Debating Technology for Dialogical Argument: Sensemaking, Engagement, and Analytics

JOHN LAWRENCE, MARK SNAITH, and BARBARA KONAT, University of Dundee
KATARZYNA BUDZYNSKA, University of Dundee and Polish Academy of Sciences
CHRIS REED, University of Dundee

Debating technologies, a newly emerging strand of research into computational technologies to support human debating, offer a powerful way of providing naturalistic, dialogue-based interaction with complex information spaces. The full potential of debating technologies for dialogical argument can, however, only be realized once key technical and engineering challenges are overcome, namely data structure, data availability, and interoperability between components. Our aim in this article is to show that the Argument Web, a vision for integrated, reusable, semantically rich resources connecting views, opinions, arguments, and debates online, offers a solution to these challenges. Through the use of a running example taken from the domain of citizen dialogue, we demonstrate for the first time that different Argument Web components focusing on sensemaking, engagement, and analytics can work in concert as a suite of debating technologies for rich, complex, dialogical argument.

1. INTRODUCTION AND MOTIVATION

Debating technology is a newly emerging research area that aims to support human debating through the use of technology. The area has been fueled by an eponymous Dagstuhl seminar [Gurevych et al. 2016],[1] where Slonim [2016] defined debating

---

[1]The meeting documented in Gurevych et al. [2016], available at http://drops.dagstuhl.de/opus/volltexte/2016/5803, took place at Schloss Dagstuhl, 13–18, December 2015. The abstract from the report reads as follows: "This report documents the program and the outcomes of Dagstuhl Seminar 15512 "Debating Technologies". The seminar brought together leading researchers from computational linguistics, information retrieval, semantic web, and database communities to discuss the possibilities, implications, and necessary actions for the establishment of a new interdisciplinary research community around debating technologies".

---

technologies as computational technologies developed directly to enhance, support, and engage with human debating.[2]

The connections between debate and argumentative dialogue are long established [Walton and Krabbe 1995] and are found in many day-to-day situations. One such domain is *citizen dialogue* (CD), which involves consultations undertaken by a public body or bodies to collect opinions and arguments from citizens. Public meetings and hearings use dialogue to generate a significant amount of argumentative data that the public body will then use when making a decision.

The debating technology community has built argumentative dialogue into its foundations [Rinott 2016; Plüss 2016; Stab and Habernal 2016a] and thus offers significant potential to assist the decision-making process in domains such as CD. However, there are several key technical and engineering challenges that need to be overcome for the full potential of debating technologies for dialogical argument to be realized.

The first of these challenges is *data structure*. Online debates are generally based on a post-then-reply metaphor that lacks argumentative structure. This means that as debates grow larger, they become increasingly difficult to keep track of, seriously impacting both their quality and the ability to perform postprocessing of the data they generate. The second related challenge is *data availability*; for many tasks, particularly in computational linguistics, annotated data is crucial, with supervised learning requiring large sets of such data. The creation of such resources, however, is extraordinarily expensive, with common datasets such as the Internet Argument Corpus [Walker et al. 2012] highlighting the value in such resources. Yet annotation and deployment are typically ad hoc with, as yet, relatively little uptake and reuse [Stab and Habernal 2016b]. The third challenge is *interoperability*. Components for analyzing natural argument and debate, processing structured and abstract argument data, and evaluating large-scale networks have existed for some time (e.g., Gordon and Karacapilidis [1997], Ravenscroft et al. [2006], Klein and Iandoli [2008], and De Liddo and Buckingham-Shum [2014]). Although these can be considered instances of debating technology, each system is nevertheless largely stand-alone. The work of the IMPACT project (e.g., see Gordon [2011]) clearly laid out the motivation and advantages of inclusive, open, and thorough interoperability to allow such tools to work harmoniously. Yet the engineering of components that build upon one another and interact with each other, although common practice in many areas of natural language processing, remains the exception in debating technology.

These challenges can be addressed by functional and efficient middleware. By defining sets of open interfaces and using common standards for representing argument and debate, interoperability can be streamlined and data can be structured and made readily available. The Argument Web [Rahwan et al. 2007; Bex et al. 2013a] is one such vision for integrated, reusable, semantically rich resources connecting views, opinions, arguments, and debates online. The growth of the Argument Web has seen the development of various services and application programming interfaces (APIs) into an infrastructure that facilitates storage, retrieval, and interaction with data generated from argument and debate [Lawrence et al. 2012b; Lawrence and Reed 2016].

Our aim in this article is to show how this Argument Web infrastructure can act as the necessary middleware to support debating technology for dialogical argument. With the help of a running example, we present a suite of tools that enrich and support debate by allowing users to do the following:

---

[2]See also the IBM Watson project at http://researcher.ibm.com/researcher/view_group.php?id=5443.

—Add structure to existing, unstructured sources (sensemaking[3])
—Use dialogue games to both navigate and provide structured contributions to new or existing debates (engagement)
—Extract valuable information from the structured data (analytics).

The article proceeds as follows. In Section 2, we introduce our running example and the domain from which is it taken (CD). In Sections 3 through 5, we present technologies for sensemaking, engagement, and debate analytics, respectively, using the running example for illustration throughout. In Section 6, we conclude this work and lay out directions for potential future work.

2. CITIZEN DIALOGUE

By CD, we understand here any form of consultation undertaken by a public body or bodies to collect opinions and arguments from citizens, concerning new legal regulations, solutions, or plans of urban development. As part of ongoing work, we annotate with argumentative–dialogical relations several transcripts of public meetings and hearings from departments of transportation (DOTs) across the United States of America.[4]

CD is a polylogue: in the large room used as a meeting space (usually a school hall, a church, or a local community center), two groups of speakers are interacting: representatives of the governmental agency (usually three to five people) and representatives of local community (numbers vary between meetings). Moreover, agreement and disagreement can occur between multiple speakers (however, in some cases, this is quite unlikely, e.g. representatives of the governmental agency would rarely disagree with each other publicly).

Excerpt (1) presents an interaction from a question and answer session taken during a public hearing in North Carolina.[5] In the example, the citizen (Ken Kelly, a pastor of the local Baptist church), is presenting his concerns about the location of a road that is planned to be built. The plan, as introduced by the local DOT, required the road to be situated across the property belonging to the church and community center. This example demonstrates an argumentative situation—the pastor is making an argument as to why the current plan is not acceptable.

(1) **Ken Kelly**: Thank you, my name is Ken Kelly. I'm the Pastor of that little Baptist Church you were pointing to there just while ago. Two years ago, we were here and, my recommendation and request would be to you, would be to just leave us alone and move on down there if you remember that two year ago meeting. [. . .] But my concern is still, my concern is still with the church, that the worst thing that you can do with us, is to take our property and our parking spaces and not take our buildings. The exchange there is to go between the buildings, or go right down the middle of the buildings, leaving us with some on either side of the road as I see it on the map.
DOT Moderator: No, actually I think the right of way goes right across your sanctuary.
Ken Kelly: That's what—well the sanctuary is the building closest to Huffine Mill Road and the educational space is the other part.
DOT Moderator: Okay.
Ken Kelly: So we would save part of our sanctuary, get my office, and the educational space and leave us with the fellowship hall down there.

---

[3]*Sensemaking* is a broad term: we use it here in the sense developed in Kirschner et al. [2003].
[4]The annotated CD corpus is publicly available as part of AIFdb at http://arg.tech/cd.
[5]The original transcript is obtained from the Web site of the North Carolina Department of Transportation (NC DOT) as an official public hearing transcript from the meeting concerning plans for an urban loop in Greensboro, which took place on May 11, 1995 (U-2525): see http://www.ncdot.gov/projects/greensborourbanloop.

Such public consultations used as a means of collecting feedback from citizens can turn into a complex and time-consuming process. The consultation scheme, from which Excerpt (1) was taken, lasted 15 years between 1995 and 2010, with construction work finally starting in 2014. The excerpt presented here is just a small part of one transcript of one meeting—typically, an entire consultation consists of several meetings, each generating transcripts of up to 15,000 words. Each of these transcripts requires scrutiny and manual summarization, with all arguments (pro and con) from citizens addressed. Unstructured information with dynamic dialogical interactions needs to be analyzed and understood to support decision making for socially important issues. In the following sections of this article, we demonstrate how the application of various tools for debating technologies, sensemaking, engagement, and analytics can support and improve the process of conducting and understanding CD.

3. SENSEMAKING

3.1. Theoretical Foundations: A Common Standard for Argument Interchange

The argument interchange format (AIF) [Chesñevar et al. 2006] was developed with the aim of creating a means of expressing argument that would provide a flexible yet semantically rich way of representing argumentation structures. The AIF was put together to try to harmonize the strong formal tradition initiated to a large degree by Dung [1995], the natural language research described at Computational Models of Natural Argument (CMNA) workshops since 2001,[6] and the multiagent argumentation work that has emerged from the philosophy of Walton and Krabbe [1995], among others (e.g., see Reed and Walton [2005], Parsons and Jennings [1996], McBurney and Parsons [2002], and Tang et al. [2009]). As originally specified, the AIF accounted for only monological argument; however, it has been subsequently extended by Reed et al. [2008] into AIF+, which is what we use in the present work. Where we refer to extensions to the core AIF, these apply equally to AIF+.

Central to the AIF+ core ontology are two types of node: information nodes (I-nodes) and scheme nodes (S-nodes). I-nodes represent propositional information contained in an argument. A subset of I-nodes represent propositional information specifically about discourse events: these are L-nodes or locutions. S-nodes capture the application of schemes (patterns of reasoning (RA-nodes), of conflict (CA-nodes), of preference (PA-nodes), of illocution (YA-nodes), or of dialogical transition (TA-nodes)). Illocutionary schemes are patterns of communicative intentions (e.g., of asserting, challenging, questioning, arguing, disagreeing) that speakers use to introduce propositional contents.[7] Dialogical transitions are schemes of interaction or protocol of a given dialogue game that determine possible relations between locutions. An individual S-node fulfills—or instantiates a specific instance of—a scheme. This core ontology of AIF+ is summarized in Figure 1, and the AIF+ notation is described in Table I.

Any single conception of argument, however, is liable to either miss aspects important for a particular philosophical or theoretical investigation (see van Eemeren et al. [2014] for a review that demonstrates the breadth of such theoretical investigations) or omit dimensions that are crucial for the development of a specific application. As such, the AIF (and by extension AIF+) offers a mechanism by which *adjunct ontologies* (AOs) can be defined to supplement the core ontology for the purposes of capturing application-specific information and concepts [Bex et al. 2013a]. These AOs are ontologies in their own right that serve to impose structure on aspects of argument not captured by the

---

[6]http://www.cmna.info.
[7]Illocutionary schemes are based on illocutionary forces defined in Searle [1969] and Searle and Vanderveken [1985]. They connect locutions with their contents.
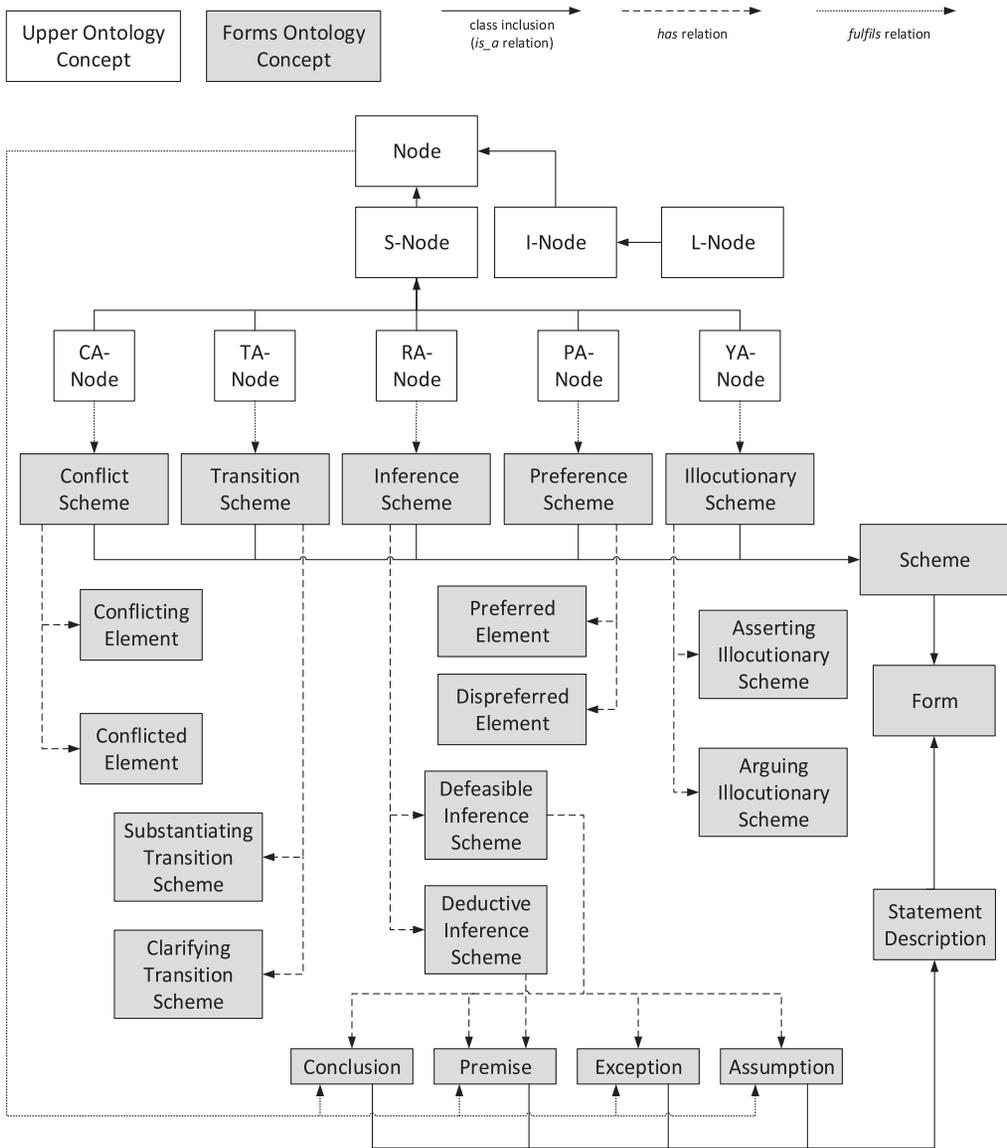
Fig. 1.   AIF+ core ontology.

Table I. AIF+ Notation

Node Type	Description
I	Propositional information contained in an argument, such as a conclusion, premise, data
L	Subset of I-nodes referring to propositional reports specifically about discourse events
RA	Application of a scheme of reasoning or inference
CA	Application of a scheme of conflict
PA	Application of a scheme of preference
YA	Application of a scheme of illocution describing communicative intentions that speakers use to introduce propositional contents
TA	Application of a scheme of interaction or protocol describing relations between locutions

core AIF. They operate under a single constraint, namely that data held according to an AO can only adumbrate and extend data held according to the AIF. In other words, when an application makes a call upon a Web service defined according to an AO, the middleware responsible for handling that AO in turn queries the core AIFdb and then adds data (typically from additional data storage) to complete the picture according to the AO. In this way, the data that is common between different applications can be shared, whereas that which is unique is not contaminated by project- or theory-specific extensions. In principle, this permits multiple AOs to be layered on top of one another, allowing for evolution of the data standards in response to the needs of the research community. Several AOs have been developed. One of the first is to handle social aspects of argument and provides an additional layer to the Argument Web (described in Section 4.1.1).

Although the AIF is used primarily as a means of representing natural argument, it has demonstrable connections to the computational models of argument that have been shaped by the influential work of Dung [1995]. Informally, Dung abstracted argument into two concepts: *arguments* and a notion of *attack* between them. Arguments have no internal structure and the nature of attack is not defined. Given a set of arguments and an attack relation between them, an *argumentation framework* is constructed. Argumentation frameworks are evaluated under several different *acceptability semantics*. Informally, an argument is accepted if all of its counter-arguments are not. Different semantics offer different determinations of acceptability ranging from highly skeptical to highly credulous.[8]

Dung's theory has been adapted, extended, and built upon. In particular, several techniques have been developed to allow structure to be introduced to the otherwise abstract atoms corresponding to arguments. One such technique and theory is ASPIC+ [Prakken 2010], which combines the work of Pollock [1987] with that of Vreeswijk [1997] to provide an account of structured argumentation from which a Dung-style framework can be obtained and evaluated for acceptability. Bex et al. [2013b] subsequently demonstrated that data represented using the AIF can be expressed in terms of ASPIC+ and vice versa, thus allowing the acceptability of natural arguments (expressed in AIF) to be determined. This connection between the AIF and ASPIC+ has been implemented in TOAST[9] [Snaith and Reed 2012], allowing for evaluation of natural arguments, similar to that found in systems such as ASPARTIX [Egly et al. 2008], Tweety [Thimm 2014], and DIAMOND [Ellmauthaler and Strass 2014].

We make use of the acceptability of natural arguments in debate analytics (see Section 5). Through comparing the acceptability of the arguments of two or more participants in a dialogue, we can compare the relative influence, or sway, of those participants.

3.2. Argument Representation: Creating AIF+ Structures

Populating the Argument Web with structured data, particularly when it was first deployed, was a challenge, with two main alternatives: manual annotation and automated text mining. To bootstrap the deployment of the Argument Web, both of these techniques have been employed.

---

[8]As well as the original semantics specified by Dung [1995], the reader is also referred to the work of Dung et al. [2007] and Caminada et al. [2012].
[9]The Online Argument Structures Tool: http://toast.arg-tech.org.

3.2.1. Manual Annotation. Many software tools exist for the manual analysis of argument; we are aware of only five that are compatible with AIF and AIF+ (although various other tools and platforms are working to provide compatibility), which are presented next.

First, Araucaria [Reed and Rowe 2004] allows a user to load a text document and map the argumentative structure contained therein. Support is provided for three types of diagram: tree based, Wigmore [1931], and Toulmin [1958].

Second, Carneades [Gordon and Walton 2006] is primarily designed for evaluating legal evidence and is both a diagramming style and underlying computational model for evaluation of arguments against various standards of proof.

Third, Rationale [van Gelder 2007] is a tool designed for building critical thinking skills. Support is provided for mapping dialogues and brainstorming.

Fourth, OVA+ (Online Visualization of Argument)[10] [Janier et al. 2014] is an online tool for analyzing and mapping arguments. It is similar to other tools such as Araucaria but permits direct analysis of Web pages and other online resources. Furthermore, it has full support for analyzing dialogical argumentation in AIF+. In addition, it facilitates real-time collaborative work that allows multiple analysts in different locations to work together on an analysis.

Fifth, the Argument AnalysisWall is a 3.2m long, 2.4m high touchscreen that allows for large-scale in-person collaborative analysis of complex debates in close to real time [Bex et al. 2013a]. A live transcript of a debate is first segmented into discrete components (or atoms) of argument which are then assembled into an AIF+ structure by a team of analysts.

OVA and the Argument AnalysisWall both natively support the AIF. Araucaria, Carneades, and Rationale use their own bespoke formats; however, all three can be translated into AIF(+) [Bex et al. 2012].

3.2.2. Argument Mining. The alternative to manual analysis of pre-existing argumentative resources outside the Argument Web is to design algorithms for automated processing (e.g., see Moens [2013], Peldszus and Stede [2013], and Lippi and Torroni [2015] for an overview of the field of argument mining). Until very recently, such automation was well beyond what was possible, but from 2014 to 2016, the availability of datasets, coupled with the increasing maturity of some text mining techniques and clear commercial demand, has led to an increasing range of preliminary research in the area. The ACL Argument(ation) Mining workshops[11] have published around 60 papers in the area, with a similar number at recent editions of COLING, ACL, and EMNLP (e.g., Levy et al. [2014], Peldszus and Stede [2015], and Rinott et al. [2015]).

Although it is early in the field of argument mining, the goal is to produce structures from natural discourse that could easily be modeled by an argument representation format such as AIF+. There are also early results demonstrating that even dialogical material can be automatically processed, and this work results directly in AIF+ output (Budzynska et al. [2014] describes this output using the corresponding theoretical structures available in inference anchoring theory, the implemented form of which is AIF+).

3.3. Manual Analysis of CD

Figure 2 shows an analysis of the CD transcript from Example (1) (see Section 2), created using OVA+ (see Section 3.2.1). The text is now represented with the use of
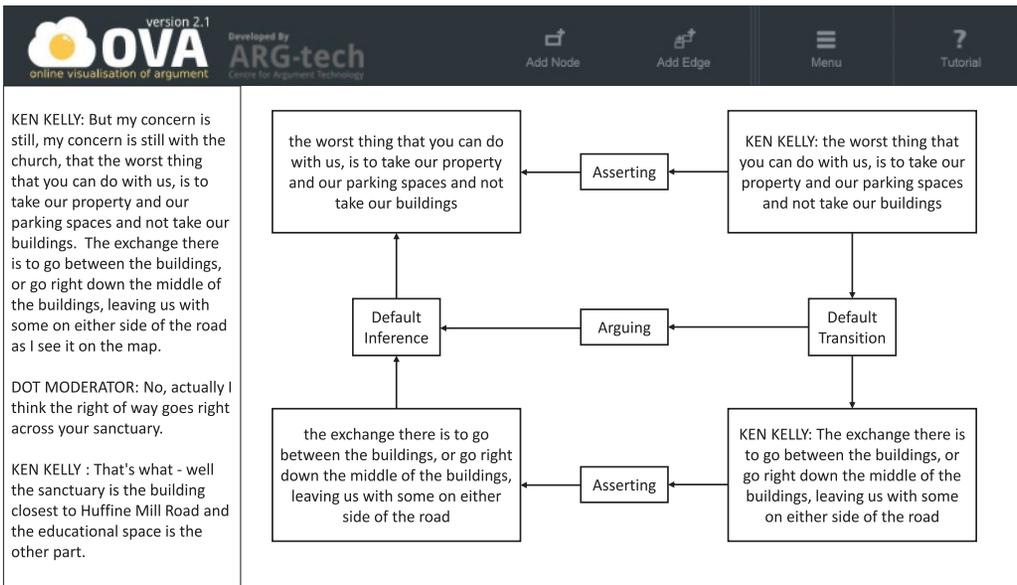

Fig. 2.   OVA+ analysis of CD map #9097—fragment.

AIF+ core ontology (described in Section 3.1 and Figure 1). Argument analysis tools such as OVA+ constitute the basic input for AIFdb. The full map, CD #9097, is available at the AIFdb Web site,[12] from where the editable version can be accessed in OVA+.

The result of manual OVA+ analysis is a computational model of an argument in the dialogical interaction. Thus, the analysis expresses the integration of argumentative (on the left side of the diagram) and dialogical (on the right side of the diagram) structures in the debating technology. On the left side of the diagram, there are I-nodes with propositional content introduced by speakers. Relations between I-nodes are represented by S-nodes, in this case a relation of *default inference*. On the right side of the diagram, there are L-nodes that contain locutions (the actual utterances from speakers). L-nodes also display information about the speakers (in this case, DOT Moderator and Citizen Kelly). L-nodes are connected by transition application nodes (TA-nodes), in this case *default transition*. The argumentative and dialogical parts of the diagram are connected by applications of illocutionary force (expressing a user's intention for a given locution, e.g., asserting, arguing, agreeing, or disagreeing), represented in AIF+ core ontology with illocutionary application nodes (YA-nodes).

Using OVA+, an analyst can add metadata, structuring argumentative–dialogical relations in the original text. The resulting argument map, added to AIFdb, increases annotated data availability in reusable format in an open database, accessible by the AIF Web service interface. In addition, the social layer (described in more detail in Section 4.1.1) is used to store information about speakers, such as to which group they belong (DOT representatives or citizens), the region they come from, their age, gender, and other information that can provide useful insights for analysis.

4. ENGAGEMENT

Engagement in complex debates requires methods to navigate the space in a natural-istic way while also allowing new structured and semantically rich contributions to
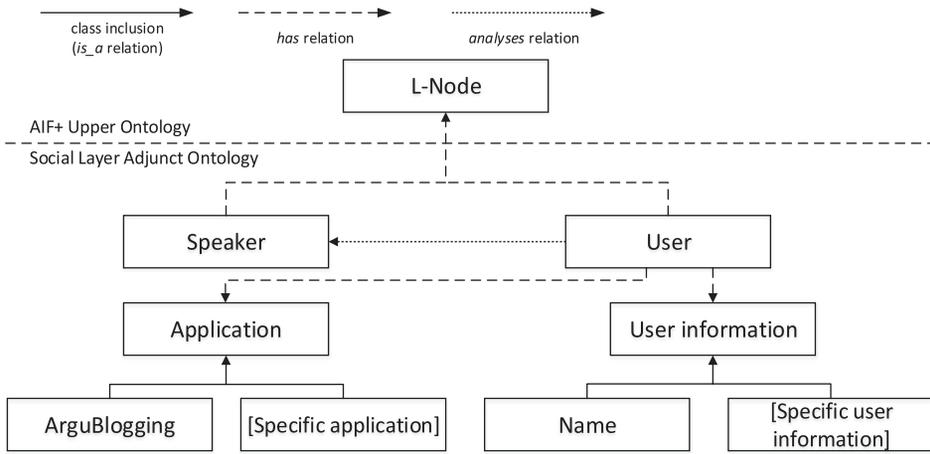
---

[12]At http://arg.tech/9097.

Fig. 3.   AIFdb social layer AO.

be gathered. In this section, we demonstrate how debating technologies for dialogical argument offer a solution to both of these challenges.

4.1. Theoretical Foundations: Social and Dialogical Aspects

4.1.1. The Argument Web Social Layer. Acting as a platform for social interaction, the Argument Web social layer ("the social layer") sits between applications and AIFdb [Snaith et al. 2016]. The social layer has two key functions within the Argument Web infrastructure: (1) managing the roles of Argument Web participants and (2) providing connections to the social World Wide Web through blogging platforms (e.g., Blogger[13] and Tumblr[14]) and social networks (e.g., Facebook[15] and Twitter[16]). The social layer processes *Social AIF* (S-AIF), AIF+ extended with an AO that connects core AIF+ resources to the agents (human, organizational, or artificial) that created them. The S-AIF AO is shown in Figure 3 (the application *ArguBlogging* is described in Section 4.4).

Snaith et al. [2016] define four different but connected roles on the social Argument Web:

(1) *Speakers*: A speaker is a person from which the content of a locution is obtained but not necessarily contributed by the person himself (e.g., a politician whose speech has been analyzed via a tool such as OVA).
(2) *Users*: A registered user of the Argument Web who makes direct contributions. Such contributions are either their own opinion (where they act as a special case of speaker) or interpretations of opinions of others.
(3) *Agents*: An agent is a software representation of speakers (and, by extension, users) in mixed-initiative dialogues. Agents are themselves speakers, with their contributions being distinct from those originally made by the people they represent. As such, agents can subsequently be represented by other agents.
(4) *Arguers*: An arguer is a participant (real or virtual) in a dialogue that is being executed using the Dialogue Game Execution Platform (DGEP; Section 4.2) and presented in an application such as Arvina (Section 4.3). All arguers are either

---

[13]https://www.blogger.com.
[14]https://www.tumblr.com.
[15]https://www.facebook.com.
[16]https://twitter.com.

(real) users or agents; any speaker who is not a user cannot be an arguer and thus to take part in a dialogue must either be represented by an agent or become a user.

When an agent is instantiated to represent a specific person, he obtains a knowledge base from the Argument Web using AIF+ resources attributed to that person (either through direct contributions or analyses contributed by a user). Using AIF+ as a knowledge base allows the agent to construct an argumentative response to questions or challenges in the dialogue.

The second function of the social layer is to manage connections to the social World Wide Web, allowing for seamless integration between Argument Web applications and existing metaphors for user-contributed content such as blogs and social networks. As well as managing user authentication, via technologies and protocols such as OpenID[17] and OAuth,[18] the social layer also translates S-AIF into a human-readable format. A practical example of this is found in ArguBlogging (described in Section 4.4).

4.1.2. Dialogue Games: Formal Systems of Interaction. In general, a dialogue consists of two or more participants (or players). To ensure that the dialogue proceeds in an orderly fashion, it requires a set of rules that specify turn taking, legal moves, and termination and outcome conditions. Influenced by the philosophical dialogue games of Hamblin [1097], Walton [1984], and Walton and Krabbe [1995], Parsons and Jennings [1996], Reed [1998], and McBurney and Parsons [2002] laid the groundwork for specifying computational protocols for argumentative interagent communication. Parsons and Jennings [1996] specify a formal protocol for negotiation between agents looking to find ways to solve problems, whereas Reed [1998] provides a computational account of the Walton and Krabbe [1995] dialogue typology (with the exception of eristics, which models physical conflict), and McBurney and Parsons [2002] formalized the modeling of dialogue games for agent communication. It is these foundations upon which protocols for debating technology are built, using the Dialogue Game Description Language (DGDL) and the DGEP, as described in more detail in Section 4.2.

4.2. Dialogue Game Description and Execution

It is an old philosophical and latterly computational idea that systems of dialogue rules can be codified. The typical approach (see the Foundation for Intelligent Physical Agents (FIPA), e.g., FIPA [1997]) is to define a system ab initio with little regard to commonalities between individual systems. Robertson [2004] has demonstrated that a language for expressing such dialogue systems in general offers significant practical and theoretical advantages and has proposed a lightweight coordination calculus (LCC) to do just this. LCC, however, is designed largely to orchestrate interaction between software agents, and to do so in principle rather than at scale. Crucially, the very lightweightness of LCC also means that a great deal of engineering is required to develop a new nontrivial system. This has limited the applicability of LCC to the implementation of debating technology systems. An alternative that sacrifices lightweight, elegant specification in favor of the advantages of a practical programming language is the DGDL.

4.2.1. The DGDL. The DGDL [Wells and Reed 2012] is a domain-specific language for capturing the properties, rules, and moves of a dialogue game. A DGDL specification consists of three main parts: composition, rules, and interactions. The composition describes the general features of the dialogue game, including the ordering of moves,

---

```
/* — COMPOSITION — */
turns{magnitude:multiple, ordering:strict}
roles{speaker, listener, Participant}
players{min:2, max:2}
player{id:Initiator}
store{id:InitialPoint, owner:shared, structure:set,
    visibility:public, {"The road should be built"}}

/* — RULES — */
rule{id:StartingRule, scope:initial,
{ assign(Initiator,speaker)
    & move(add, future, AgreeReason, {p,q}, Initiator, {inspect(in,{p},InitialPoint)})
    & move(add, future DisagreeReason, {p,q}, Initiator, {inspect(in,{p},InitialPoint)})
    & move(add, future, Question, $Participant, {p}, Initiator,
        {inspect(in, {p}, InitialPoint)})
    & move(add, future, Challenge, $Participant, {p}, Initiator, {inspect(in, {p},
        InitialPoint)})
    }
}

/* — INTERACTIONS — */
interaction{Question, $Participant, {p}, Questioning, {p}, "Do you agree with $p?",
{ move(add, next, AgreeReason, {p,q}, Target)
    & move(add, next, DisagreeReason, {p,q}, Target)
    & move(add, next, Agree, {p}, Target)
    & move(add, next, Disagree, {p}, Target)
    & move(add, next, NoOpinion, {p}, Target)
    & assign(Target, speaker)
    }
}
```

Listing 1. Excerpt from the DGDL specification for the MM game.

the number of moves that can take place in each turn,[19] the roles in the dialogue, and commitment stores. The rules describe specific actions that are to occur either at the end of each move, at the end of each turn, at the beginning of the dialogue, or at the end of the dialogue. Most dialogue games specified in the DGDL will have a starting rule, describing how a dialogue actually commences. The interactions define the moves that participants can make and the effects they have, both on the subsequent progress of the dialogue and the underlying AIF+ structure the dialogue generates.

Listing 1 shows an excerpt from the DGDL specification for the game *MM,*[20] illustrating the component parts. MM is based on the protocol employed in the BBC Radio 4 program *The Moral Maze*[21] but can also be employed as a general-purpose protocol for navigating and contributing to existing debates in a range of domains. It is this protocol that we will use in our running example, introduced in Section 2.

4.2.2. The DGEP. Having specified a dialogue game in the DGDL, this specification can then be processed by the DGEP [Bex et al. 2014]. The DGEP allows participants to take part in dialogues following the rules specified by a DGDL protocol.

The DGEP maintains the legal move list for each participant, based on the defined rules and interactions: a rule is executed when it is in scope and an interaction when it is moved by a player during the game. Once all requirements of a rule can be confirmed, the DGEP instantiates the effects, possibly binding variables with, for

---

[19]In the DGDL, a turn can contain one or more moves. For example, it is player A's turn, and he first makes a claim, immediately followed by an argument.

[20]The full DGDL specification for MM is available at http://arg.tech/mm.

[21]The basic format of *The Moral Maze* is that four "witnesses" are interrogated by a panel to establish the truth or otherwise of the program topic: http://www.bbc.co.uk/moralmaze. The dialogue is a hybrid of inquiry (the participants are attempting to establish the truth of the program topic) and persuasion (the participants all have their own opinions on the topic and attempt to persuade the others of it).

example, matches from the store. These effects may be commitment store operations, role and status updates, and, perhaps most importantly, move availability, indicating moves that become either mandatory or possible at some point in the future. For all legal moves at a given turn, there may be many instantiations (depending on, e.g., the commitments and knowledge base of the participants).

A participant selects an appropriate move from its legal move list and executes it by passing the move identifier back to the DGEP. Participants' replies are handled asynchronously both because of increased robustness and because some dialogue games, such as those in which interruptions are permitted, and those, such as auctions, where many agents take on identical roles, do not impose rigid turn-taking rules.

In addition to updating the legal move list for participants, the execution of a move also updates AIF+ structures in AIFdb. This means that user contributions made using applications built on the DGEP capture not only support and conflict between individual statements but also the argumentative function, thus connecting argument resources with the conversational structure that created them. In this way, as use of Argument Web debating technology tools expands, so do the resources available to make it attractive to new users also grow.

4.2.3. Mixed-Initiative Argumentation with the DGEP. The DGEP allows for *mixed-initiative argumentation* [Snaith et al. 2010], a type of collaborative intelligence [Epstein 2015] or human-agent collective [Jennings et al. 2014]. Mixed-initiative dialogues contain both real (human) and virtual (agent) participants; however, crucially, the DGEP does not distinguish between them. This creates a level playing field in the dialogue and, in principle, allows virtual agents to take the lead and steer the conversation.

To ensure that the DGEP does not distinguish between real and virtual participants, the agents in a mixed-initiative dialogue are not handled by the DGEP itself but instead by a separate agent management system (AMS). The AMS is responsible for creating agents, populating their knowledge bases, and ensuring that they respond at the correct junctures in the dialogue. The connections between the DGEP, the applications it supports, and other Argument Web infrastructure and components are summarized architecturally in Figure 4.

The DGEP provides a range of Web service interfaces, allowing a user to both perform interactions and get information about the current dialogue state (e.g., their list of available moves). These Web services can then be used by either software agents playing the roles of specific participants or by graphical interfaces allowing human users to take part in the discussion.

The DGDL and DGEP allow for rapid development of applications for dialogue-based engagement in argument and debate. Here, we use our running example to demonstrate two such applications: Arvina, which uses mixed-initiative argumentation for navigating and contributing to complex information spaces, and ArguBlogging, which taps into existing metaphors for social interaction on the World Wide Web. Both tools represent two examples of human interfaces to Argument Web infrastructure, which offer unique opportunities to harvest openly available structured data from users as a by-product of something they are doing anyway.

4.3. Arvina: Engagement Through Mixed-Initiative Argumentation

Human interaction with mixed-initiative debates can be supported via an application built upon the DGEP's Web services. One such application is Arvina[22] [Lawrence et al. 2012a]. A human user selects a dialogue game and can invite additional real or virtual participants to take part. Arvina provides two benefits to users; first, it allows them to

---
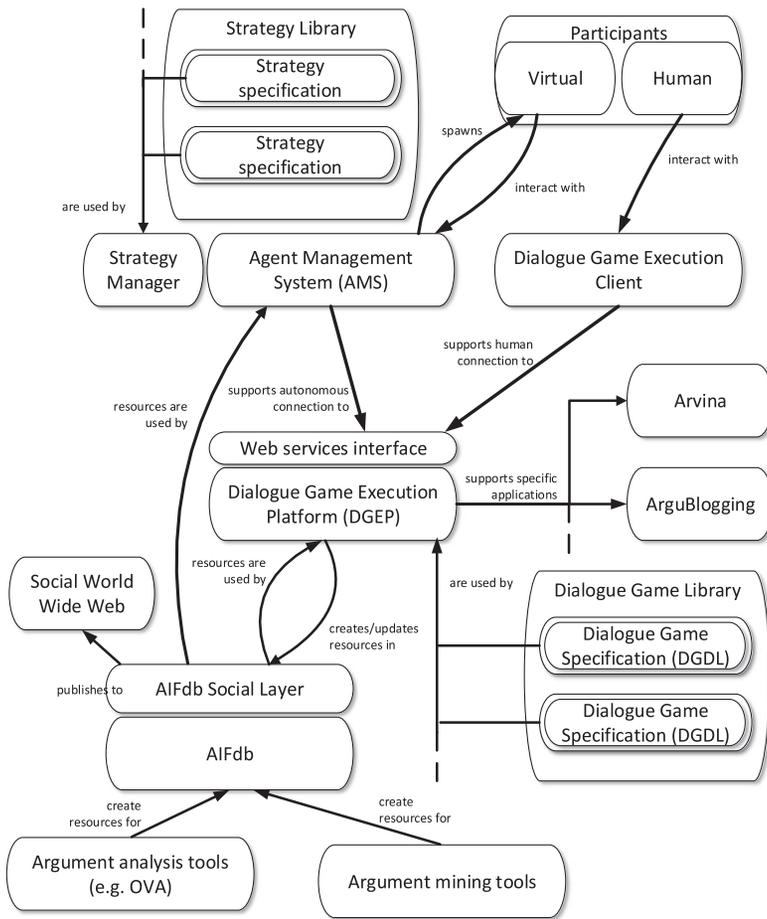
[22]http://arg.tech/arvina-project.

Fig. 4.   The DGEP application architecture.

navigate and engage with a complex debate space in an intuitive way by "interrogating" software agents; second, it allows them to contribute their own opinion to the debate, which in turn can allow them to be represented as a virtual participant in future use by another user.

Figure 5 shows the Arvina interface configured for engaging in CD. The original participants in the dialogue (Ken Kelly and the moderator) are represented by agents, whereas the current (human) user (Jane Smith) has a set of moves available that will allow her to contribute her own opinion and interrogate the agents. Interactive analysis of CD using mixed-initiative argumentation in Arvina can be used by both members of the public and governmental agency specialists alike. Figure 6 presents the result of the contribution from Arvina stored in AIFdb.

For members of the public, mixed-initiative argumentation allows them to navigate a potentially complex debate in a naturalistic way, and allows them to contribute their own opinions. Specialists need to take into account not only the opinions of citizens but also numerous other factors (funding, technical, and legal issues, etc.). By supporting interaction with an analysis of complex information spaces, mixed-initiative argumentation can enhance engagement and has the potential to assist in better decision making.
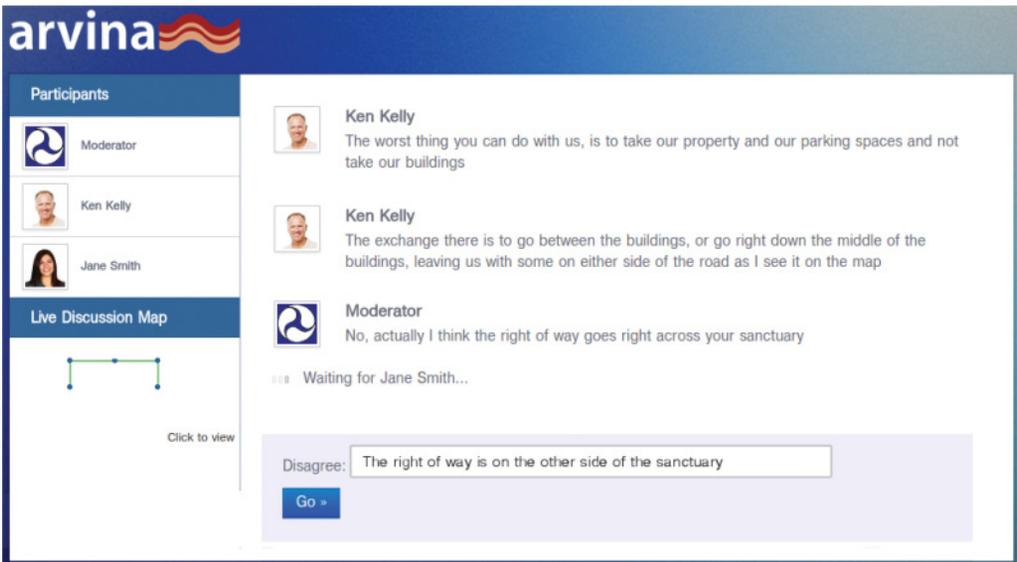
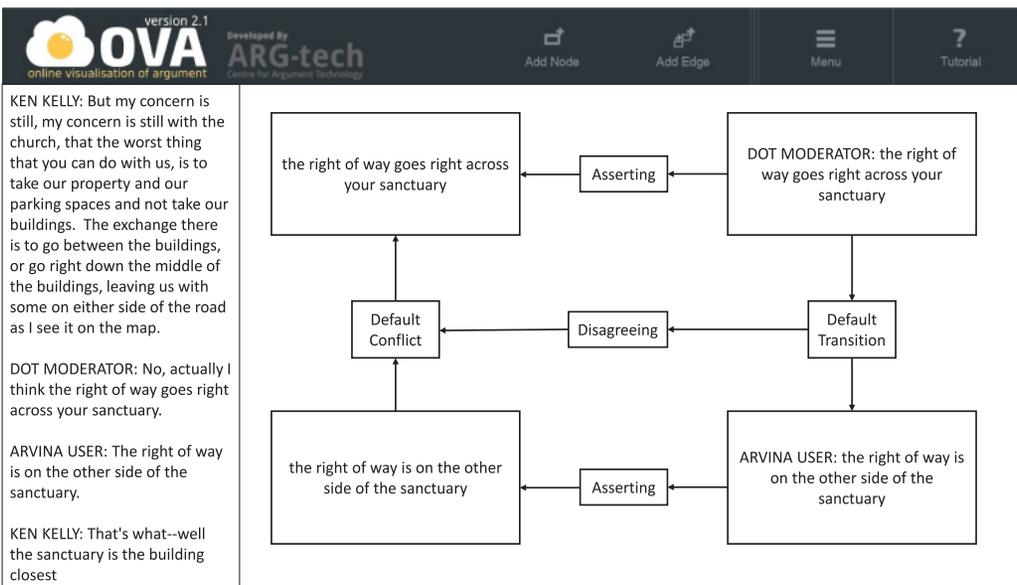Fig. 5.   Arvina interface for a mixed-initiative dialogue.

Fig. 6.   Fragment of CD modified by Arvina interaction and displayed in OVA+.

4.4. ArguBlogging: Lowering the Barrier to Engagement

ArguBlogging[23] is a lightweight, easy-to-use[24] application that aims to open up the Argument Web to regular Internet users, particularly bloggers [Bex et al. 2014; Snaith et al. 2012]. It provides a new interface for bloggers to respond to online opinions while

---

[23]http://argublogging.com.
[24]Based on a survey carried out in Bex et al. [2014].

building there as he swung his arch there, but to tie it all together I think he did an
excellent job being an engineer myself, I do appreciate the handiwork there.
But my concern is still, my concern is still with the church, that the worst thing that you
can do with us, is to take our property and our parking spaces and not take our
buildings. The exchange there is to go between the buildings, or go right down the
middle of the buildings, leaving us with some on either side of the road as I see it on
the map.

Moderator: No, actually I think the right of way goes right across your
sanctuary.

Ken Kelly: That's what-- ...osest to Huffine
Mill Road and the educat...

Moderator: Okay.

Ken Kelly: So we would ...ice, and the
educational space and l...there. And I guess I
know that as long as we ...a long ways off, I guess
my concern is in the lin...t'll we not do. We have a
ball-field there, do we ...e this, do we stop that, and
all these kind of things ...r a septic tank for our
bathrooms, will be left, ...e road of all of our major
property, where it looks ...also a concern. But with all
those in consideration ...l, I, that's my concern is
being in limbo, can you ...

Moderator: Well, you can ...and get approvals
and so forth, if there ...like that would definitely
create a hardship for th...at is called Advanced
Acquisition, so that you ...could be making plans
for what they need to do ...nute, till the right of way is
acquired under the norma...ays make that request to
our Right of Way Branch ...ion.

Ken Kelly: Now is that a...to approval or?

ArguBlogging

I [Agree ‡] because

It would be a mere shell of the church
that was left.

A mere shell

☑ Tumblr ☐ Blogger [**More**]

Send

Fig. 7.   ArguBlogging widget.

simultaneously contributing to the Argument Web via the social layer. This allows
debate and discussion to be constructed across blogs, linking existing and new online
resources to form distributed, structured conversations.

The ArguBlogging tool is installed as a bookmarklet[25] in the user's browser. The user
highlights the relevant piece of text on a Web page and clicks the bookmarklet, which
causes the ArguBlogging widget to be rendered on the Web page. The widget, shown in
Figure 7, presents a simple interface that allows the user to submit a response to the
highlighted text along with a title for the blog post the response creates. Two blogging
platforms are currently supported: Tumblr[26] and Blogger.[27]

When a user publishes to his blog(s) with ArguBlogging, an "Argue" button is ap-
pended to the post(s), which when clicked launches the ArguBlogging tool for respond-
ing to that post. This promotes further use of the tool and with it a continuation of the
debate in a structured fashion.

At present, ArguBlogging executes a simple dialogue protocol that allows users to
agree or disagree with specific sections of text in online content. In principle, however,
the modular approach to the Argument Web means that it could provide a more so-
phisticated protocol in the future, such as permitting challenges to a published opinion
(i.e., seeking justification) as opposed to simple disagreement.

---

[25]Bookmarklets are small pieces of software that extend the capabilities of Web browsers: http://book
marklets.com.
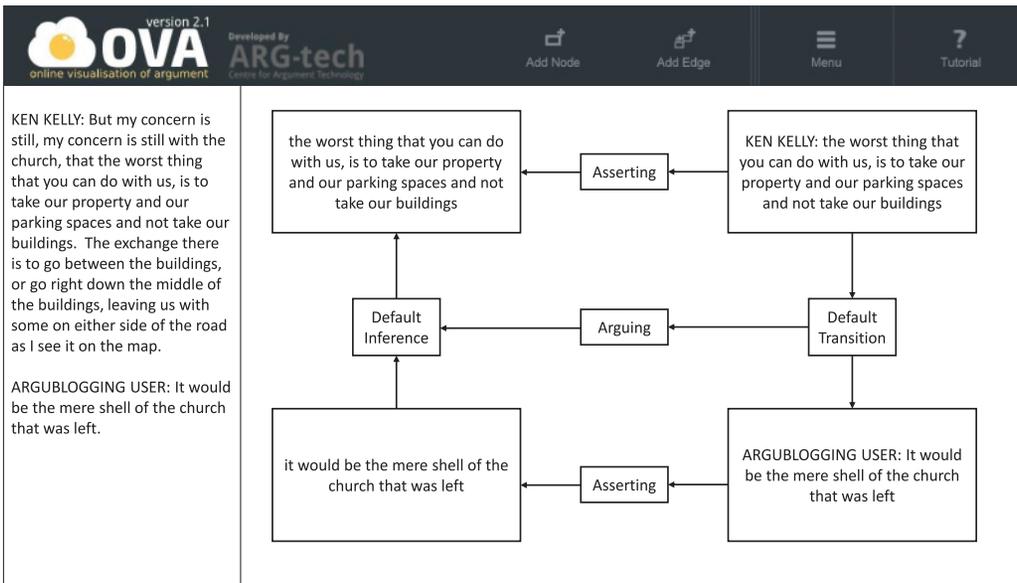[26]https://tumblr.com.
[27]https://blogger.com.

Fig. 8.   Fragment of CD modified by ArguBlogging interaction and displayed in OVA+.

In the context of CD, the ArguBlogging tool creates expanded debate space, where face-to-face engagement can be enriched further by online interaction, providing already structured data and building argument network. When the transcript of a public hearing is published online, further arguments and opinions can be contributed by citizens (whether they attended the hearing or not). These contributions can in turn be examined by the public body as part of the decision-making process. Figure 7 shows ArguBlogging being used to contribute to the CD from Example (1) (see Section 2) by interacting directly with the online transcript of the original debate. Figure 8 presents the result of this interaction, stored in AIFdb. The argument structure has changed with the introduction of an additional premise and associated dialogical information from the ArguBlogging user. Moreover, this premise is marked as coming from a different speaker, which acts as an indicator that the case is important to more citizens. ArguBlogging in this case allows for joint argument creation by various groups of citizens, including those not present at the original public meeting.

5. DEBATE ANALYTICS

Large and complex debates generate vast quantities of information that is best understood when presented in a visual form. One of the most prominent systems for visualizing argument and debate is DebateGraph,[28] which has been used by the U.S. White House, UK Foreign and Commonwealth Office, and CNN for visualizing debates that are complex and wide ranging. Furthermore, Kirschner et al. [2003] tackle sensemaking through visualization of argument, whereas more recently the Election Debate Visualisation (EDV) project has sought to address these issues in the context of televised election debates [Plüss and De Liddo 2015; Coleman and Moss 2016], and the VisArgue project in the context of deliberative democracy [El-Assady et al. 2016].

Although these tools present argument and debate in an easy-to-navigate way, what is lacking is a the ability to support insight into the character of debates as a whole.

---

[28]http://debategraph.org.

There are no tools to naturally identify the key points of the debate, to monitor the level of participation of different contributors, or quite simply to indicate which side "won" (if any). Furthermore, complex phenomena such as divisiveness and popularity are not immediately obvious but are important in understanding the dynamics, and ultimately the outcome, of debates. Although attempts have been made to analyze user interactions and roles in debates [Jain and Hovy 2013; Jain et al. 2014], these do not provide insight into the debates themselves. To address this problem, we have developed a suite of debate analytics, grouped into three categories: simple, semantics based, and graph based.

5.1. Simple Analytics

Simple analytics examine basic features of an AIF+ graph such as the number of edges into or out of a node.

5.1.1. Statistics. From an AIF+ graph, we extract simple statistics: the total number of I-nodes and S-nodes, the total number of words, and the average node length (i.e., total number of words/total number of I-nodes).

5.1.2. Key Points. A key point in a debate will have many responses, represented in an AIF+ graph by many incoming edges to the I-node containing that point. We count the incoming edges to each I-node and identify the top five as the key points in the debate.

5.1.3. Participation. For each participant, count the number of locutions they have made and represent them in a pie chart. This provides an easy way of identifying which participants were most and least active in the debate.

5.1.4. Stimulating. A point of debate is stimulating if it receives many responses, either to agree or disagree. Given an AIF+ graph, we count the number of responses to each locution

5.2. Semantics-Based Analytics

Semantics-based analytics use Dung-style semantics to determine the acceptability of a participant's arguments. An AIF+ graph is translated into ASPIC$^+$, then using TOAST, a Dung-style abstract argumentation framework is derived and evaluated.

5.2.1. Defended. For a given argument A, if all of A's counterarguments are successfully attacked, then A is defended. This is a direct use of Dung's concept of argument defense.

5.2.2. Sway. Given two participants $P1$ and $P2$, if $P1$ has more acceptable arguments than $P2$ (under the same semantics), then $P1$ is said to carry more sway.

5.3. Graph-Based Analytics

The rich expressivity of AIF+ graph structures offers a wide variety of analytical insights. We focus here on two examples, each of two different kinds of such insights. The first are two examples of properties of arguments: divisiveness and popularity; the second are two examples of properties of arguers: trollishness and sycophancy. In each case, different intuitions about these notions give rise to different definitions that can be captured and distinguished formally.

To obtain operational definitions and provide measurable criteria for these metrics, we found our concepts on properties of AIF+ graphs.

We construe the argumentation analysis as a directed graph, $G = (V, E)$, in which vertices (V) are either propositions, locutions (a special subclass of propositions), or relations between propositions, and those relations are either support (pro arguments),

conflict (con arguments), illocution, or transition,[29] captured by a function R that maps $V \mapsto \{prop, locution, support, conflict, illocution, transition\}$ and edges exist between them $E \subset V \times V$. Every relation may be further subtyped (classifying different types of support or illocution, etc.); however, to keep the notation uncluttered, we use a separate set of functions $R_{support}, R_{conflict}, R_{illocution}, R_{transition}$ (abbreviated R_s, R_c, R_i, and R_t) to encapsulate these taxonomies. Finally, in same cases, we need to identify who is responsible for a given locution. For this, we require a function quite separate to the graph, which acts on vertices for which $R(v) = locution$ and identifies such utterers: $U \mapsto \{u_1, u_2, \ldots u_n\}$. For syntactic convenience, we refer to the number of edges at (i.e., the order of) a vertex v as $|v|$ and add a superscript to indicate whether we are interested in the number of incoming or outgoing edges, and a subscript to indicate constraints on the values of $R(v')$ of the vertex v' to which v is connected in each case. Thus, for example, $|v|^{in}_{R(v')=support}$ is the number of edges incoming to v originating at vertices of type $support$.

5.3.1. Divisiveness. A first estimate, D_1, of the divisiveness of a proposition v might be based upon the relative number of incoming supports and conflicts. A proposition with many of both might be taken to be divisive, whereas few of either might suggest only limited divisiveness. Such an intuition might straightforwardly be calculated as follows:

$$Divisive_1(v) = |v|^{in}_{R(v')=support} * |v|^{in}_{R(v')=conflict}. \tag{1}$$

Alternatively, divisiveness might be taken to be intrinsically concerned with a pair of propositions v_1 and v_2 that are in conflict and might be a measure of the amount of support on both sides. Thus, given $(v_1, v_c) \in E$ and $(v_c, v_2) \in E$ and $R(v_c) = conflict$,

$$Divisive_2(v_1, v_2) = |v_1|^{in}_{R(v')=support} * |v_2|^{in}_{R(v')=support}. \tag{2}$$

Our third concept of divisiveness combines the concepts introduced in D_1 and D_2, as it calculates the value for a single proposition but considers every other with which it is in conflict. By this measure, every proposition v_2 that is in conflict with v (i.e., for which there is an edge either outgoing from v through a conflict v_c to v_2, or in the other direction, or both) is considered using $Divisive_2$ and the sum over all such v_2 is calculated:

$$Divisive_3(v) = \sum_{\substack{\forall v_2 s.t. \\ ((v_2,v_c),(v_c,v)\in E, or \\ (v,v_c),(v_c,v_2)\in E) and \\ R(v_c)=conflict}} Divisive_2(v, v_2). \tag{3}$$

5.3.2. Popularity. As with divisiveness, there are several different interpretations of popularity. The first is that a popular notion is one that has been claimed many times. We take "claiming" to be an illocution of type *assert*-ing as given by $R_{illocution}$.

$$Popular_1(v) = |v|^{in}_{R_i(v')=assert} \tag{4}$$

This will not do, however, because the same person could assert something many times, without it being at all popular. We need to further constrain distinctness among speakers and directly use the cardinality of the set of appropriate utterers as our

---

[29]Informally, a transition is what occurs between dialogue moves.

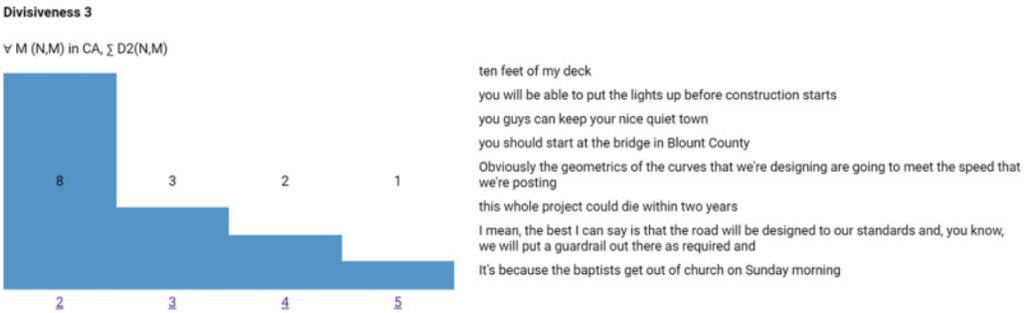Fig. 9.   Excerpt of the statistics for the CD corpus.

Fig. 10.   Divisiveness score for the CD corpus.

measure of popularity:

$$Popular_2(v) =_{def} |\{u \mid U(l) = u, \forall l \text{ subject to } (l, r), (r, v) \in E \text{ and } R(r) = locution \text{ and } R_i(r) = assert\}|.$$

(5)

A final extension reflects the fact that popularity may be judged through several illocutionary forces:

$$Popular_3(v) =_{def} |\{u \mid U(l) = u, \forall l \text{ s.t. } (l, r), (r, v) \in E \text{ and } R(r) = locution \text{ and } R_i(r) \in \{assert, agree\}\}|.$$

(6)

5.4. Using Analytics to Understand CD

The measures proposed in this section, when implemented, can serve as a tool for gaining quick insight into the structured data. Figure 9 shows an excerpt from the simple analytics statistics for an entire corpus of CDs. We observe, for instance, that assertive questions (AQ), in which a speaker conveys her own opinion in a question, are much more common than pure questions (PQ), in which she does not.

From the graph analytics, Figure 10 shows the results of implementing the $Divisive_3$ score to compute over the CD corpus. We can see that the eight propositions shown have the highest divisiveness scores (≥ 2). This tool allows the representatives of a governmental agency to quickly spot the most divisive issues, supporting understanding of the citizens' input.

6. CONCLUSIONS

Argument is complex and multifaceted, with challenges at philosophical, linguistic, and computational levels, involving a wide variety of phenomena (interaction, emotion, inference, setting, and character, to mention only a few). It is no wonder that the catalysis of the research area that tries to bring all of these together is challenging, and even less surprising that the infrastructure for supporting this area is complex and wide ranging. The Argument Web takes a slice through many other research areas, aiming to unify those parts that are apposite and sufficiently mature to be incorporated into engineered systems, while recognizing that it is merely a first step and must therefore admit of extension, expansion, and evolution.

The goal here has been to show how the foundations of the Argument Web are laid out, how they lead naturally to infrastructure that is applicable to both more and less academic domains, and then how that infrastructure opens a range of new and exciting avenues in areas such as mixed-initiative argumentation and large-scale debate analytics. This article has demonstrated for the first time that the different components that focus on sensemaking, engagement, and analytics can work in concert as a suite of debating technologies for dialogical argument, with each providing motivation and a rich environment for the others.

Although still in its early days, the ecosystem provided by the Argument Web provides many distinct niches and is attracting contributions from a wide range of researchers in AI, computational linguistics, software engineering, and beyond, and the very diversity of that mix augurs well for its longevity and success.

ACKNOWLEDGMENT

REFERENCES

Floris Bex, Thomas Gordon, John Lawrence, and Chris Reed. 2012. Interchanging arguments between Carneades and AIF. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA'12)*. 390–397.

Floris Bex, John Lawrence, and Chris Reed. 2014. Generalising argument dialogue with the dialogue game execution platform. In *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA'14)*. 141–152.

Floris Bex, John Lawrence, Mark Snaith, and Chris Reed. 2013a. Implementing the Argument Web. *Communications of the ACM* 56, 10, 951–989.

Floris Bex, Sanjay Modgil, Henry Prakken, and Chris Reed. 2013b. On logical reifications of the argument interchange format. *Journal of Logic and Computation* 23, 5, 66–73.

Katarzyna Budzynska, Mathilde Janier, Juyeon Kang, Chris Reed, Patrick Saint-Dizier, Manfred Stede, and Olena Yaskorska. 2014. Towards argument mining from dialogue. In *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA'14)*. 185–196.

Martin W. A. Caminada, Walter A. Carnielli, and Paul E. Dunne. 2012. Semi-stable semantics. *Logic and Computation* 22, 5, 1207–1254.

Carols Chesñevar, Jarred McGinnis, Sanjay Modgil, Iyad Rahwan, Chris Reed, Guillermo Simari, Matthew South, Gerar Vreeswijk, and Steven Willmott. 2006. Towards an argument interchange format. *Knowledge Engineering Review* 21, 4, 293–316.

Stephen Coleman and Giles Moss. 2016. Rethinking election debates: What citizens are entitled to expect. *International Journal of Press/Politics* 21, 1, 3–24.

Anna De Liddo and Simon Buckingham-Shum. 2014. New ways of deliberating online: An empirical comparison of network and threaded interfaces for online discussion. In *Proceedings of the International Conference on Electronic Participation*. 90–101.

Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77, 321–357.

Phan Minh Dung, Paolo Mancarella, and Francesca Toni. 2007. Computing ideal sceptical argumentation. *Artificial Intelligence* 171, 10–15, 642–674.

Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. 2008. ASPARTIX: Implementing argumentation frameworks using answer-set programming. In *Proceedings of the 24th International Conference on Logic Programming (ICLP'08)*. 734–738.

Mennatallah El-Assady, Valentin Gold, Carmela Acevedo, Christopher Collins, and Daniel Keim. 2016. ConToVi: Multi-party conversation exploration using topic-space views. *Computer Graphics Forum* 35, 3, 431–440.

Stefan Ellmauthaler and Hannes Strass. 2014. The DIAMOND system for computing with abstract dialectical frameworks. In *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA'14)*. 233–240.

Susan L. Epstein. 2015. Wanted: Collaborative intelligence. *Artificial Intelligence* 221, 36–45.

FIPA. 1997. *FIPA 97 Specification Part 2: Agent Communication Language*. FIPA, Geneva, Switzerland.

Thomas F. Gordon. 2011. The policy modelling tool of the IMPACT argumentation toolbox. In *Proceedings of the Jurix Workshop on Modelling Policy-Making (MPM'11)*. 29–38.

Thomas F. Gordon and Nikos Karacapilidis. 1997. The Zeno argumentation framework. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*. ACM, New York, NY, 10–18.

Thomas F. Gordon and Doulgas N. Walton. 2006. The Carneades argumentation framework. In *Proceedings of the 1st International Conference on Computational Models of Argument (COMMA'06)*. 195–207.

Iryna Gurevych, Eduard H. Hovy, Noam Slonim, and Benno Stein. 2016. Debating technologies (Dagstuhl Seminar 15512). *Dagstuhl Reports* 5, 12, 18–46. DOI:http://dx.doi.org/10.4230/DagRep.5.12.18

Charles L. Hamblin. 1970. *Fallacies*. Chaucher Press, Bungay, Suffolk, UK.

Siddharth Jain, Archna Bhatia, Angelique Rein, and Eduard Hovy. 2014. A corpus of participant roles in contentious discussions. In *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC'14)*. 1751–1756.

Siddharth Jain and Eduard Hovy. 2013. Determining leadership in contentious discussions. In *Proceedings of the International Workshop on Social Multimedia Research (SMMR'13)*. IEEE, Los Alamitos, CA, 1–6.

Mathilde Janier, John Lawrence, and Chris Reed. 2014. OVA+: An argument analysis interface. In *Proceedings of the 5th International Conference on Computational Models of Argument (COMMA'14)*. 463–464.

N. R. Jennings, L. Moreau, D. Nicholson, S. D. Ramchurn, S. Roberts, T. Rodden, and A. Rogers. 2014. Human-agent collectives. *Communications of the ACM* 57, 12, 80–88.

Paul A. Kirschner, Simon J. Buckingham-Shum, and Chad S. Carr (Eds.). 2003. *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Springer-Verlag, London, UK.

Mark Klein and Luca Iandoli. 2008. Supporting collaborative deliberation using a large-scale argumentation system: The MIT Collaboratorium. MIT Sloan Research Paper. MIT, Cambridge, MA.

John Lawrence, Floris Bex, and Chris Reed. 2012a. Dialogues on the Argument Web: Mixed initiative argumentation with Arvina. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA'12)*. 513–514.

John Lawrence, Floris Bex, Chris Reed, and Mark Snaith. 2012b. AIFdb: Infrastructure for the Argument Web. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA'12)*. 515–516.

John Lawrence and Chris Reed. 2016. Argument mining using argumentation scheme structures. In *Proceedings of the 6th International Conference on Computational Models of Argument (COMMA'16)*.

Ran Levy, Yonatan Bilu, Daniel Hershcovich, Ehud Aharoni, and Noam Slonim. 2014. Context dependent claim detection. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING'14)*. 1489–1500.

Marco Lippi and Paolo Torroni. 2015. Argumentation mining: State of the art and emerging trends. *ACM Transactions on Internet Technology* 16, 2, 10:1–10:25.

Peter McBurney and Simon Parsons. 2002. Dialogue games in multi-agent systems. *Informal Logic* 22, 3, 257–274.

Marie-Francine Moens. 2013. Argumentation mining: Where are we now, where do we want to be and how do we get there? In *Proceedings of the 5th Forum on Information Retrieval Evaluation (FIRE'13)*. ACM, New York, NY, 2:1–2:6.

Simon Parsons and Nicholas Jennings. 1996. Negotiation through argumentation—a preliminary report. In *Proceedings of the 2nd International Conference on Multiagent Systems (ICMAS'96)*. 267–274.

Andreas Peldszus and Manfred Stede. 2013. From argument diagrams to argumentation mining in texts: A survey. *International Journal of Cognitive Informatics and Natural Intelligence* 7, 1, 1–31.

Andreas Peldszus and Manfrede Stede. 2015. Joint prediction in MST-style discourse parsing for argumentation mining. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 938–948. http://aclweb.org/anthology/D15-1110

Brian Plüss. 2016. Communication of debate aspects to different audiences. In *Report of Dagstuhl Seminar 15512, Debating Technologies*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 29.

Brian Plüss and Anna De Liddo. 2015. Engaging citizens with televised election debates through online interactive replays. In *Proceedings of the ACM International Conference on Interactive Experiences for TV and Online Video (TVX'15)*. ACM, New York, NY, 179–184.

John L. Pollock. 1987. Defeasible reasoning. *Cognitive Science* 11, 481–518.

Henry Prakken. 2010. An abstract framework for argumentation with structured arguments. *Argument and Computation* 1, 2, 93–124.

Iyad Rahwan, Fouad Zablith, and Chris Reed. 2007. Laying the foundations for a world wide Argument Web. *Artificial Intelligence* 171, 897–921.

Andrew Ravenscroft, Simon McAlister, and Enzian Baur. 2006. *Development, Piloting and Evaluation of InterLoc: An Open Source Tool Supporting Dialogue Games in Education*. Final Project Report. Learning Technology Research Institute, London Metropolitan University, UK & Joint Information Systems Committee, Bristol, UK. http://www.interloc.org.uk/publications_files/ILoc_FR_Fc_v3.pdf.

Chris Reed. 1998. Dialogue frames in agent communication. In *Proceedings of the 3rd International Conference on Multi-Agent Systems (ICMAS'98)*. IEEE, Los Alamitos, CA, 246–253.

Chris Reed and Glenn Rowe. 2004. Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools* 13, 961–980.

Chris Reed and Douglas Walton. 2005. Towards a formal and implemented model of argumentation schemes in agent communication. *Autonomous Agents and Multi-Agent Systems* 11, 2, 173–188.

Chris Reed, Simon Wells, Joseph Devereux, and Glenn Rowe. 2008. AIF+: Dialogue in the argument interchange format. In *Proceedings of the 2nd International Conference on Computational Models of Argument (COMMA'08)*. 311–323.

Ruty Rinott. 2016. The role of evidence in debates. In *Report of Dagstuhl Seminar 15512, Debating Technologies*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 31.

Ruty 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*. 440–450. https://aclweb.org/anthology/D/D15/D15-1050

David Robertson. 2004. Multi-agent coordination as distributed logic programming. In *Proceedings of the 2004 International Conference on Logic Programming*. 416–430.

John R. Searle. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, UK.

John R. Searle and Daniel Vanderveken. 1985. *Foundations of Illocutionary Logic*. Cambridge University Press, Cambridge, UK.

Noam Slonim. 2016. What is debating technologies. In *Report of Dagstuhl Seminar 15512, Debating Technologies*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 29.

Mark Snaith, Floris Bex, John Lawrence, and Chris Reed. 2012. Implementing ArguBlogging. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA'12)*. 511–512.

Mark Snaith, John Lawrence, and Chris Reed. 2010. Mixed initiative argument in public deliberation. In *Proceedings of the 4th International Conference on Online Deliberation (OD'10)*. 2–13.

Mark Snaith, Rolando Medellin, John Lawrence, and Chris Reed. 2016. Arguers and the Argument Web. In *Argument Technologies*, F. Paglieri, C. Reed, F. Bex, N. Green, and F. Grasso (Eds.). College Publications, London, UK.

Mark Snaith and Chris Reed. 2012. TOAST: Online ASPIC+ implementation. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA'12)*. 509–510.

Christian Stab and Ivan Habernal. 2016a. Detecting argument components and structures. In *Report of Dagstuhl Seminar 15512, Debating Technologies*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 32.

Christian Stab and Ivan Habernal. 2016b. Existing resources for debating technologies. In *Report of Dagstuhl Seminar 15512, Debating Technologies*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 32.

Yuqing Tang, Timothy J. Norman, and Simon Parsons. 2009. A model for integrating dialogue and the execution of joint plans. In *Proceedings of the 8th International Conference of Autonomous Agents and Multiagent Systems*. 883–890.

Matthias Thimm. 2014. Tweety—a comprehensive collection of Java libraries for logical aspects of artificial intelligence and knowledge representation. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*. 528–537.

Stephen E. Toulmin. 1958. *The Uses of Argument*. Cambridge University Press, Cambridge, UK.

Frans H. van Eemeren, Bart Garssen, Erik C. W. Krabbe, A. Francisca Snoeck Henkemans, Bart Verheij, and Jean H. M Wagemans. 2014. *Handbook of Argumentation Theory*. Springer, Netherlands.

Tim van Gelder. 2007. The rationale for rationale. *Law, Probability and Risk* 6, 1–4, 23–42.

Gerard A. Vreeswijk. 1997. Abstract argumentation systems. *Artificial Intelligence* 90, 225–279.

Marilyn A. Walker, Pranav Anand, Jean E. Fox, Rob Abbott, and Joseph King. 2012. A corpus for research on deliberation and debate. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC'12)*. 812–817.

Douglas N. Walton. 1984. *Logical Dialogue—Games and Fallacies*. University Press of America, Lanham, MD.

Douglas N. Walton and Erik C. W. Krabbe. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, New York, NY.

Simon Wells and Chris A. Reed. 2012. A domain specific language for describing diverse systems of dialogue. *Journal of Applied Logic* 10, 4, 309–329.

John H. Wigmore. 1931. *The Principles of Judicial Proof* (2nd ed.). Little, Brown & Co., Boston, MA.