

A Framework for Enriched, Controlled On-line Discussion Forums for e-Government Policy-making*

Adam Wyner¹ and Tom van Engers²

¹University College London, Department of Computer Science
adam@wyner.info

²University of Amsterdam, Leibniz Center for Law
vanEngers@uva.nl

March 29, 2010

Abstract

The paper motivates and proposes a framework for enriched on-line discussion forums for e-government policy-making, where pro and con statements for positions are structured, recorded, represented, and evaluated. The framework builds on current technologies for multi-threaded discussion lists by integrating modes, natural language processing, ontologies, and formal argumentation frameworks. With modes other than the standard “reply comment”, users specify the semantic relationship between a new statement and the previous statement; the result is an argument graph. Natural language processing with a controlled language constrains the domain of discourse, eliminates ambiguity and unclarity, allows a logical representation of statements, and facilitates information extraction. However, the controlled language is highly expressive and natural. Ontologies represent the knowledge of the domain. Argumentation frameworks evaluate the argument graph and generate sets of consistent statements. The output of the system is a rich and articulated representation of a set of policy statements which supports queries, information extraction, and inference.

Keywords: eParticipation, policy-making, controlled language, ontology, argumentation

1 Introduction

eGovernment is mostly aimed at providing services that make it easier for citizens to comply with the law and to execute their rights as well as for public administrations to effectively and efficiently enforce the law. Effective law enforcement requires that the laws and regulations are supported by a sufficient

*Under review at eGOV 2010. Corresponding author: Adam Wyner, adam@wyner.info.

number of addressees, since lack of support usually leads to non-compliant behaviour, which can only be countered at great cost or with repressive actions. In the European Union, it is widely recognised that it is important to promote and reinforce democratic institutions and build support for regulations by, in part, broadening participation and making policy efficiently. As the application of policy depends on compliance, consultation and dialogue about policy would encourage natural compliance as the resulting policy would more accurately relate to and bear on the people which the policy affects. The result would be better, more effective, and less expensive enforcement of policy.

One means to broaden participation and make policy efficiently is to develop online forums for consultations and participatory discussions ([10] and [2]). However, online forums generate a wealth of information which needs to be structured, represented, extracted, reasoned with, and analysed. Moreover, as participants use natural language to express and understand policy, the tools must process natural language to some extent in order to keep the discussion well-structured, to provide fine-grained information about the contents, and to allow information extraction. In other words, current forums contain a great deal of information expressed in natural language, and yet they are relatively uninformative in the sense that it is difficult to keep the discussion on topic and to glean much useful information from them without detailed manual reading, review, reformulation, and summarisation. This explains why forums are not used to co-create policies and regulations to the extent they could be. To support these tasks mentioned above, a range of emerging technologies can be applied and integrated such as multi-threaded discussion forums with controlled languages, ontologies, and argumentation frameworks.

In this paper, we propose and outline a framework which extends multi-threaded discussion forums and integrates these technologies. Our approach is incremental and systematic in that we address some aspects of argumentation and natural language processing, leaving other aspects to future development. We first give some background to the proposal and a small sample policy making discussion, indicating issues to be resolved. We then outline some of the key components which provide a generic structure for a policy making support tool, schematically illustrating their relationships. Finally, we outline future work and issues that must be address.

The objective of the tool is scoped as our primary objective is to provide a better basis for policy-making by providing a well structured, coherent, and analysable set of statements that underlie the policy. There are a range of additional elements of policy-making that we do not consider such as drafting rules, explanatory portions for motives and goals, measures of intended and unintended effects to determine the effectiveness of the policy, bridging strategies during implementation of the new policy, and so on.

2 Background

While in many cases governments already consult stakeholder representatives and experts in legislation drafting phase and policy making processes, the most used tools typically are face to face meetings and consultation/commenting rounds, thereby limiting the number of participants and missing the opportunity to leverage “the wisdom of the crowds” and building support while creating the

regulations.

The popularity of unthreaded or multi-threaded online forums (U/MT) such as those which appear on slashdot.com, the BBC's Have Your Say¹, and discussion forums on craigslist.org demonstrate that these systems can be used to exchange information, whether to comment on a blog or to exchange information on a discussion board. They are easy to install, maintain, and use. In U/MTs, users can respond to particular topics, building a flat or hierarchical tree of topics and comments. Standardly, users only need to indicate that they wish to respond to a topic in order to open a text box for text entry, which is then saved. A database records not only the input statements, but information about the relationships among statements.

While U/MTs clearly have advantages in terms of gathering and storing information, they are not in and of themselves highly informative, but require analysis to identify relationships between statements other than the "reply", extract information, draw inferences, or create argument maps ([19], [14], and [13]). For example, consider the BBC's Have Your Say¹, in which a problem or topic of discussion is presented on line for members of the public to respond to; thus it represents an informal example of a policy making discussion. Participants write their comments in. However, the semantic relationships between the statements is not, in many cases, made explicit, but must be inferred, where the semantic relationships might be agreement, disagreement, introduction of a premise or exception, refinement, pronominal anaphora, and others. While one comment may refer to an immediately preceding comment or use some index to a preceding comment, often this is not the case. Consequently, it is difficult to reconstruct the overall argument, the relationships among statements, and the conclusions that might follow from the give set of statements. Indeed, the longer the list of comments, the more problematic the task. In addition, it is difficult to prevent or filter out redundancy, to require grammatically well-formed statements, or to encourage the discussion to be bound within a given domain of discourse.

To indicate such semantic relationships, online systems have been developed for debate such as Debategraph², Debatepedia³ and for argument representation such as Araucaria⁴ [18], Carneades⁵ [6], and ArguMed⁶ [20]. These systems have a range of rich resources. Debatepedia allows columns of statements to be made pro or con a given statement. Debategraph is a dynamic system that supports incremental development of argument structure; however, it is proprietary and has no argument theory to draw general conclusions. With Araucaria, one takes a given fragment of text and annotates it with an argument property (e.g. premise or conclusion) and indicates the relationship to other statements, resulting in a graph of statements in relationships. Carneades has an argumentation theory guided by heuristics rather than an underlying logic. ArguMed provides a formal theory and implementation for evaluating arguments, but is not available for online use with multiple participants. In all systems, the participant is responsible for indicating the semantic relationship between one statement and

¹http://news.bbc.co.uk/1/hi/talking_point/default.stm

²<http://debategraph.org/>

³http://debatepedia.idebate.org/en/index.php/Welcome_to_Debatepedia!

⁴<http://araucaria.computing.dundee.ac.uk/doku.php>

⁵<http://carneades.berlios.de/>

⁶<http://www.ai.rug.nl/~verheij/aaa/argumed1.htm>

the next, e.g. whether a statement is a premise, an exception, a conclusion, or a contradiction with respect to some other statement. Online systems are not underwritten by an argumentation logic. And finally, no system supports natural language processing with respect to a controlled language or ontology.

Online U/MTs and debate representation systems have limitations. First, they all rely on the participant to have a clear, intuitive understanding of the relationships between statements. Debategraph and Debatepedia do not have a clear semantics associated with these relationships or a formally specified semantics of argumentation. Araucaria has an intuitive semantics, but it is not given formally. Second, no system allows context dependent relationships where, for example, one and the same statement may serve as the conclusion of one argument and the premise of another. Third, no system supports a modular architecture, where different relationships or debate components may be added systematically. Fourth, for all systems, the linguistic content of the statements is unanalysed and unconstrained; that is, the statements are not parsed, or given a semantic interpretation, or required to be relevant and novel to the current discussion, or constrained in terms of terminology and length. Finally, the systems do not have a formally specified semantics of argumentation such that we could determine sets of consistent statements (positions on the policy) and the supporting arguments for conclusion ([1] and [16]).

We want a context dependent, flexible, modular system to represent different notions of semantic relationships among statements, where the sentences are linguistically analysed and the debate is formally evaluated. In our proposed system, we address these limitations. Before we turn to our proposal, we introduce an example to ground the discussion.

3 Example

The example is derived from a public discussion list concerning recycling. It is an adaptation of the BBC's Have Your Say discussion of the question Should people be paid to recycle?⁷ While we might like to consider unrestricted natural language, e.g. the full range of expressions found in the discussion, doing so introduces complexities that we are unable to address at this time. Rather, we take a piecemeal and incremental approach wherein we first develop a working fragment that we extend as needed. Therefore, we consider a restricted domain and normalised the language.

We follow approach to argument graphing as in [15], where ordinary legal disputes are formalised. Each statement is represented as a node, claims (conclusions) and premises are represented with continuous arrows between nodes, while contradictions or conflicts between statements are represented with dashed arrows. Each statement is made separately by an individual on the discussion list, and the order in which a statement appears (given by a number) in the discussion list may be different from the order below. Thus, discussions are by nature dynamic.

- (1) Every householder should pay tax for the garbage which the householder throws away.
- (2) No householder should pay tax for

⁷<http://newsforums.bbc.co.uk/nol/thread.jspa?forumID=7269&edition=2&ttl=20100218141845>

the garbage which the householder throws away. (3) Paying tax for garbage increases recycling. (4) Recycling more is good. (5) Paying tax for garbage is unfair. (6) Every householder should be charged equally. (7) Every householder who takes benefits does not recycle. (8) Every householder who does not take benefits pays for every householder who does take benefits. (9) Professor Resicke says that recycling reduces the need for new garbage dumps. (10) A reduction of the need for new garbage dumps is good. (11) Professor Resicke is not objective. (12) Professor Resicke owns a recycling company. (13) A person who owns a recycling company earns money from recycling. (14) Supermarkets create garbage. (15) Supermarkets should pay tax. (16) Supermarkets pass the taxes for the garbage to the consumer.

For example, an individual makes statement (1), another gives (4) as a reason or premise for (1), yet another makes (3) as an additional reason for (3), which can be understood to lend greater strength to the claim that (1) should hold. (9) supports the claim in (4). However, (4) is criticised by the claim that the Professor is not objective, so the implication one might draw from his statement does not hold. In (2), we have a counter-proposal with a range of supporting reasons; the counter-proposal can be understood as the negation of a claim to the previous argument in favour of taxing garbage. (16) attacks (15), which is one of the premises of the argument in favour of (2), so is the negation of a premise.

We would want users to enter only well-formed grammatical statements which can be semantically translated. The lexicon ought to be clear and from a well-defined domain. The semantic relationships among lexical terms could best be encoded in an ontology, which represents domain knowledge. The semantic relationships between statements (premise for a claim, contradiction of a statement, etc) should be recorded. Moreover, we would want to enable extraction of information about entities (e.g. people, organisations, objects), their properties (e.g. Resicke is a professor), their relations (e.g. ownership, payment), and anaphoric relationships (what elements co-refer in the discussion).

This “discussion” is not to be taken as a full, complete representation of all possible statements and counter-statements. Moreover, in some cases, there is an intuition that one statement attacks another statement – (16) attacks (15) – but much is left implicit. Argumentative discussion proceeds by such partial steps with missing premises which can be filled in later. Finally, it is our presumption that the system is used in “high value” contexts by participants who are willing and able to adapt to some of the constraints (topic, expressivity, explicit marking of statement relations) in order to gain the advantages (clarity and reasoning support).

4 The Proposal

In Section 2, we introduced four limitations of current debate systems. Arguably, the first – that the system must rely on the user to correctly specify the relationship between statements - can be addressed in large measure by what follows from the debate construction; that is, participants introduce statements




Figure 1: Forum Design

which support or undermine claims of semantic relationships among statements using explicit linguistic indicators or a menu. We focus, then, on the second, third, and fourth limitations. The proposal has four components: multi-modal multi-threaded discussion forums, a controlled language, an ontology (related to the controlled language), and an argumentation framework. We outline these components and then discuss them briefly in subsections below.

In Figure 1, we have a user and a discussion forum. The user reads the statements on the forum and identifies some part of the discussion which she wants to contribute to by adding a statement in relation to an existing statement. The user selects a mode to indicate the relationship between the existing statement and the statement to be input. A text box opens into which the user enters text. The input of text is guided by a controlled language (labelled ACE) which guides the user to input a well-formed statement given a restricted grammar and vocabulary. The vocabulary is linked to an underlying ontological representation. The controlled language parses and semantically represents the sentence. In addition, as the statement is associated with a mode, an argumentation graph is created, which relates one statement to the next; the graph is input to an argumentation framework (AF), which calculates sets of consistent statements out of the debate (which contains inconsistent statements). The resulting statements in their relationships constitute a policy from which inferences, answers to queries, and extracted information may be given.

4.1 Multi-modal Multi-Threaded Discussion Forums

In order to reuse existing, widely available technology, we build on open source multi-threaded discussion forums such as in [22] or on slashdot.com rather than redevelop existing debate systems. In current forums, there is a single “reply” mode, where the user only indicates that a current input statement is a reply to an existing statement. In contrast, we allow the introduction of a menu of alternative choices to specify the relationship between an existing statement and a new statement, which is a multi-modal multi-threaded discussion forum. For example, the menu would list “contradicts”, “is premise of”, or “is an exception to” relations. Thus, given an existing statement P and a user input statement Q, the user selects among the available relations; in particular, suppose the user wants to indicate that Q contradicts P, and so selects that relation from the menu. From such choices, one can construct an “argument graph”.

The proposal is similar to existing debate systems such as Arcauria or Carneades, but is based on a widely available open source system and is easily extended to directly make use of controlled languages, ontologies, and argumentation frameworks. The proposal would adapt discussion facilities such as on slashdot.com, which allow valuations of (not relationships between) statements, “weighting” of statements, and a range of participant roles. In our proposal, the menu of relationship indicators is context dependent. We adopt a modular framework where, for example, different menus of indicators can be installed to serve other policy-making purposes.

4.2 Controlled Language Input

In addition to modes, the input statements are parsed and semantically interpreted using Attempto Controlled English (ACE), which ensures that only grammatically well-formed statements are input, translated to first order logic representations, and linked by discourse referents ([5] and [9]).⁸ The user introduces a new statement into the discussion by opening a textbox; the textbox is linked to ACE; the result of parsing and semantic interpretation is stored in the database along with the information about the mode.

ACE is a controlled natural language, which means that it is a formal language with a restricted lexicon, restricted range of syntactic forms, and correlated semantic interpretations. Though the language is controlled, it is nonetheless highly expressive, and the sentences are read and understood as natural language, e.g. English text such as Every woman is a human or the following small discourse A man tries-on a new tie. If the tie pleases his wife then the man buys it. The vocabulary is extensive and extensible. The grammar allows the formation of complex noun phrases with relative clauses and quantifiers, possessives, pronouns, composite sentences with conjunctions, negation, conditionals, and subordination. Given a knowledge base comprised of ACE statements, questions can be addressed that return a truth-value or the list of entities which satisfy the question. ACE bars entry of ungrammatical or ambiguous sentences using a predictive editor, which guides the user on input. Moreover, ACE comes packaged with an inference engine and a redundancy checker. Thus, ACE serves as a knowledge representation and query language.

⁸<http://attempto.ifi.uzh.ch/site/>

While ACE as a controlled language is more limited in certain respects, it is nonetheless highly expressive and flexible, allowing users to enter natural and grammatical English sentences. Though using ACE in a discussion forum may not suit just any user with any sort of English, in the context of valuable policy discussions, the advantages strongly outweigh the disadvantages. Thus, with ACE, we address a range of linguistic issues about policy-making forums.

4.3 Ontologies

An ontology explicitly and formally defines the concepts and relations along with the restrictions among them that may exist in a given domain [7] and [8]). An ontology provides a common vocabulary and framework, allowing researchers to share, test, and modify the conceptualisation. Ontology frameworks have been developed to suit the representation of web-based content – the Ontology Web Language 2 (OWL 2) – which consists of classes, the subsumption relations between the classes, object properties that relate to instances of classes, and restrictions on what properties may hold of these instances. In a well-designed ontology, one can draw inferences about classes as well as the types of individuals (i.e. the classes to which they belong) and their properties. Given instances of the classes, we have a knowledge base.

ACE provides a plug-in to the ontology development tool in the Protege ontology editor and knowledge acquisition system.⁹ Thus, using the same terminology to control the input as in the ontology enables one to use ontology to reason about the domain.

4.4 Argumentation Frameworks

Given well-formed propositions provided by ACE, we have a database of propositions and their relations which can be used by further systems such as an argumentation framework such as [3], [16], or [1] as well as argument graphing systems such as Araucaria and Carneades.

From a representation of the relationships among the propositions given by the modes, we can define arguments, where an argument is essentially statements in support of a claim such as we provided in our example. From such a network of statements in relationships, we calculate consistent sets of propositions as in argumentation frameworks, which are formal graphs of nodes and arcs with respect to which one can define notions of attack and defeat. The key point is that consistent sets of propositions which are drawn out from a debate form a well-formed policy, at least in terms of the statements gathered at the time.

While there are a number of argumentation frameworks and graphing systems, none of them control and structure the natural language input so as to support the logical translation. Rather, the presumption is that the input to the argumentation framework are simply expressions of well-formed propositional or first-order logic. Our proposal is unique in providing a systematic means to input natural language, translate it into logic, and then input the correlated proposition into an argumentation framework to calculate consistent sets of propositions. Thus, we maintain a link between the language of the policy,

⁹<http://protege.stanford.edu/>

the logical representation of those statements, the logical representation of the arguments as a whole, and evaluation of the arguments.

5 Discussion and Conclusion

Despite the lack of structure, there have been efforts to apply information extraction techniques to the contents of blogs and their comments in order to determine relationships among postings. [19] overviews blog search engines and their search capabilities, which are limited to familiar search capabilities such as date, phrase control (e.g. exact match, inclusion or exclusion of words), and specified headers (information in the blog title, URL, or other indices). Some (e.g. Google Blog Search) have meme trackers which allow search for and organization of blogs according to concepts or terms that they share. [14] and [13] show how subjects of online debates can be successfully tracked in short documents according to co-occurring terms. However, performance falls in relation to the length of each document. [12] discuss sentiment analysis in online forums. [23] outlines the limited success in identifying argument structure (premises, rule, claim) from texts so long as overt indicators of argument structure are present (e.g. indicators such as supposing and therefore). Performance falls where argument structure must be inferred. So far as we know, there are no studies of forums or blogs which identify the point-counterpoint structure that one might find in a debate.

While our proposal is ambitious, it is practical in the sense that we extend and integrate a range of existing technologies, each serving a key role in the overall input to and evaluation of a policy. The prototype system would be tested for usability and expressivity in realistic user context. In addition to realising the core functionalities, we anticipate extending the system to address dialogue types, argumentation schemes, and critical questions [21], which further articulate the ways that participants introduce statements into the policy discussion. Discussion protocols may be used as one way to introduce context dependent, fine-grained structure to the discussion ([11] and [4]). The database of statements ought to be compatible with emerging Argument Interchange Formats [17]. Finally, we will investigate the extent to which natural language argument indicators can themselves be used to mark intrasentential relationships rather than a menu.

Acknowledgements

During the writing of this paper, the first and third authors were supported by the IMPACT Project (Integrated Method for Policy making using Argument modeling and Computer assisted Text analysis) FP7 Grant Agreement No. 247228.

References

- [1] Philippe Besnard and Anthony Hunter. *Elements of Argumentation*. MIT Press, 2008.

- [2] Dan Cartwright and Katie Atkinson. Using computational argumentation to support e-participation. *IEEE Intelligent Systems*, 24(5):42–52, 2009.
- [3] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.
- [4] Raquel Fernández and Ulle Endriss. Abstract models for dialogue protocols. *Journal of Logic, Language and Information*, 16(2):121–140, 2007.
- [5] Norbert E. Fuchs, Kaarel Kaljurand, and Tobias Kuhn. Attempto controlled english for knowledge representation. In Cristina Baroglio, Piero A. Bonatti, Jan Maluszynski, Massimo Marchiori, Axel Polleres, and Sebastian Schaffert, editors, *Reasoning Web*, volume 5224 of *Lecture Notes in Computer Science*, pages 104–124. Springer, 2008.
- [6] Thomas Gordon, Henry Prakken, and Douglas Walton. The carneades model of argument and burden of proof. *Artificial Intelligence*, 171:875–896, 2007.
- [7] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [8] Rinke Hoekstra. *Ontology Representation – Design Patterns and Ontologies that Make Sense*, volume 197 of *Frontiers of Artificial Intelligence and Applications*. IOS Press, Amsterdam, June 2009.
- [9] Tobias Kuhn. *Controlled English for Knowledge Representation*. PhD thesis, Faculty of Economics, Business Administration and Information Technology of the University of Zurich, 2010.
- [10] Ann Macintosh. Moving toward "intelligent" policy development. *IEEE Intelligent Systems*, 24(5):79–82, 2009.
- [11] Peter McBurney, David Hitchcock, and Simon Parsons. The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems*, 22(1):95–132, 2007.
- [12] Rudy Prabowo and Michael Thelwall. Sentiment analysis: A combined approach. *Journal of Informetrics*, 3(1):143–157, 2009.
- [13] Rudy Prabowo, Michael Thelwall, Iina Hellsten, and Andrea Scharnhorst. Evolving debate in online communication: A graph analytical approach. *Internet Research*, 18(5):520–540, 2008.
- [14] Rudy Prabowo and Mike Thelwall. Finding and tracking subjects within an ongoing debate. *Journal of Informetrics*, 2(2):107–127, 2008.
- [15] Henry Prakken. Formalising ordinary legal disputes: a case study. *Artificial Intelligence and Law*, 16:333–359, 2008.
- [16] Henry Prakken. An abstract framework for argumentation with structure arguments. *Argument and Computation*, ??-?, 2010. To appear.

- [17] Chris Reed Rahwan, Iyad and Fouad Zablith. On building argumentation schemes using the argument interchange format. In *Working Notes of the 7th Workshop on Computational Models of Natural Argument (CMNA 2007)*, Hyderabad, 2007.
- [18] Chris Reed and Glenn Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools*, 13(4):961–980, 2004.
- [19] Michael Thelwall and Laura Hasler. Blog search engines. *Online Information Review*, 31(4):467–479, 2007.
- [20] Bart Verheij. Argumed - a template-based argument mediation system for lawyers. legal knowledge based systems. In J.C.Hage, T.J.M. Bench-Capon, A.W. Koers, C.N.J. de Vey Mestdagh, and C.A.F.M. Grtters, editors, *JU-RIX: The Eleventh Conference*, pages 113–130, Nijmegen, 1998. Gerard Noodt Instituut.
- [21] Douglas Walton. *Argumentation Schemes for Presumptive Reasoning*. Erlbaum, Mahwah, N.J., 1996.
- [22] Luke Welling. *PHP and MySQL web development*. Pearson Education Inc., 2009.
- [23] Adam Wyner, Raquel Mochales-Palau, Marie-Francine Moens, and David Milward. Approaches to text mining arguments from legal cases. In Simonetta Montemagni, Daniela Tiscornia, and Enrico Francesconi, editors, *Semantic Processing of Legal Texts*, Lecture Notes in Computer Science. Springer, 2009. To Appear.