Towards Web-based Mass Argumentation in Natural Language

Adam Wyner[1] and Tom van Engers[2]

[1] Department of Computer Science
University College London
Gower Street, London WC1E 6BT, United Kingdom
adam@wyner.info
[2] Leibniz Center for Law
Oost-Indisch Huis, Kloveniersburgwal 48, Room E 1.09A
PO Box 1030, 1000 BA Amsterdam, The Netherlands
VanEngers, K.Bahreini@uva.nl

Abstract. Within the artificial intelligence community, argumentation has been studied for quite some years now. Despite progress, the field has not yet succeeded in creating support tools that members of the public could use to contribute their views to discussions of public policy. One important reason for that is that the input statements of participants in policy-making discussions are put forward in natural language, while translating the statements into the formal models used by argumentation scientists is cumbersome. These formal models can be used to automatically reason with, query, or transmit domain knowledge using web-based technologies. Making this knowledge explicit, formal, and expressed in a language which a machine can process is a labour, time, and knowledge intensive task. To make such translation and it requires expertise that most participants in policy-making debates do not have. In this paper we describe an approach with which we aim at contributing to a solution of this knowledge acquisition bottle-neck. We propose a novel, integrated methodology and framework which adopts and adapts existing technologies. We use semantic wikis which support mass, collaborative, distributive, dynamic knowledge acquisition. In particular, ACEWiki incorporates NLP tools, enabling linguistically competent users to enter their knowledge in natural language, while yielding a logical form that is suitable for automated processing. In the paper we will explain how we can extend the ACEWiki and augment it with argumentation tools which elicit knowledge from users, making implicit information explicit, and generate subsets of consistent knowledge bases from inconsistent knowledge bases. To a set of consistent propositions, we can apply automated reasoners, allowing users to draw inferences and make queries. The methodology and framework take a fragmentary, incremental development approach to knowledge acquisition in complex domains.

1 Introduction

In research on formal argumentation, systems have been developed to reason with inconsistent, defeasible knowledge bases ([1], [2], [3] and [4]). However, the examples used to demonstrate the systems are simplified or abstract. Translating arguments which are expressed in natural language into formal argumentation frameworks is highly labour and knowledge intensive, drawing on implicit knowledge. Moreover, tools have not been developed to support users while debating some issue. Such tools could find widespread application in important domains. In this paper, we address this issue by proposing a natural language interface tool to formal argumentation systems. Such a tool could help us to overcome the knowledge acquisition bottleneck and allow laypersons who have no expertise in knowledge engineering to contribute to the argument using AI-based argumentation tools. The tools would point out inconsistencies, elicit from the user information about implicit contradictions that rely on background information, and support the user in the construction of complex arguments.

One such application is to support policy-making. In the European Union, the promotion and reinforcement of democratic institutions is widely recognised to be important. Consultations and dialogues with members of the public concerning public policy are a key means to encourage *natural compliance* so that the policy relates more closely to the people the policy affects. Typically, the consultations involve face-to-face meetings and written position reports. However, these tools limit the number of participants, do not leverage “the wisdom of the crowds”, and are not structured representations of knowledge to which further analytic processes can be applied.

To broaden participation and make policy more efficiently, online forums have been used ([5] and [6]); however, these are limited in their functionality and representation. In particular, some online forums generate a wealth of information which needs to be structured, represented, extracted, reasoned with, and analysed (see the BBC’s "Have Your Say" and similar, UK Parliamentary debates, and online graphical/debate tools such as Compendium or DebateGraph which allow users to add statements one or more at a time). At this level, argumentation is interactive, dialectical, dynamic, and explicit (at least in the sense that we can take the information as is as well as subject it to query and subsequent refinement by others). Other systems do not allow users to contribute to the structure of the consultation. In any case, 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; much of the users’ knowledge is implicit, informal, and unstructured. This explains in part why forums are not used to co-create policies and regulations to the extent they could be.

In addition to issues related to natural language processing and knowledge engineering, a policy-making support system must explicitly deal with *disagreement* and *divergent* information. The participants in a forum can (and typically do) contribute statements for and against previous statements in the forum – a *dialectical discussion*. While the goal of a policy-making discussion might well be to arrive at a consistent statement of policy, the dialectical discussion will give rise to inconsistent knowledge bases, which cannot then be further reasoned with.

To improve public policy-making, our goal is to improve tools for knowledge acquisition and maintenance. The knowledge that people have of a domain is often implicit, informal, and expressed in natural language (the source form). To automatically reason with, query, or transmit using web-based technologies, the domain knowledge must be made explicit, formal, and expressed in a language which a machine can process (the target form). However, translation of domain knowledge from the source form to the target form is a labour, time, and knowledge intensive task [7], creating a "knowledge acquisition bottleneck" which has limited the adoption and use of powerful AI-technologies [8]. In addition, as we pointed out above, domain knowledge may be uncertain, inconsistent, distributed in a community, evolve over time.

To make online forums and debate tools more useful for policy-making, [9] and [10] propose and outline a framework which extends multi-threaded discussion forums, integrating NLP, ontologies, and argumentation. In contrast to existing debate and argumentation support systems, this support tool for policy-making makes the semantic *content* of comments formal and explicit as well as to makes formal and explicit the range of fine-grained *relationships between the statements* which are relevant for the construction of dialectical arguments.

The novelty of this paper is that we propose and outline a framework to take advantage of and extend an existing system for a semantic wiki which takes sentences of natural language as input, parses and semantically interprets them, and generates an ontology. The sentences can be edited and augmented by other users. To this system, we propose to add features to support knowledge acquisition for argumentation – the elicitation of contradiction and implicit premises.

Our approach is incremental, systematic, and scoped. We address some aspects of argumentation and natural language processing, leaving other aspects to future development. In particular, we work with a *controlled natural language* (CNL), described in Section 5, which has an extensive vocabulary and an expressive, yet normalised, syntax; the sentences written in this language appear as natural English sentences. They can be automatically parsed and (with some limits) translated into first order logic, giving rise to a knowledge base. The language can be used to express an ontology and rules. As we intend to deploy the system in "high value", the advantages of using a controlled language outweigh the disadvantages (See [11] for a related exercise and [12] for useability studies). In addition, as the objective of the tool is to provide a better basis for policy-making by providing a well structured, coherent, and analysable set of statements that underlie the policy, we do not consider a range of policy-making issues such as drafting rules, explanatory text, measures of the effectiveness of the policy, and so on.

The structure of the paper is as follows. In Section 2, we discuss some of the relevant background to our proposal. In Section 3, we ground our presentation with a sample policy-making. In Section 4, we outline the modules of our policy-making support tool, schematically illustrating their relationships and setting the context for the argumentation and NLP. The NLP functionality is reviewed in Section 5 and exercised on the example sentences. We discuss formal approaches to argumentation in Section 6, with the question being how such systems would handle arguments from natural language. In Section 7, we review our main points and questions, proposing a framework to integrate natural language into formal argumentation frameworks.

2 Background

In addition to standard online forums[3], there are tools for debates such as Debategraph[4] and for argument representation such as Araucaria [13]. These systems have a range of rich resources. 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.

In existing systems, the participant is responsible for indicating the semantic relationship between one statement and the next, e.g. whether a statement is a premise, an exception, a conclusion, or a contradiction with respect to some other statement. The linguistic content of the statements is unanalysed; that is, the statements are not parsed or given a semantic interpretation. The terminology and syntactic forms are unconstrained. 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 ([3] and [4]).

To scope the problems and make progress on a debate tool for public-policy making in natural langauge, we want a system which represents the relevant semantic relationships among statements, linguistically analyses and constrains the statements so as to provide rich, well-formed input to an argumentation system. In our proposed system, we address these limitations. Before we turn to our proposal, we introduce an example to ground the discussion.

3 Example

[9] provide an example, which is adapted from the BBC's Have Your Say online discussion of the question *Should people be paid to recycle?*[5] In making the adaptation, we have restricted terminology and normalised the language. This is realistic as we take a fragmentary and incremental approach, working a fragment which can be extended with terminology and syntactic structures. The presumption is that each statement is made separately by an individual and posted to a commonly available web-page; 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, non-linear (statements made earlier can be modified later), and incremental (adding statements over time). Our initial set of statements is:

> (1) Every householder should pay tax for the garbage which the householder throws away. (2) No householder should pay tax for the garbage which the householder throws away. (3) Paying tax for garbage increases recycling. (4)

[3] `www.slashdot.com`, among many others.
[4] `http://debategraph.org/`
[5] `http://newsforums.bbc.co.uk/nol/thread.jspa?forumID=7269&edition=2&ttl=20100218141845`

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) Tom says that recycling reduces the need for new garbage dumps. (10) A reduction of the need for new garbage dumps is good. (11) Tom is not objective. (12) Tom 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.

We follow an approach to argument graphing as in [14], where ordinary legal disputes are formalised (using the Arcauria tool). 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. However, rather than *reconstruction* of a argument *after* it has been made, we presume that users actively contribute to the *construction* of the argument as it is being made, different individuals taking turns to introduce a statement and its relationship to prior statements. For instance, one individual makes statement (1), while another gives (4) as a reason or premise for (1); another participant makes (3) as an additional reason for (3); and (15) is contradicted by (16), and so on (see [9] for a fuller discussion of the relations). In actual debates, the relations between these statements are often not explicit (see discussion of missing argument indicators [15]). These implicit relations must be made explicit. Nor do argument graphing approaches justify the selection of relations. For instance, even if contradiction were given between (15) and (16), there are background notions that ought to be elicited.

Among researchers in formal argumentation, it is common practice to construct the argument is to construct a *graph*, where the statements are nodes and the relationships arcs. In the example of such a graph in Figure 1, solid lines in the graph represent *support* relations (premises of a rule) while dashed lines represent *attack* relations (contradiction); rules themselves can be attacked, represented here as an attack on the arc between (4) and (9) (see [16] on reification of rules).

We aim at building a policy-making support tool. Such a tool should support dynamic, collaborative, evolutionary construction of the argument over time by multiple users. To bridge between formal argumentation systems and arguments in natural language, the tool constrains the users to input only well-formed grammatical statements using a limited terminology (the controlled language). The tool also elicits from users information about the semantic relationships between the statements. The "discussion" is a partial representation of what could be a more complicated discussion. Argumentative discussion proceeds by such partial steps with missing premises which can be filled in later. Finally, 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).

In the next section, we outline the framework, focusing the discussion on the syntactic and semantic analysis of the input sentences.

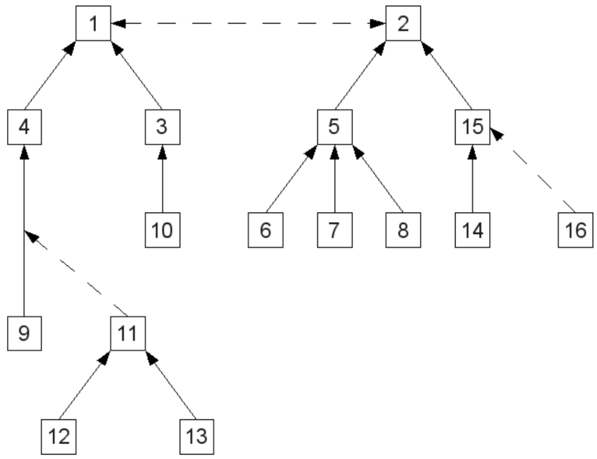

Fig. 1. Graphic Representation of Recycling Debate

4 Framework Outline

In this section, we outline the framework in [9] represented in Figure 2. We have a *user* and a *discussion forum*. For our purposes, the forum is *ACEWiki*, a semantic wiki which is described in Section 5. ACEWiki allows one to build knowledge bases in natural language, wherein the statements are translated into and out of First-order Logic. ACEWiki presents the user with statements in natural language. The participant reads the statements on the wiki, selects a statement to respond to, and chooses either to enter a new statement or to edit an existing statement. In either case, an editing window opens which provides a *predictive editor* which guides the user to input well-formed statements using the grammar and a given lexicon. After entering the statement, it is parsed and semantically interpreted as a statement of First-order Logic; in addition, a built in inference engine (Pellet) tests if the input sentence is consistent with the existing knowledge base. If it is, then it is entered into the knowledge base; otherwise, the syntactic form is recorded in the wiki, but with an indication that it is not part of the knowledge base.

Using ACEWiki, we can incrementally construct knowledge bases which correspond to portions of the graph in Figure 1. For instance, given (1), a user could edit it, adding (4) and (3) as premises; another user could add (10) as a premise of (3); yet another user could add (9) as a premise of (4). However, the system would not allow a user to enter (2) as a contradiction of (1) since it is inconsistent. Yet, it is fundamental to policy-making to allow such inconsistent knowledge bases to be made and reasoned with. For this reason, we add to ACEWiki functionalities associated with instantiated argumentation frameworks (AF) in Section 6 which can calculate sets of consistent statements out of an knowledge base comprised of inconsistent statements. This paper proposes a framework to overcome this limitation.

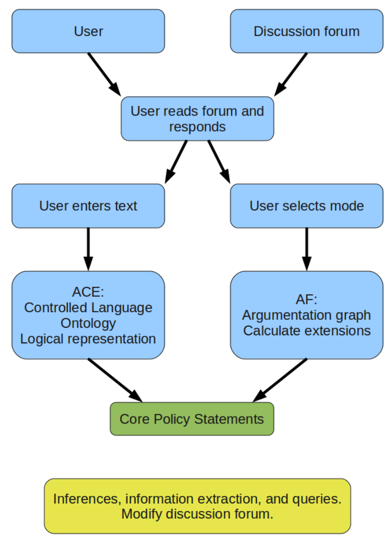

Fig. 2. Flow of Input

5 Attempto Controlled English

In this section, we outline aspects of *Attempto Controlled English* (ACE), ACEWiki, and how we have used it to translate the 16 example sentences into First-order Logic. 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 [12].[6] 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. ACEWiki uses ACE to provide a semantic wiki, which is a web-based wiki containing statements processed with ACE.

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

[6] http://attempto.ifi.uzh.ch/site/description/

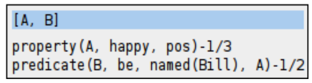

Fig. 3. Semantic Representation of *Bill is happy*.

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.

Consider a simple sentence such as "Bill is happy." A user enters in the sentence to an ACE interface on an ACEWiki; different representations can be requested such as a syntactic phrase structure tree or semantic representation. In the semantic representation, we use a *Discourse Representation Stucture* (DRS see [17] and [18]), which is a variant of first order logic and supports the semantic representation of aspects of discourse such as pronominal anaphora. In Figure 3, discourse referents (objects) A and B are introduced (some of the details of the representation are discussed further below). With respect to these objects, A is indicated to be the property *happy*, and B is the predicate *be* which predicates the property *happy* of an entity named *Bill*. Within a box, the statements are interpreted as conjuncts. A sentence such as "Every household creates some garbage." is interpreted as a conditional rule. As a first-order logical representation, a DRS can be used for reasoning.

ACE supports a large lexicon, a range of grammatical constructions, and correlated semantic interpretations: negation on nouns or verbs, conjunction, disjunction, conditionals, quantifiers, adjectives, relative clauses, discourse anaphora, modals ("necessity", "possibility", "permission", and "recommendation"), possessives, prepositional phrases, verbs with three arguments, and verbs with subordinate clauses.

[10] translate the 16 sentences into ACE, modifying the sentences to yield ACE compatible formats with the intended interpretations; agreeing with the broad results of [12], learning to use ACE accurately is not unduly burdensome. [10] discuss considerations that went into the revisions of the sentences as well as present and briefly discuss the semantic representations.

(1) Every household should pay some tax for the household's garbage.
(2) No household should pay some tax for the household's garbage.
(3) Every household which pays some tax for the household's garbage increases an amount of the household's garbage which the household recycles.
(4) If a household increases an amount of the household's garbage which the household recycles then the household benefits the household's society.
(5) If a household pays a tax for the household's garbage then the tax is unfair to the household.
(6) Every household should pay an equal portion of the sum of the tax for the household's garbage.
(7) No household which receives a benefit which is paid by a council recycles the household's garbage.

(8) Every household which does not receive a benefit which is paid by a council supports a household which receives a benefit which is paid by a council.
(9) Tom says that every household which recycles the household's garbage reduces a need of a new dump which is for the garbage.
(10) Every household which reduces a need of a new dump benefits the household's society.
(11) Tom is not objective.
(12) Tom owns a company that recycles some garbage.
(13) Every person who owns a company that recycles some garbage earns some money from the garbage which is recycled.
(14) Every supermarket creates some garbage.
(15) Every supermarket should pay a tax for the garbage that the supermarket creates.
(16) Every tax which is for some garbage which the supermarket creates is passed by the supermarket onto a household.

ACE translates each of these forms, draws inferences, and indicates inconsistencies where they arise. However, ACE can only reason with consistent sets of statements and has no functionality to support argumentative reasoning. Intuitively relating consistent knowledge bases of ACE and AF preferred extensions, we can say that the set of propositions derived from a preferred extension correlates to a consistent knowledge base of ACE (e.g. {a4, a1} correlates to the set of statements {(1), (3), (4), (9), (10), (16)}). Moreover, a consistent knowledge base of ACE correlates to an argument (e.g. {(1), (3), (4), (9), (10)} correlates to the argument a1). Finally, given a consistent knowledge base, if the addition of a statement induces contradiction, then the argument of which the statement is a part constitutes an attack on the consistent knowledge base. However, while these intuitions seem clear, formally specifying them and implementing them remains for future work.

6 Formal Argumentation

In this section, we introduce relevant aspects of Argumentation Frameworks (AFs). Abstract Argumentation Frameworks (AFs) ([1], [2], [19], among others) have been developed to explore. They are powerful in large part because arguments are atomic nodes and only one undifferentiated attack holds between them. Another line of research instantiates AFs with concrete examples ([3] and [4]). In such systems, a knowledge base is given which is made up of facts along with strict and defeasible inference rules in a Defeasible Logic (*DL*) ([20]). The facts and rules of the knowledge base intially appear as linguistic expressions, *If Bill is single, then he goes out on a Friday night*, which are then informally translated into a logical form, $P \rightarrow Q$, where P is *Bill is single* and *Bill goes out on a Friday night* is Q. The nodes of an abstract AF are interpreted as instantiations of arguments from the knowledge base.

However, all known approaches face the knowledge acquisition bottleneck – the arguments as expressed in natural language are manually translated into the knowledge base, the abstract nodes needed to calculate with an AF must be identified (either manually or automatically), and input into the AF. This is time-consuming, labour intensive, and requires the elicitation of implicit knowledge; these limiting the application of formal argumentation approaches. Moreover, without some formalisation of the translation method, the translation cannot be systematically validated. In the following, we briefly outline AFs and one approach to argument construction; the discussion highlights key problems to address.

For our purposes, we consider the AF of [1], where there is one set of undifferentiated objects, *arguments*, which are nodes in a graph as well as one undifferentiated relationship between the nodes, the *attack* relation, which can be represented as a graph in which attacks are arcs between nodes representing the arguments.

Definition 1. *An argumentation framework* AF *is a pair* $< \mathcal{X}, \mathcal{R} >$*, where* $\mathcal{X}$ *is a set of* objects*,* $\{a_1, a_2, \ldots, a_n\}$ *and* $\mathcal{R}$ *is an* attack *relation between objects. For* $\langle a_i, a_j \rangle \in \mathcal{R}$ *we say the the object* a_1 *attacks object* a_2*. We assume that no object attacks itself.*

Some of the relevant auxiliary definitions are as follows, where S is a subset of $\mathcal{X}$:

Definition 2. *We say that* $p \in \mathcal{X}$ *is* acceptable with respect to S *if for every* $q \in \mathcal{X}$ *that attacks* p *there is some* $r \in S$ *that attacks* q*. A subset,* S*, is* conflict-free *if no argument in* S *is attacked by any other argument in* S*. A conflict-free set* S *is* admissible *if every* $p \in S$ *is acceptable to* S*. A* preferred extension *is a maximal (w.r.t.* $\subseteq$*) admissible set. The object* $p \in \mathcal{X}$ *is* credulously accepted *if it is in at least one preferred extension, and* sceptically accepted *if it is in* every *preferred extension.*

Given the graph of statements in Figure 1, we abstract the "arguments" and represent the relationship as an AF. We assume, for the moment, that argument *a1* is comprised of statements {1, 3, 4, 9, 10}, *a2* of {11, 12, 13}, *a3* of {2, 5, 6, 7, 8, 14, 15}, and *a4* of {16}. Rebuttals, premise defeats, and undercutting are abstracted into the one abstract argument relation *attack*. The resulting graph is in Figure 4.

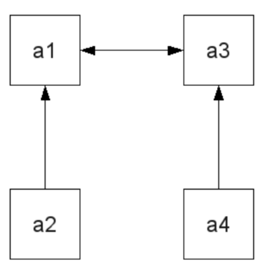

Fig. 4. Recycling Debate in an AF

In this AF, there are several preferred extensions depending on what is asserted to be true: if neither of a2 or a4 hold, then {a1} and {a3}; if a2 holds, but a4 not hold, then {a2, a3}; if a4 holds, but a2 does not, then {a4, a1}.

Given extensions of arguments, we would want to extract the sets of consistent statements which follow from a given set of assertions. For instance, given {a4, a1} and the assertions(9) and (10), then a consistent set of statements is {(1), (3), (4), (9), (10), (16)}; this would be a policy. While [3] and [4] (among others) formally specify the construction of abstract arguments from a knowledge base, they do not provide systematic means to construct arguments from natural language whether using explicit or implicit relationships among the statements that comprise the argument. Moreover, [3] discuss a range of computational problems and solutions to the construction of abstract arguments from a rich knowledge base; while we do not have the space here to elaborate, we assume the solutions of [3] can be applied to our knowledge base to construct arguments and extract consistent sets of propositions.

7 Discussion and Conclusion

To this point we have presented a means to provide a fine-grained analysis of natural language statements which can be used to argue about policy. Furthermore, we have presented a formal argumentation framework which can determine consistent sets of propositions from an inconsistent knowledge base. This brings us to the question of how we can further integrate natural language input with formal argumentation frameworks so as to encourage rich, complex argumentation.

As discussed in Section 5, using ACE can support the user to input grammatically well-formed and semantically interpreted sentences. Yet, this is not an altogether straightforward task, but requires a degree of specialist knowledge and understanding. In comparison to other knowledge acquisition tools (e.g. ontology editing tools such as Protege), ACE is easier to use. Yet further work needs to be done to extend the expressivity of ACE to cover additional constructions and interpretations, while checking that the resultant interpretations are as intended. One strategy might be to supply the user with questions derived from the sentences, where the idea is that this might indicate to the user that unintended interpretations appear and that the sentence needs revision.

Another global area to address is collaborative ontological construction. ACE has a tool, ACEView, which is integrated with Protege for building ontologies using natural language expressions. Where ontology construction is distributed and web-based as in ACEWiki, one must be concerned to maintain a common conception of the terms and relationships being introduced. In part, this is accomplished in ACEWiki by listing related available individuals, concepts, and statements. But, whether this is useful to a broader audience is an issue yet to be examined.

The language of ACE is limited to First-order Logic. This constraint on expressivity is problematic in policy discussions which often turn on normative concepts – how items are 'usually' understood, thought there are exceptions, or what parties are obligated to do, though they might violate the obligation. Such modal concepts are currently beyond the scope of ACE and, by the same token, reasoning with argumentation frameworks. One strategy might be to find equivalent first-order expressions, sacrificing brevity for expressivity. Nonetheless, this does not address reasoning with modal concepts.

In Section 6, we discussed issues about the relationship of formal argumentation frameworks and arguments in natural language. To provide a policy-making support tool, it will be essential to systematically translate from one to the other and back again.

The implicit knowledge which logicians use to translate from natural language into an argument such as the relationships of premise, claim, and contradiction, must be further empirically examined. An argumentation tool which encourages users to enter statements and annotate relationships among them would contribute to a database which would form the basis of such an empirical investigation. A separate issue is to identify a formalism which best supports the mapping, whether that of [3], [16], [4], or others.

Finally, in addition to feedback on the form and interpretation of individual sentences, the policy-making tool ought to provide concurrent support for participants to understand the implications, inconsistencies, or missing information from complex argument. This requires that natural language be input to the system, but that the formal argumentation framework also provide meaningful results which participants can then use to determine additional input. We know this is complicated and difficult, which is why we take a fragmentary and incremental approach to overcoming the knowledge acquisition bottleneck.

Acknowledgements

We thank the European Commission for supporting this research in the IMPACT Project (Integrated Method for Policy making using Argument modeling and Computer assisted Text analysis) FP7 Grant Agreement No. 247228.

References

1. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2) (1995) 321–358
2. Bench-Capon, T.J.M.: Persuasion in practical argument using value-based argumentation frameworks. Journal of Logic and Computation **13**(3) (2003) 429–448
3. Besnard, P., Hunter, A.: Elements of Argumentation. MIT Press (2008)
4. Prakken, H.: An abstract framework for argumentation with structure arguments. Argument and Computation **?** (2010) ?–? To appear.
5. Macintosh, A.: Moving toward "intelligent" policy development. IEEE Intelligent Systems **24**(5) (2009) 79–82
6. Cartwright, D., Atkinson, K.: Using computational argumentation to support e-participation. IEEE Intelligent Systems **24**(5) (2009) 42–52
7. van Engers, T.: Legal engineering: A structural approach to improving legal quality. In Macintosh, A., Ellis, R., Allen, T., eds.: Applications and Innovations in Intelligent Systems XIII, Proceedings of AI-2005, Springer (2005) 3–10
8. Forsythe, D.E., Buchanan, B.G.: Knowledge acquisition for expert systems: some pitfalls and suggestions. In: Readings in knowledge acquisition and learning: automating the construction and improvement of expert systems. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1993) 117–124
9. Wyner, A., van Engers, T.: A framework for enriched, controlled on-line discussion forums for e-government policy-making. In: Proceedings of eGov 2010, Lausanne, Switzerland (2010) Under Review.
10. Wyner, A., van Engers, T., Bahreini, K.: From policy-making statements to first-order logic. In: Proceedings of eGOVIS 2010. (2010) Under review.

11. Shiffman, R.N., Michel, G., Krauthammer, M., Fuchs, N.E., Kaljurand, K., Kuhn, T.: Writing clinical practice guidelines in controlled natural language. In Fuchs, N.E., ed.: Proceedings of the Workshop on Controlled Natural Language (CNL 2009). Lecture Notes in Computer Science, Springer (2010)
12. Kuhn, T.: Controlled English for Knowledge Representation. PhD thesis, Faculty of Economics, Business Administration and Information Technology of the University of Zurich (2010)
13. Reed, C., Rowe, G.: Araucaria: Software for argument analysis, diagramming and representation. International Journal on Artificial Intelligence Tools **13**(4) (2004) 961–980
14. Prakken, H.: Formalising ordinary legal disputes: a case study. Artificial Intelligence and Law **16** (2008) 333–359
15. Wyner, A., Mochales-Palau, R., Moens, M.F., Milward, D.: Approaches to text mining arguments from legal cases. In Montemagni, S., Tiscornia, D., Francesconi, E., eds.: Semantic Processing of Legal Texts. Lecture Notes in Computer Science. Springer (2009) To Appear.
16. Wyner, A., Bench-Capon, T., Dunne, P.: Instantiating knowledge bases in abstract argumentation frameworks. In Parsons, S., ed.: Proceedings of the AAAI Fall Symposium 2009: The Uses of Computational Argumentation, AAAI Press (2009) Forthcoming
17. Kamp, H., Reyle, U.: From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language: Formal Logic and Discourse Representation Theory. Springer (1993)
18. Asher, N.: Reference to Abstract Objects in Discourse. Kluwer Academic Publishers (1993)
19. Caminada, M., Amgoud, L.: On the evaluation of argumentation formalisms. Artificial Intelligence **171**(5-6) (2007) 286–310
20. Pollock, J.: Cognitive Carpentry: A Blueprint for How to Build a Person. MIT Press (1995)