and corresponding webpage (<http://www.argumentinterchange.org>) explore the first full prototype, discussing a mature version of the argument ontology, implementation of the underlying linked infrastructure, and tools allowing interaction with the Argument Web.

Connecting Opinions

One of the main functions of the Web

is enabling people to share, comment on, and argue on a variety of topics, from the newest video game to who should be the next president. Popular blogs can have large numbers of followers who comment on the blog and on each other; discussion forums (such as Reddit) allowing people to share and discuss topics and incorporate mechanisms for rating contributions and users; and Twitter and Facebook allowing users to quickly share and comment on their friends' opinions. While online interaction is facilitated and promoted, online critical discussion usually is not. Comments are rated almost solely according to popularity, not according to

whether they present a valid rational argument; a funny picture of President Barack Obama is more likely to end up on the front page of a website than a cogent argument for why one should vote for or against him. It may be possible to “like” a comment, but one cannot “agree” or “disagree” with a comment. Moreover, the only type of relation between two statements is usually “reply” (statement 1 replies to statement 2) and more specific argumentative relations (such as statement 1 supports statement 2 and statement 1 opposes statement 2) are not available.

The need for better online argument and debate is recognized by

websites supporting online critical thinking and structured debate. Some offer databases of structured argument for users to explore. Debatabase^a (formerly Debatepedia) offers numerous high-quality debates about a range of topics, including those for and against a particular topic. Archelogos^b concentrates on arguments from ancient philosophy (such as Aristotle and Plato), making explicit their argumentative structure

a <http://idebate.org/debatabase>
 b <http://archelogos.com>

and logical interconnections. Other websites are more interactive, allowing users to construct arguments in a structured way. Truthmapping^c is a more interactive discussion board that introduces slightly more structure to a debate, requiring new statements either support or oppose an existing statement. In Debategraph,^d debates are visualized as radial trees, with a central topic node surrounded by related statement nodes. Users

c <http://truthmapping.com>
 d <http://debategraph.org>

click nodes to expand elements of a debate and connect new nodes to the tree via relations (such as respond, support, and oppose). A graphical presentation like the one in Debategraph can give a quick overview of complex debates,^e helping users make sense¹¹ of a large amount of information; other examples of websites offering visual representations of arguments are Argunet^f and Cohere.

These Web technologies all have an explicit semantic structure that connects the statements and arguments in a debate, allowing for far better navigation and analysis of a debate; for example, we can render the structure as a graph or as an outline report, count the number of arguments for and against a claim and thus evaluate argument strength, analyze the graph and identify circular reasoning or gaps in chains of arguments, and discover inconsistencies among arguments and agreements among disputants.

The main limitation of these websites is there is little or no integration between them and they do not fully connect to the Web as a whole. Links to sources on the Web can be incorporated into arguments, and debates can be shared by providing links to separate discussion threads or debate graphs. However, these links contain no explicit semantics. Hence, each site provides a silo of structured and semantically rich argumentative content, but these silos are not connected to one another or to the rest of the Web, at least not through semantic links.

To illustrate the added value of connected arguments across the Web, consider the issue of Bashar al-Assad's Syria and the morality of potential Western intervention in the summer of 2012. Say you ask, "Should we invade Syria?" Ideally, you would find a number of arguments, along with links to their sources:

Video. A YouTube video of a press conference in which the U.K. Prime

e For examples of argument and debate analyses, see the "Can Computers Think" argument at <http://www.macrovu.com/CCTGeneralInfo.html> and http://debategraph.org/can-computers_think and the U.K. nuclear weapons debate at <http://www.florisbex.com/NuclearArgMap.pdf>

f <http://argunet.org>

Figure 1. Example of linked argument data.

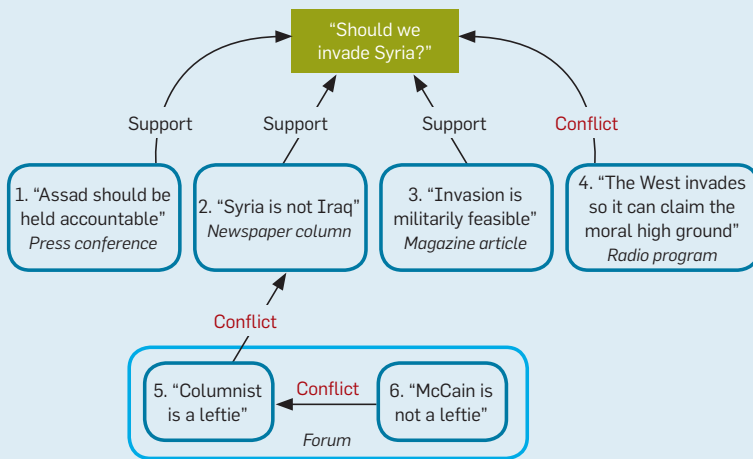
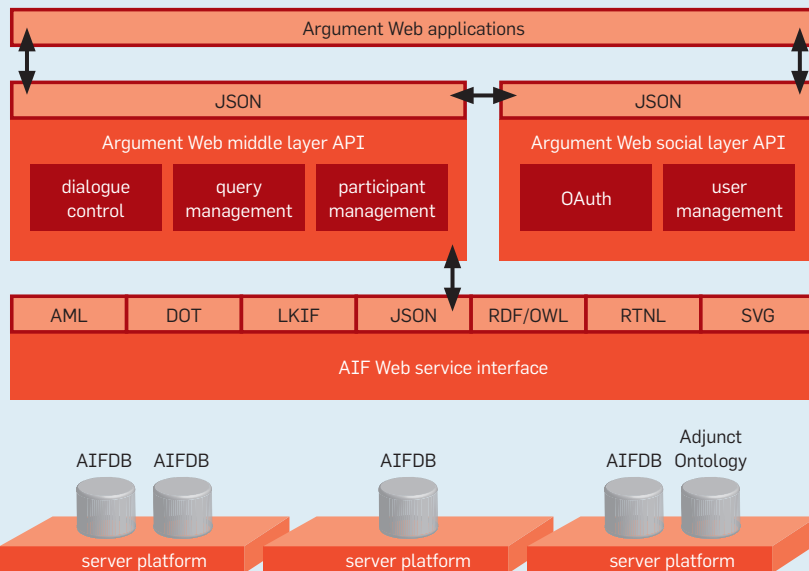


Figure 2. The Argument Web architecture.



Minister says Assad should be held accountable for war crimes;^g

Newspaper. A column in the online edition of *The Guardian* claiming Syria is different from Iraq;^h

Magazine. An article in *Foreign Policy* exploring the military feasibility of attacking Assad's Syria;ⁱ and

Radio. A point about the morality of intervening put forward in the BBC Moral Maze program;^j


(You focus on the newspaper column, which links to other arguments posted on the Web agreeing or disagreeing with the fact that Syria is different from Iraq.)

Forum. One counterargument posted on a forum^k says the writer of the column is a “leftie” with no idea what is the difference between Syria and Iraq or what real war means; and


On the same forum, another user disagrees with this argument, saying “Arizona Senator John McCain is for intervention and is no leftie.”

Figure 1 outlines the small Web of linked arguments, along with support and conflict links. Note the statements and arguments can be on different webpages, but a counterargument can also be given on the same page, as in the forum. In order to realize linked arguments across the Web we need to conform to the demands of linked data,⁴ in that each piece of data, or the nodes and links between them, is addressable through a unique URI that stands in well-defined relationships to other URI-addressable data. Only two projects—Cohere^{l,7} and the Argument Web—allow for linked argument data.

Cohere aims to link ideas on the Web. Individual online statements can be referred to through a URL, supporting a range of semantic relationships between components (such as “explains” or any other relation the user wants to define). While this breadth is helpful for Cohere, it makes it diffi-



It is possible to use the Argument Web to explore, say, mathematical aspects of arguments phrased in natural language from various sources on the Web.



cult to build tools with specific applications; for example, argumentation requires a fixed set of argumentative relationships that can then support computational processing (such as visualization, navigation, and evaluation) not easily supported if the relationships being captured are dynamic.

Infrastructure of the Argument Web

The Argument Web aims to create a Web infrastructure that allows for storage and automatic retrieval and analysis of linked argument data.⁴ It is based on a common ontology for argument called the Argument Interchange Format⁶ (AIF) that ties together natural linguistic models of argument (such as models that see argumentation as a language activity¹⁷) with abstract mathematical models of argument.^{2,9} It is possible to use the Argument Web to explore, say, mathematical aspects of arguments phrased in natural language from various sources on the Web.

AIF Ontology. At its core, the AIF ontology distinguishes between information (such as propositions and sentences, or the nodes in Figure 1) and general patterns of reasoning that, applied to specific information, provide the individual relations between information (the links in Figure 1). Links can be classified according to the scheme they fulfill.¹³ The AIF scheme taxonomy is based on argumentation schemes¹⁸ that are generally accepted for scholarly investigation of argument; for example, the counterargument in Figure 1 saying the columnist is a “leftie” is a typical *ad hominem* argument,¹⁷ an argument against a person. Here, it is not what the columnist says—Syria is different from Iraq—that is being countered but the credibility of the speaker; that is, people who are left-of-center politically should not be taken seriously when commenting on invading other countries.

In addition to argument structures, or static structures representing information and the support and attack links between them, as in Figure 1, the AIF ontology is also able to capture the argument processes, or the dynamic discussions in which people put forward and challenge claims and reasons. In recent work we showed the AIF ontology can be used as a

g <http://youtube.com/watch?v=OBPcv7qSIi4>

h <http://guardian.co.uk/commentisfree/2012/feb/10/syria-not-iraq-wrong-intervene>

i http://foreignpolicy.com/articles/2012/01/10/the_syrian_invasion

j <http://bbc.co.uk/programmes/b01kkp5q#synopsis>

k http://amazon.com/forum/politics?_encoding=UTF8&cdForum=Fx1S3QSZRUL93V8&cdThread=TxIAO1K6FZPOSI

l <http://cohere.open.ac.uk>

general framework for capturing dialogue protocols.³ Such protocols provide rules that determine which types of responses can be given to which types of statements or questions; for example, a challenge statement like “Why is invading Syria militarily feasible?” can force other parties to give reasons for their claim, as in, “We should invade Syria because it is militarily feasible.”

Note there is no restriction on the content of an argument; for example, argument contents may themselves be expressed in a common formal ontology, meaning they yield to additional computational processing. On the other hand, this increases the complexity of the ontology, meaning common understanding of the ontology is reduced, and potential users are largely restricted to scholars and experts. Hence, the AIF core ontology is kept as basic as possible, and the Argument Web does not demand explicit semantic characterization of argument content.

Semantic Web. The Argument Web essentially represents a large-scale deployment of Semantic Web^{1,13} technology, taking a pragmatic approach to issues of infrastructure by specifying the ontology not only in the formal languages of the Semantic Web (such as the Web Ontology Language, or OWL-DL) but as an instance of a relational database schema adumbrated by a set of Web services acting as an RDBMS for interacting with the data as if it were an RDF triple store. This approach builds on highly scalable, mature, robust, commercially accepted database systems while still conforming to the main principles and demands of linked data.⁴

Argument Web resources are distributed across multiple databases, instances of what are called AIFdb (a database solution for the Argument Web¹⁰) in reference to the underlying AIF ontology (see Figure 2). Each AIFDB instance provides interfaces through a variety of formats for programmatic access, Semantic Web processing, visualization, and compatibility with existing argumentation tools.¹¹ Core argumentation functionality is then provided by the Argument Web Middle Layer in three areas:

Query management. Allowing developers to query, say, what argument(s)

Even rhetorically and linguistically sophisticated maneuvers (such as ad hominem argumentation) can be captured in the Argument Web.

supports a given claim or whether a given position conflicts with any other(s);

Participant management. Allowing developers to query, say, which users have put forward a given argument; and

Dialogue control. Allowing developers to execute dialogue protocols and, in that context, query whose turn it is in a given dialogue and what legal moves are available to the participants. The Argument Web social layer provides developers a straightforward library through which to construct applications that interact with the Argument Web.

Because the Argument Web, like the Web itself, is designed for public use, the data being produced is noisy, thus presenting a challenge when trying to build computational systems able to reason over Argument Web resources. One approach to dealing with noisy data is engineering oriented: Design the reasoning systems in the full expectation the data over which they work is noisy, and therefore results are either only as good as the data from which they are drawn or else drawn from subsets preprocessed to reduce noise.

The Argument Web’s core concept—argumentation—also provides another approach to dealing with noisy data. Take claim identity. In any large distributed knowledge base updated by many different individuals, a major challenge is how to identify when two items in the knowledge base are the same. However, the claim that, say, “the domineering Western aggressors invading Syria” is the same as “Western liberating forces intervening in Syria” is precisely the sort of claim that is liable to be the source of a dispute. Thus, what seems to be an engineering challenge—identifying duplicates—becomes a subject of discussion on the Argument Web.

Argumentation Tools and Interfaces

A number of tools have been developed as part of the Argument Web implementation. Some (such as for argument visualization and analysis) are geared mostly toward educational uses. However, in order to open up the Argument Web to a broader audience, familiar interfaces (such as blogging and instant messaging) can

be adapted to support argumentation.⁸ Moreover, the solid theoretical and infrastructural basis of the Argument Web allows us to experiment with new technologies for discourse and argument analysis, or Argument Web app library,^m discussed in the following sections.

Argument visualization and analysis.

As noted, Web technologies for argumentation focus in part on visualizing arguments as graphs or diagrams easily navigated by users. The idea of online argument visualization stems from offline argument-diagramming tools used in, say, courses on critical thinking and for teaching complex legal argumentation.¹² The Argument Web implementation recognizes existing approaches to argument visualization by defining and implementing import-export functions^{2,8} for existing argument visualization tools.ⁿ

Using a visual interface to AIFDB (see Figure 3a) lets users view and navigate the Argument Web; like Debategraph, nodes can be clicked to expand new parts of a debate. New arguments can be constructed using the Online Visualization of Argument, or OVA, allowing direct analysis of Web content; text from a website can be selected and input directly to an argument graph that is then saved to the Argument Web.

For the University of Dundee’s partnership with the BBC, we wanted to analyze a discussion in real time. The BBC’s Radio 4 program “The Moral Maze” broadcasts 45 minutes of debate on a moral issue each week. The idea is if an analysis is done in real time and saved directly to the Argument Web, listeners would be able to interact with the discussion as it happens. However, a trained analyst can take weeks to analyze even one hour of debate, so analysis must be performed in parallel. We therefore designed and built a collaborative workspace—the AnalysisWall, a touchscreen measuring 11 feet by 7 feet running bespoke analysis software¹⁴ (see Figure 3b). As the program was broadcast live, stenographers provided a text feed, segmentation analysts broke the text into its component parts, and eight analysts collaborated

Figure 3. Argument visualization and analysis.

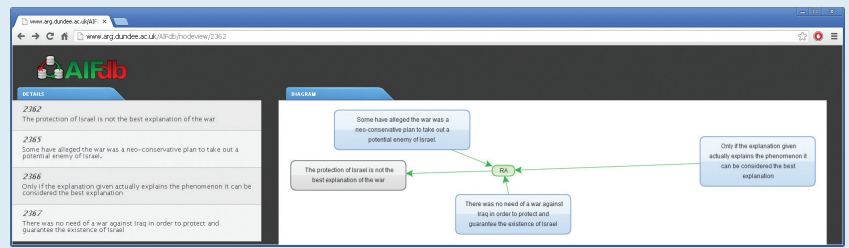


Figure 4. The Arvina 2 debate interface.

The screenshot shows the Arvina 2 web interface. At the top, there is a navigation bar with 'New Topics', 'Popular Topics', 'Create a New Topic', and 'Logout'. Below this is a chat feed with several messages from users like 'Clifford' and 'Matthew'. On the right side, there is a 'Topic' sidebar with a Union Jack icon and text about 'The British Empire'. At the bottom, there is a 'Live Discussion Map' and a 'Participants' section. The interface is labeled (c).

m <http://www.argumentinterchange.org/library>
 n <http://argumentinterchange.org/developers>

to tease apart the structure of what is being said, directly inserting it into the Argument Web.

Direct discussion and mixed-initiative argumentation. Direct discussion between two or more people on the Web takes place not just via email and instant messaging but also on forums and message boards. However, these technologies offer only the most basic of structural tools, and the inferential structure of an argument in a discussion is easily lost. The Web-based discussion software Arvina¹⁰ (see Figure 4) allows participants to debate a range of topics in real time in a way that is structured but at the same time is unobtrusive. Users can ask questions (such as “Do you agree?” and “Why do you think this is the case?”) of other participants in the discussion, as well as express their own agreement or disagreement with a particular point and provide supporting reasons for their views. Moreover, Arvina can use dialogue protocols to structure the discussion between participants.

Arvina can take on a multi-agent system populated by agents representing authors whose opinions are available on the Argument Web. It is thus possible to question the participants of, say, past Moral Maze radio programs about their opinions, and agents representing participants can be added to a discussion and questioned about what they think about the topic discussed during the program. An agent then answers by giving the opinions originally expressed during the broadcast by a particular

participant. In this way people are able to question any opinion expressed on the Argument Web, whether originally added through OVA, ArguBlogging, or other tools.

Argument blogging. A final example of how argumentation technologies based on the Argument Web facilitates online debate concerns blogging, a highly popular form of online communication. If one wants to reply to an opinion presented somewhere on the Web in a blog, the usual way is to provide a simple hyperlink to the article or page in which the opinion is expressed. However, the resulting structure of supporting and competing opinions is easily lost due to lack of semantic information in the links.

We built a simple tool called ArguBlogging, for Argument Blogging,¹⁰ to improve rational debate through this popular form of expressing opinions online. It ensures that, if desired, opinions on blogs and other webpages can be linked through the underlying infrastructure of the Argument Web, allowing users to agree or disagree with opinions anywhere on the Web, simultaneously posting it to their blog (connecting to two popular blogging platforms, Blogger and Tumblr) and adding it to the Argument Web; Figure 5 is the ArguBlogging window (or widget) rendered on a webpage, providing options for responding.

A typical blog post through ArguBlogging contains an “argue” button that, clicked, brings the widget onscreen. It allows users to respond to the blog post as a whole through ArguBlogging and is similar to the

“like” button on Facebook and “share” button on Twitter.” This way, the ArguBlogging tool uses familiar interaction styles to enable critical argumentation on the Web.


Evaluating the Argument Web

The kind of continuous evaluation needed by the Argument Web involves usability in two distinct domains—public and academic—and expressivity and computational flexibility through metrics drawing on formal, computational methods, as well as on more pragmatic engineering principles.

In terms of raw usage through compatible tools (such as Rationale and Araucaria), the Argument Web today includes tens of thousands of users worldwide. The native applications (such as OVA) released in the past few years naturally involve fewer, but the core AIFdb includes more than 11,000 argument components, the second largest semantically rich argumentation dataset (after Debategraph), including in seven languages; for example, the Archelagos repository is fully imported into the Argument Web. More significant, the infrastructure and analysis tools have been used in at least seven research projects involving 10 labs in France, Poland, the U.K., and the U.S., each with different foci, including generative and analytical, multilingual and monolingual, and dialogical and monological.

In terms of expressivity, the Argument Web balances well-defined formal properties and pragmatic solutions to engineering problems; for example, the underlying ontology is demonstrably more expressive than one of the foremost formal accounts of defeasible argumentation, ASPIC+,² but a well-defined subset can be used to induce abstract frameworks on which evaluation can be performed.¹⁶ Ongoing work in discourse analysis⁵ demonstrates that even rhetorically and linguistically sophisticated maneuvers (such as ad hominem argumentation) can be captured in the Argument Web. At the same time, formal description-logic analysis of the AIF ontology¹³ shows powerful abstractions can support advanced search and evaluation of arguments involving, say, specific schemes of presumptive reasoning.

Figure 5. The ArguBlogging widget.



The Argument Web's ultimate objective is to improve the quality of online argument and debate. Keeping in mind that evaluation of natural argument is philosophically thorny, it is possible to propose objective metrics through which to assess natural arguments, including, say, consensus about which arguments are best (one might disagree with an argument while still appreciating its merits); exhaustiveness, speed, volume of content, signal to noise, and structural complexity; argument richness (in terms of the range of argument types used); and dialogue richness (range of dialogical moves used).

Conclusion

The Argument Web represents the first technology linking debate, disagreement, and argument structures from a variety of tools applied in different domains. The approach has strong potential both academically and practically. Along with developer and user interest and increasing Argument Web resources, the academic community gains access to a valuable resource that, particularly for computer scientists, can function as a testbed for new theories of argument-acceptability applications and as a rich dataset with which to deploy new applications; for linguists and philosophers, it offers a unique corpus of discourse activity, replete with detailed annotation and commentary.

By solving theoretical problems involving how argument structures can be navigated and extended through dialogical processes, the Argument Web opens up a new class of application in which intuitive, dialogically based interfaces (such as Arvina and ArguBlogging) can be used to explore and improve large-scale debates. There is evidence¹⁹ that debate is a good way to navigate and support engagement with complex issues involving disagreement (such as abortion, climate change, and military intervention) that are not just important but that define our time. By supporting and facilitating engagement in debates (otherwise daunting, leading even to disengagement and disempowerment) the Argument Web promises to play not only a technological but an important societal role as well. 



Debate is a good way to navigate and support engagement with complex issues involving disagreement (such as abortion, climate change, and military intervention) that are not just important but that define our time.



References

1. Berners-Lee, T., Hender, J., and Lassila, O. The Semantic Web. *Scientific American* 284, 5 (2001), 28–37.
2. Bex, F., Modgil, S., Prakken, H., and Reed, C. On logical reifications of the Argument Interchange Format. *Journal of Logic and Computation*. Aug. 2012; doi: 10.1093/logcom/exs033
3. Bex, F. and Reed, C. Dialogue templates for automatic argument processing. *Frontiers in Artificial Intelligence and Applications* 245 (2012), 366–377.
4. Bizer, C., Heath, T., and Berners-Lee, T. Linked data: The story so far. *International Journal on Semantic Web and Information Systems* 5, 3 (2009), 1–22.
5. Budzynska, K. and Reed, C. The structure of ad hominem dialogues. *Frontiers in Artificial Intelligence and Applications* 245 (2012), 410–421.
6. Chesñevar, C., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., Vreeswijk, G., and Willmott, S. Towards an Argument Interchange Format. *Knowledge Engineering Review* 21, 4 (2009), 293–316.
7. De Liddo, A., Sándor, Á., and Buckingham Shum, S. Contested collective intelligence: Rationale, technologies, and a human-machine annotation study. *Computer Supported Cooperative Work* 21, 4–5 (2012), 417–448.
8. de Moor, A.D. and Aakhus, M. Argumentation support: From technologies to tools. *Commun. ACM* 49, 3 (Mar. 2006), 93–98.
9. Dung, P.M. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence* 77, 2 (1995), 321–357.
10. Lawrence, J., Snaith, M., Bex, F.J., and Reed, C. Demonstration papers on the AIFdb, ArguBlogging, Arvina, and TOAST. *Frontiers in Artificial Intelligence and Applications* 245 (2012), 511–516.
11. Okada, A., Buckingham Shum, S.J., and Sherborne, T., Eds. *Knowledge Cartography: Software Tools and Mapping Techniques*. *Advanced Information and Knowledge Processing*. Springer, London, 2008.
12. Pinkwart, N., Ashley, K., Lynch, C., and Aleven, V. Evaluating an intelligent tutoring system for making legal arguments with hypotheticals. *International Journal of Artificial Intelligence in Education* 19, 4 (2009), 401–424.
13. Rahwan, I. Mass argumentation and the Semantic Web. *Web Semantics* 6, 1 (2008), 29–37.
14. Rahwan, I., Zablith, F., and Reed, C. Laying the foundations for a World Wide Argument Web. *Artificial Intelligence* 171, 10–15 (2007), 897–921.
15. Reed, C., Bex, F., Lawrence, J., and Snaith, M. DEMO: The Argument Analysis Wall. In *Proceedings of Digital Futures: The Third Annual Digital Economy All Hands Conference* (Aberdeen, U.K., 2012).
16. Snaith, M. and Reed, C. TOAST: Online ASPIC+ implementation. *Frontiers in Artificial Intelligence and Applications* 245 (2012), 509–510.
17. van Eemeren, F.H. and Grootendorst, R. *A Systematic Theory of Argumentation: The Pragma-Dialectical Approach*. Cambridge University Press, Cambridge, U.K., 2004.
18. Walton, D.N., Reed, C.A., and Macagno, F. *Argumentation Schemes*. Cambridge University Press, Cambridge, U.K., 1995.
19. Yuan, T., Moore, D., Reed, C., Ravenscroft, A., and Maudet, N. Informal logic dialogue games in human-computer dialogue. *Knowledge Engineering Review* 26, 3 (2011), 159–174.

Floris Bex (f.j.bex@rug.nl) is a researcher and instructor at the University of Groningen, Groningen, The Netherlands.

John Lawrence (johnlawrence@computing.dundee.ac.uk) is a Ph.D. student at the School of Computing of the University of Dundee, Dundee, U.K.

Mark Snaith (marksnaith@computing.dundee.ac.uk) is a postdoctoral research assistant at the School of Computing of the University of Dundee, Dundee, U.K.

Chris Reed (c.a.reed@dundee.ac.uk) is a professor at the School of Computing, University of Dundee, Dundee, UK.

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.