

An Ontology for Argumentation on the Social Web: Rhetorical Extensions to the AIF

Tom BLOUNT ^a, David E. MILLARD ^a and Mark J. WEAL ^a
^a *Web and Internet Science, University of Southampton, UK*

Abstract. A key area in the research agenda of modelling argumentation is to accurately model argumentation on the social web. In this paper we propose additional extensions to our ontology for argumentation on the social web (which integrates elements of the Argument Interchange Format and the Semantically Interlinked Online Communities project) for the purposes of modelling social and rhetorical tactics used in eristic or irrational arguments. We then present a review of these extensions from a panel of experts in the fields of argumentation modelling, web science, philosophy and open and linked data and discuss the value of modelling social argument, the challenges faced to create usable and accurate models and the completeness, clarity and consistency of our proposed additions.

Keywords. argumentation, rhetoric, social web, social media, ASWO, AIF, SIOC

1. Introduction

The social web and social media describe the relationships and communities that form over the world wide web, and the way in which people share content, ideas and information. As the social web becomes more and more ubiquitous, the potential for using it to investigate how truly massive communities interact, communicate and argue increases dramatically. A key area in the research agenda of modelling argumentation is to accurately model argumentation on the social web [1].

Currently, the majority of argumentation modelling tools and ontologies are primarily geared towards formal, rather than informal, argumentation. This approach is highly suited towards AI-based methods and can allow for reasoning over arguments to determine the final outcome, or the correct course of action. However, it neglects the set of informal social argumentation that, while virtually impossible to reason over, represents an equally valuable area of argumentation, particularly on the web. Rising levels of e-bile make understanding how otherwise civil discussions can evolve to turn abusive and toxic an important topic to consider [2].

In this work we build on our previous work of bringing together the Semantically Interlinked Online Communities (SIOC) project and the Argument Interchange Format (AIF) [3] and our extensions of this model in the Argumentation on the Social Web Ontology (ASWO) to incorporate rhetorical attacks and declarations of support [4] with additional features to capture some of the extra-logical tactics used in informal argumen-

tation. We then conducted an expert review of these additions to determine how they affected the clarity, completeness and consistency of the ontology, and the overall inherent value in attempting to model this form of social argumentation.

2. Background

Argumentation can, very broadly, be separated into two categories: dialectic, and eristic. The terms dialectic and eristic were coined in Ancient Greece to describe modes of argumentation with different goals and were popularised in Plato's *Republic* [5], and more recently by Walton and Krabbe [6,7]. A dialectic argument takes the form of two or more parties engaged in rational discourse with the aim of either discovering the truth behind a particular matter, or formulating a solution or resolution for a set of circumstances [8]. For example, an academic presenting their findings to an audience of their peers and rationalising that they are indeed valid, is an example of a dialectic argument, but so too can be a group of friends trying to decide on the best place to have lunch. These arguments tend to rely heavily on the weight of facts and evidence, although personal preference can still hold some sway (for example, a free market vs. protectionism or take-away vs. a restaurant). In contrast, an eristic argument is an argument in which there is no clear resolution in the minds of the participants: they are not motivated by solving a problem, or convincing their opponent [8]. Instead, they may be quarrelling for its own sake as a form of catharsis [9], or to be seen to "win" the argument in the eyes of any spectators [10]. As a result, these arguments favour more emotive language and facts may be deliberately distorted to serve a participant's agenda.

Many theoretical models of argumentation are based on the assumption of a dialectic argument, which is useful when building systems to aid automated reasoning to discover the final resolution to a discussion. However, on the social web there is a clear proliferation of eristic argumentation that often will not have a resolution. Nonetheless, this style of argumentation is also important to consider.

The social web presents a number of challenges for extracting and analysing arguments, particularly due to the lack of clear "indicators" of argument or structure. This problem is compounded by the type of language used; often consisting of highly informal language, incorporating quickly evolving slang and irregular punctuation and grammar [11].

3. Existing Models

3.1. Argument Interchange Format

The Argument Interchange Format (AIF) is a framework for representing argumentation as a directed graph [12], modelling information "nodes" and the relationships (such as inference or conflict) between them. In their work on an extension to the AIF, dubbed AIF+, Reed et al. differentiate between these logical relations and the actual words spoken during the debate [13] and introduce a web-based tool, Online Visualisation of Argument (OVA+), to annotate, display and share argumentation on the web [14].

Information nodes (*I-nodes*) represent a (purported) piece of information, data, or claim. Scheme nodes denote a logical connection between information nodes,

whether an inference (*RA-node*), a conflict (*CA-node*), or a value preference (*PA-node*). Illocutionary-Anchor nodes (*YA-nodes*) tie the information and logical structure of an argument with the spoken or written locution. Locution nodes (*L-nodes*) represent the actual words that are spoken or written by participants. Transition nodes (*TA-nodes*) represent links between locutions. However, this is adapted by the ASWO to instead denote locutions that do not add information nodes, but still further the debate (such as prompting for more details, evidence, etc.).

3.2. *Semantically Interlinked Online Communities*

The Semantically Interlinked Online Communities project (SIOC), a semantic-web vocabulary for representation social media, aims to enable the cross-platform, cross-service representation of data from the social web [15]. This allows for semantic representations of Sites, which hold Forums, which contain Posts, authored by a UserAccount (explicitly *not* a person, as a person can own and manage more than one UserAccount). SIOC also allows the modelling of replies between posts.

4. Proposed Additions

Previously, we examined how to link the AIF and SIOC to provide further contextual information about arguments on social media [3]. We now propose several additional nodes to aid modelling rhetorical or “extra-logical” argument with the ASWO.

One of the key additions is the Persona node (*P-nodes*): this represents the “character” that a person assumes during the discussion. For example, a person may argue in a different fashion in a debate about music than they would about technical expertise. This allows one UserAccount to have many Personas where necessary. The inverse, linking one Persona to multiple UserAccounts, is also possible and can be used to represent a participant attempting to artificially solidify their position by creating multiple accounts.

Faction and Audience nodes (*F-* and *A-nodes*) represent abstract groups of Personas; a Faction is any grouping of Personas and can potentially include those outside the Thread, whereas the Audience represents all Personas currently participating in, or observing, the discussion.

Personal Support and Personal Conflict nodes (*PS-* and *PC-nodes*) allow a means of representing support and attack that does not rely on logic and instead uses rhetorical force, social pressure or some other form of “extra-logical” tactic.

Implication nodes (*Im-nodes*) allow analysts to represent a participant implying a relationship between two (or more) nodes, such as Personas. These can be combined with the Personal Support/Conflict nodes to indicate whether the implication is positive or negative.

5. Expert Review

Six experts, from the fields of argumentation systems, web science, philosophy, and linked data, were chosen to review these proposed additions to the model. Experts A and B have a background in argumentations systems and modelling argumentation, and are familiar with the AIF. Expert A is a computer science lecturer whose research is con-

1. **User 1:** *The tech industry is often biased against women*
User 2: *@User1 You would say that, you're a woman*
User 3: *@User1 **** off and die you ****ing nazi before I come and **** you up*

2. **User 1:** *Guns killed 33,000 people last year; they need to be banned*
User 2: *@User1 And a lot of those were minors*
User 3: *@User2 According to who?*

3. **User 1:** *What does Barack Obama call illegal aliens? Undocumented democrats!*
User 2: *@User1 You're so stupid you probably went to the library to find Facebook*

Figure 1. The three argumentation samples the experts were asked to model

cerned with argumentation-based models of communication and formal reasoning, with interests in AI and behaviour change. Expert B is a post-doctoral researcher with degrees in library and information science, mathematics, and liberal arts whose thesis focused on the problem of analysing, integrating, and reconciling information in online discussions. Expert C is a web-science graduate student, researching the relation between social structures in virtual worlds and the real world, with a focus on practices of gender and power. Expert D is a philosophy graduate student, specialising in ethics, moral obligations and with a background in argumentation and formal logic. Experts E and F are specialists in the area of open and linked data working in web and data innovation and development. Expert E is an institutional open data specialist and Expert F is a senior technical specialist.

Each expert was provided with a document describing the background of this area and an overview of the existing models. They were then asked to model three argumentation samples shown in Figure 1, illustrating a variety of different rhetorical structures, by speaking aloud and/or sketching with pen and paper. They were then shown the additions to the model, and asked to model the three argumentation samples again. They were then asked a series of semi-structured question aimed to evaluate their thoughts on how best (and whether) to model social (and anti-social) argumentation, the completeness of the ontology, the clarity of the ontology and the consistency of the ontology.

5.1. Results and Analysis

Table 1 shows an overview of the key points discussed by the experts along the themes of modelling social argumentation, completeness, clarity and consistency (and relevant sub-themes).

5.1.1. Social Argumentation

Each of the experts agreed that there was value in modelling social argumentation, Expert F going so far as to say they believed there was no argument that didn't have social components. Expert D discussed how understanding the nuances of how people argue socially could lead to ways of helping or encouraging them to argue "better", in a more cooperative or polite manner.

The challenges of modelling social argumentation the experts foresaw were mostly a question of scale. In part, the sheer volume of data in a social media discussion can be

Table 1. Summary of experts' opinions on key aspects of ASWO

Theme	Sub-theme	Comments
Social Argumentation	<i>Value</i>	<p>"...if we're going to have a realistic model of how people argue, we've got to look at how people really argue rather than how our "ideal reasoner" would argue" – Expert A</p> <p>"I think modelling social argumentation is very important...I want to say it's useful in trying to help people argue 'better'." – Expert D</p>
	<i>Challenges</i>	<p>"Even in quite a simple back-and-forth argument, there's quite a lot going on...scale is a challenge" – Expert C</p> <p>"...enthymemes, humour, there's lots of missing information, there's lots of playing to particular audiences...there are lots of things that are current events or would only make sense to a particular group" – Expert B</p>
	<i>Abuse/Threats</i>	<p>"I, personally, tend to ignore all of those because I'm...focusing on the informal proof structures" – Expert A</p> <p>"...it's hard to exclude them...if you think about what you're going to do with the model...do you want to retrieve threatening and abusive comments? Well you might want to exclude them from being retrieved, which also makes it relevant to model that" – Expert B</p>
Completeness	<i>Implicit/Explicit Premises</i>	<p>"I think when people model arguments it's pretty common to infer the reading, and what's interesting is that there can be multiple readings. So it wouldn't be wrong to...put in some interpretation, as long as it's clear it's an interpretation and there can be others. " – Expert B</p>
	<i>Social Meta-Data</i>	<p>"One other thing... is other people's opinions of statements. A lot of systems have thumbs up and thumbs down...what you need is, I think, an audience response" – Expert F</p>
Clarity	<i>Generalisation</i>	<p>"If anything I think maybe your default conflict is a superclass - everything is a conflict, and one of the subclasses is a...rational argument. But then you've also got personal attack, ad hominem...these are all alternatives to rational argument, but at the default it might be worth allowing modelling of a conflict. Not a conflict as it is in the original model, but as a superclass of interaction." – Expert F</p>
Consistency	<i>Internal consistency</i>	<p>"whenever you try to model anything in a formalised system...if you give two people the same thing...unless it's something really simple, they will always find two different ways of modelling it" – Expert E</p> <p>"...rather than having the minimal number of nodes and encouraging people to just misuse them, I would rather say 'Here's a definite type of argumentation we want to capture and share...'" – Expert A</p>
	<i>External consistency</i>	<p>"Consistent with [the AIF], maybe not, but building on? Definitely" – Expert C</p>

overwhelming, particularly when considering the speed with which it can grow, but also in terms of the variety of information, which is often contextual, such as references to current events, or cultural “in-jokes”.

Experts A and D explained that they would not consider abusive argumentation as a valid when modelling an argumentation structure (as they focused broadly on dialectic arguments and that was the current standard for their domain), although they agreed it was a potentially valuable area to explore. Expert B explained that it depended very much on the purpose of the model — in some cases it may be important to model threatening and abusive attacks specifically so they can be excluded when presenting the model to users. Expert E also noted that excluding this type of argument can lead to confusion if a particular abusive comment changes the course of the argument, or causes the quality of the rest of the discussion to degenerate.

5.1.2. Completeness

Experts A and B both made explicit mention of the ability to mark certain posts as being in direct response to other participants in the discussion as a useful addition to argumentation frameworks.

Expert B noted that as many annotations have the potential to be subjective, it would be possible to extend this to include further subjective annotations such as an analyst’s confidence in a particular reading of an inference. Expert C had similar views and discussed including mappings of a participant’s agreement or disagreement with key positions in the dialogue as well.

Expert F discussed the potential for an “activity” score for each locution, derived from the social meta-data of each post (e.g. number of replies, number of up- or down-votes or number of retweets); this metric could be derived on a per-purpose basis to allow analysts to correctly categorise different platforms for their own needs, and to highlight key areas of the discussion that had solicited or stimulated large amounts of discussion.

Broadly, all experts agreed that to adequately model social argument that it was necessary to include further context about the participants, such as demographic information where available, such as by linking the SIOC UserAccount to a FOAF Agent, or additional information about key events related to the discussion to maintain relevance of the model for future analysis, and to limit the number of assumptions needed to be made by analysts.

5.1.3. Clarity

Expert D was concerned that, when faced as an analyst with a statement that appeared ambiguous (for example, a statement of support that could be interpreted as genuine or sarcastic) they may struggle to accurately and objectively model it, and proposed a means of allowing analysts to mark such relations as existing without committing to associating them with either a support or an attack.

Expert F proposed a similar solution, by means of generalising the model to include super-classes of Support and Conflict. “Personal” conflict, for example, is perhaps too specific a name for all non-logical conflicts: there are rhetorical attacks that can target institutions or accounts run by software, but also, importantly, positions and information. These Support and Conflict super-classes would encompass Logical Support/Conflict and Rhetorical Support/Conflict and could then be further sub-classed to provide more specific instances of each, where apparent, allowing analysts to defer when unsure.

5.1.4. Consistency

The majority of experts felt that these additions to the ASWO were consistent with the nodes used in the AIF. However, Experts C and F disagreed, pointing to the fact that the ASWO was intentionally inconsistent with the AIF because they were developed for different purposes.

In terms of inter-rater reliability — whether two analysts attempting to model the same argument would reach the same result — the experts were much more divided. While they agreed that the objective parts of the model (i.e. the locutions, user account and, in most circumstances, the persona) could be modelled identically (and in most cases, automated), Experts C and B felt that both analysts would reach the same conclusion overall with minor deviations, whereas Experts A, D and E disagreed, stating there was too much subjective information to model identically. Expert A felt that the analyst would naturally perceive the argument through their own lens of cultural and social context and Expert D noted the different levels of detail an analyst may choose to use, whether focusing only on premises that have been explicitly stated, or including additional implicit information.

How important this is was also a matter of some debate: Experts B and C felt that it was likely there would (and should) be one “correct” representation of an argument. Experts D and F agreed to an extent, citing their proposals for handling ambiguous content being able to aid annotators in this regard, so that if the model could not be complete, it could be consistent. Expert A felt that ideally analysts should reach the same conclusion but in practice, the subjective nature of the task might make this impossible. Expert E felt the consistency of annotators would, in practice, be less important and would be a factor of the intended purpose of the model.

6. Conclusions

In this paper we provide further extensions to the ASWO to incorporate other modes of rhetorical persuasion that contrast with logical argument. We conducted an expert review that highlighted some key strengths of this model, such as the ability to model directed replies, the ability to model the audience and the ability to model instances of irrational and eristic argument that were previously difficult or impossible to achieve with the AIF alone.

This review highlights some current limitations of the ASWO framework as it stands that will need to be addressed to further improve the model. Firstly, the issue of scalability: annotating web-based argumentation in this manner remains a high-cost affair in terms of knowledge and time. Future work will examine how suitable a crowd-sourced annotation approach is, with respect to accuracy and inter-rater reliability. Secondly, automation: because social argumentation can rely heavily on nuanced contextual information (such as the ability to recognise humour, sarcasm or references to current events) it is likely impossible to model it in such a way that it could be automatically reasoned over. However, because the ASWO provides additional information about rhetorical tactics in use, it provides human analysts the means to explore the resulting structure in greater detail and context. This can also potentially be used to highlight areas of particular interest, or assist in community decision-making environments.

The review also highlights useful directions of further work, such as including further contextual information such as participant demographics or social meta-data, or generalising the ontology further. It also lays groundwork for an investigation into how such rhetorical structures are used on different social web platforms. By using this extended framework, we aim to determine if and how the perceived contribution and value of a comment correlates to the dialectic and eristic content.

Our hope is that these developments lead to a means of more accurately modelling social argumentation which in turn provides a path to creating tools to allow social media users to critically analyse discussions in progress and to encourage them to engage with debates in good faith.

Acknowledgements

This work is funded by a scholarship provided by the UK Engineering and Physical Sciences Research Council. Thanks also go to the reviewers for their detailed, constructive and helpful reviews.

References

- [1] J. Schneider, T. Groza, and A. Passant, "A review of argumentation for the Social Semantic Web," *Semantic Web*, vol. 4, no. 2, pp. 159–218, 2013.
- [2] E. A. Jane, "'Your a Ugly, Whorish, Slut' Understanding E-bile," *Feminist Media Studies*, vol. 14, no. 4, pp. 531–546, 2014.
- [3] T. Blount, D. E. Millard, and M. J. Weal, "Towards Modelling Dialectic and Eristic Argumentation on the Social Web," in *14th workshop on Computational Models of Natural Argument*, 2014.
- [4] T. Blount, D. E. Millard, and M. J. Weal, "An Investigation into the Use of Logical and Rhetorical Tactics within Eristic Argumentation on the Social Web," in *ACM Conference on Hypertext and Social Media*, 2015.
- [5] Plato, *Book V. The Republic*, Basic Books, 380BC. (Bloom, A.D. Trans. 1991).
- [6] D. Walton and E. C. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, NY: State University of New York Press, 1995.
- [7] D. N. Walton, *The New Dialectic: Conversational Contexts of Argument*. University of Toronto Press, 1998.
- [8] G. B. Kerferd, *The Sophistic Movement*. Cambridge University Press, 1981.
- [9] J. Schneider, S. Villata, and E. Cabrio, "Why did they post that argument? Communicative Intentions of Web 2.0 Arguments," in *Arguing on the Web 2.0*, (Amsterdam), SINTELNET, European Network for Social Intelligence, 2014.
- [10] C. Jørgensen, "Public Debate – An Act of Hostility?," *Argumentation*, vol. 12, no. 4, pp. 431–443, 1998.
- [11] J. Schneider, B. Davis, and A. Wyner, "Dimensions of argumentation in social media," *Lecture Notes in Computer Science*, vol. 7603, pp. 21–25, 2012.
- [12] C. Chesñevar, J. McGinnis, S. Modgil, I. Rahwan, C. Reed, G. Simari, M. South, G. Vreeswijk, and S. Willmott, "Towards an argument interchange format," *Knowledge Engineering Review*, vol. 21, no. 4, pp. 293–316, 2006.
- [13] C. Reed, S. Wells, J. Devereux, and G. Rowe, "AIF+: Dialogue in the Argument Interchange Format.," *FRONTIERS IN ARTIFICIAL INTELLIGENCE AND APPLICATIONS*, vol. 172, p. 311, 2008.
- [14] M. J. J. L. C. REED, "Ova+: An argument analysis interface," in *Computational Models of Argument: Proceedings of COMMA*, vol. 266, p. 463, 2014.
- [15] J. G. Breslin, S. Decker, A. Harth, and U. Bojars, "SIOC: An Approach to Connect Web-Based Communities," *International Journal of Web Based Communities*, vol. 2, no. 2, pp. 133–142, 2006.