CHRIS REED
Division of Applied Computing
University of Dundee, United Kingdom
chris@computing.dundee.ac.uk

Preliminary results from an argument corpus

Abstract. As reported in (Katzav et al., 2003), the University of Dundee has been developing a small corpus of examples of argumentation from a variety of domains (newspaper editorials, advertising, parliamentary records, judicial summaries, etc.) and a variety of regions (including India, Japan, South Africa, UK, Australia, US and others). This corpus has been analysed according to theories of argument structure (van Eemeren et al., 1996) as part of a project examining the role and structure of argumentation schemes – linguistic forms expressing stereotypical patterns of reasoning that form the 'glue' of interpersonal rationality. The corpus represents the first resource of its kind, and it is now being utilised by software systems in both teaching and research contexts. After explaining briefly the motivation and methodology adopted by the data collection and analysis work, this paper presents the first results of preliminary analyses of the corpus as a whole, and explores two distinct areas. The first is a straightforward investigation of surface features of the analysed arguments. Through such investigation, general differences between types of argument are identified. The second area is then a deeper exploration of scheme use, assessing links between scheme cladistics and their domain of use. This represents the first empirical assessment of real-world use of a complex set of argumentation schemes.

Introduction

Argumentation theory aims to better understand the way in which people argue, in situations of dialectical conflict, of dialogic co-operation, and of monological exposition (see (van Eemeren et al., 1996) for a textbook overview). It crosses traditional disciplinary boundaries in drawing upon linguistics, communication studies, psychology, rhetoric, law and philosophy. Increasingly its theories are also being adopted and extended in computer sciences including the computational theory, distributed computing, computational linguistics and artificial intelligence (Reed and Norman (2003) offer good examples of this interdisciplinary breadth).

In both theoretical and practical strands within the field, the topic of diagramming argument has been attracting increasing attention, as it both quickly uncovers interesting theoretical issues, and also forms a useful tool for students learning argumentation and critical thinking skills. Increasingly, software tools for supporting the task of diagramming are being deployed in both pedagogic and professional situations (Kirschner et al., 2003). A problem with many of these tools is their lack of argumentation theoretical input, which has meant that in some cases the approaches have been very ad hoc and therefore less appealing to the academic community. The *Araucaria* software for argument analysis and diagramming (Araucaria, 2004) tries to tackle this problem by tying recent theoretical advances to the software development process. The result is now in use in schools, universities, law practices and judiciaries around the world, but is also of use in academic work (Reed and Rowe, 2005). One example of the theoretical facet of *Araucaria* is its handling of argumentation schemes.

Argumentation schemes are becoming increasingly prominent in both argumentation theory and its applications in artificial intelligence. Schemes represent stereotypical forms of reasoning that though practically useful and frequently employed are nonetheless non-deductive and invalid on traditional grounds. Recent research has been trying to better understand, identify, classify and evaluate these schemes (Kienpointner, 1992; Walton, 1996; Katzav and Reed, 2004a). *Araucaria* supports argument analysis involving schemes, and saves resulting analyses in an open interchange format, the Argument Markup Language (AML). With a simple way of performing analysis, and storing the results for subsequent recall, manipulation and exchange, *Araucaria* offers an opportunity to build a resource of textual arguments and their analyses. Such a resource has applications in both teaching (where classroom exercises can be based upon a wide range of real world – rather than textbook – arguments) and research (where the time-consuming process of collecting examples is a tedious and expensive business) (Katzav et al., 2003).

The corpus construction process was conducted using a simple methodology, whereby two dozen or so online (and therefore semi-permanently accessible) resources were accessed on a regular basis and the first argument encountered at each site was stored and analysed. The sources were categorised by geographical region (Australia, India, Japan, South Africa, UK, US), and by broad domain (Cause Information, Discussion Forum, Legal, Magazine, Newspaper, Parliamentary Record). The corpus itself is freely available for both access and update (Araucaria, 2004), so in this paper we restrict investigation to those analyses conducted in 2003.

At around 150 extracts (and ca. 300 argument scheme instantiations), the 2003 corpus is probably large enough to support some limited statistical analysis. The aim here, however, is not to pre-empt deep, rigorous exploration of a much enlarged corpus, but rather to offer preliminary analyses in the way of observations and trends supported and suggested by the raw data. The aims of such investigation are

(i) to demonstrate that a corpus of analysed argument can indeed support interesting observations about argument usage in domains of discourse and cultural communities

(ii) to identify a set of issues in argument usage that can form a focus of future study

(iii) to lay a foundation for a methodology by which sets and taxonomies of argument schemes might be evaluated

With these objectives in mind, the next section draws on the raw data from the corpus in making a set of observations and generalisations.

Observations

The first, and most prominent, feature of the dataset is the pre-eminence of normative argument, and specifically, of the two schemes in the (Katzav & Reed, 2004b) taxonomy, *Argument from the Constitution of Positive Normative Facts* and its counterpart, *Argument from the Constitution of Negative Normative Facts*. Across the corpus as a whole, these two occur in a little over one quarter (26%) of all arguments.

Such normative arguments conclude with what should be the case or what should happen – a simple example is given here. This argument is taken from the Indian Parliament, House of the People, Synopsis of Debates, 9 August 2002. Individual argument components (roughly, complex propositions) are shown in boxes, with arrows indicating analysed relationships between them. The dashed box indicates a reconstructed premise – this example, like most in the corpus, is enthymematic. The scheme is marked by a coloured area around the argument diagram components from which it is composed, and named at its conclusion.

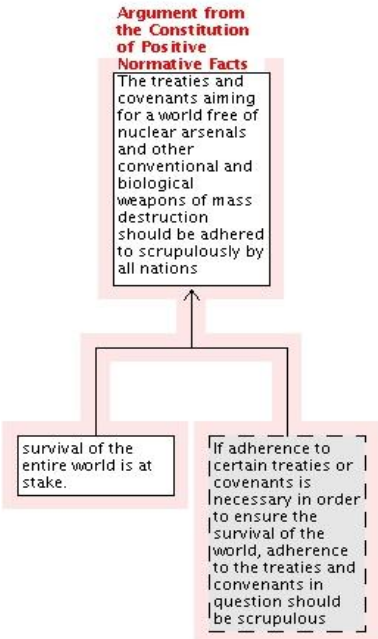

Figure 1. An example of an *Argument from the Constitution of Positive Normative Facts.*

It is perhaps unsurprising that normative argument should be so common in the "wild" - argument in many of the domains from which the corpus is drawn is used normatively, i.e. to shift opinion on what should be the case. Reflecting on our own experience, newspaper editorials often make a case for what should happen with respect to some hot news topic; parliamentary debate often involves arguing for what should be an appropriate course of action; legal argument discusses what someone's fate should be; discussion forums involve heated debate about what should happen. In fact (perhaps as an indication of the unreliability of such reflection) the corpus suggests that normative argument is much more prevalent in newspapers and parliamentary debate than it is in the law courts. But nevertheless, it is encouraging that our intuitions accord with the corpus data.

Perhaps less obviously, it is interesting that normative arguments with a clearly positive conclusion (i.e. that use *Argument from the Constitution of Positive Normative Facts*) are much more common that those with a clearly negative conclusion (i.e. that use *Argument from the Constitution of Negative Normative Facts*) – by a factor of around two and one half (18% of arguments positive by comparison to 7.5% for negative). This may be as a result of a rhetorical rule based at least in part in the social psychology of message adoption (McGuire, 1974) – positive conclusions are more likely to be accepted than their negatively phrased counterparts. (Indeed the negative expression of even very simple facts has, through a venerable series of psychological experiments, shown to confuse subjects' reasoning capabilities, (Wason, 1966)). This strong bias holds across the entire corpus, and is manifest in each domain. Some domains, however, show distinct identities in terms of the argumentation schemes that are employed.

A good example is the scheme *Argument from Implication*, which explicitly builds a deductive structure. Although not entirely uncommon, occuring in 14% of arguments in the corpus, it is worth noting that the distribution of that 14% is not at all even – there is one example of a parliamentary record using it and three legal examples, whilst the remainder (11 further examples) all occur in newspaper and magazine editorials. An instantiation of this scheme is shown below – taken from *Mail & Guardian Online* (South Africa).

One possible explanation for the disproportionately high frequency of the scheme *Argument from Implication* in popular press editorials concerns expectation and appearance. Editorials are supposed to be strongly argumentative, with a clear standpoint in the pragma-dialectical sense (van Eemeren et al., 1992). One of the ways of conveying such clarity and of developing a strong, characteristic argumentative flavour, is to use relationships between discourse components which themselves have clear argumentational roles. *Argument from Implication* fits this bill admirably. Further support for this contention is offered by the fact that *Argument from Implication* is often associated with strong clue words such as *therefore*, *because*, and *as a result* which signpost an argument, making its structure clearer to the reader - and thereby also making clearer the fact that it is an argument. Of course, this role for clue words is well known both in (computational) linguistics (Knott, 1997) and in argumentation theory (Snoeck Henkemanns, 2003) – in the latter, it is often used as a mechanism for helping students learn first to identify and then to analyse instances of argumentation (see, e.g. a textbook such as (Wilson, 1986) pp17-23). It is also enlightening to review the full text extract of the argument above:

The notion that there is a media vendetta to prove that black people are inherently corrupt is fallacious. The simple fact is that this country is run by a black government and the upper rungs of public service are mainly peopled by blacks. And another truth beyond doubt is that the same government runs one of the more competent and forward-looking administrations on the planet. It is, therefore, demographically logical that its successes are directly attributable to black people at the helm. And it is also demographically logical that when wrongdoing takes place in the ranks of government, the probabilities are that it will be the black people running the show who will be fingered. That is simple logic.

Mail & Guardian Online (South Africa)
Editorial, "Facts not Fallacy" 6 June 2003

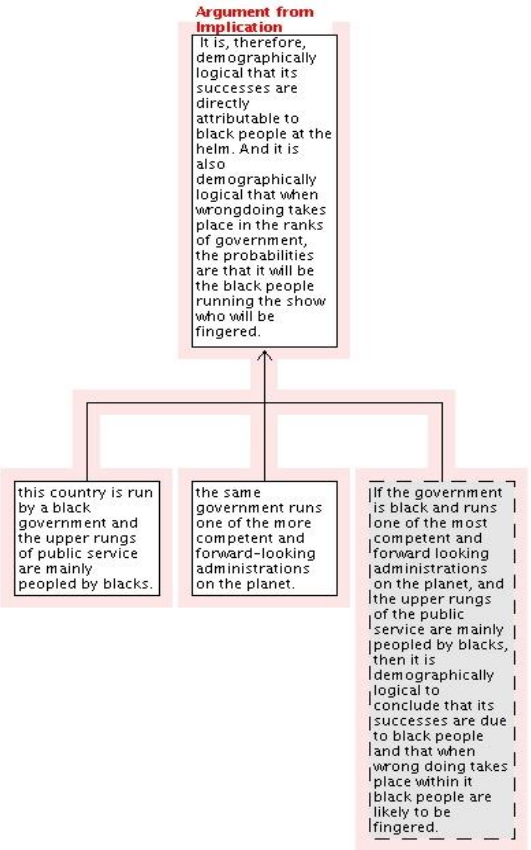

Figure 2. An example of an *Argument from Implication.*

The text not only includes several strong clue words, but also closes with a clear indication that the author is emphasising the argumentational structure and character of the text – and perhaps it is just such emphasis that *Argument from Implication* conveys, which is why it is common in editorials.

In the legal extracts in the corpus, of which there are 15 (drawn from UK and US courts), the same *Argument from Implication* scheme occurs relatively frequently (in one fifth). It may be that this is explicable in similar terms as for newspaper editorials, namely, that the strong argumentational character is a vital component of examples in the domain. Such a claim would need more data to make convincingly, but seems plausible enough. Much more interesting, however, is the observation that two thirds (61%) of legal arguments involve the scheme from the (Katzav and Reed, 2004b) set defined as *Argument from Constitution of Properties*. The template for the scheme clarifies its role somewhat:

<u>*Argument from Constitution of Properties*</u>

(1) *A*
(2) *A* constitutes the fact that object *B* has property F
(3) Therefore, *B* has property *F*

One of the simplest examples of the use of this scheme in the corpus is shown right (taken from Supreme Court of the United States, Opinions, United States, et al, Petitioners v. Thomas Lamar Bean, "On Writ of Certiorari to the US Court of Appeals of the Fifth Circuit", Cite 537 U.S.__(02), Docket No. 01-704, 10 Dec 2002).

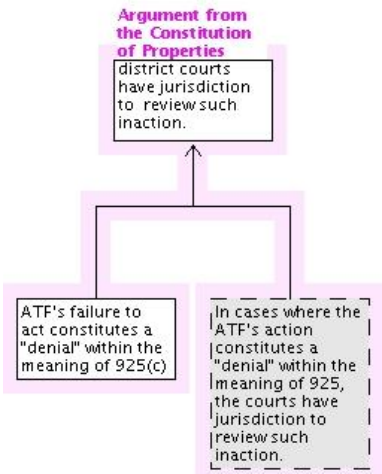

Figure 3. A judicial example of an *Argument from Constitution of Properties.*

Perhaps it is simply the case that legal argumentation makes heavy use of this form of argument as an intrinsic part of its domain. But it is also possible that the scheme – or rather the taxonomy of schemes from (Katzav and Reed, 2004b) – is somewhat lacking with respect to legal argument, in that only a relatively abstract, underspecified scheme such as *Argument from Constitution of Properties* is appropriate for capturing a wide range of legal argumentation. Empirical data of this form can therefore be used as a driver of theoretical research: the (Katzav and Reed, 2004b) taxonomy could be further refined in the area of *Argument from Constitution of Properties* to better handle the range of legal discourse.

Legal argument in the corpus is thus heavily characterised by the use of a single scheme in this particular taxonomy. But the corpus also offers an even stronger relationship between domain and scheme, whereby the *only* observations of the scheme occur within that one domain. The domain is summarised as “Discussion Forums”, and includes various online newsgroups, noticeboards and fora in which the public can contribute comments in both moderated and unmoderated forms. One of the sources is a discussion board provided as a service by the Christian Apologetics & Research Ministry (http://www.carm.org/boards.htm) . All of the arguments drawn from that source, and none others in the corpus, use the scheme *Argument from Non-Causal Law*. Though the scheme lies in the taxonomy to catch uses of laws of nature in argument that are not causal (and therefore, in the taxonomy, “external”), all instances in this domain use the same type: all are built on reference to divine laws.

A good example (Christian Apologetics & Research Ministry, Boards, Atheism, Topic #25743, In response to reply #16, 7:40 AM PST, 10th July 2003) is shown right.

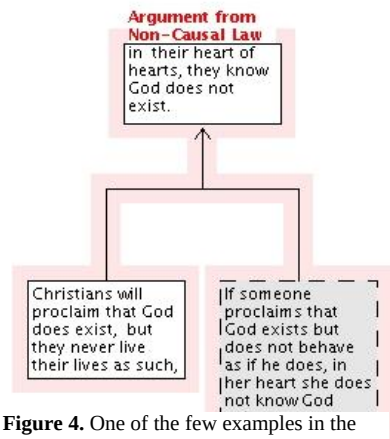

Figure 4. One of the few examples in the corpus of *Argument from Non-Causal Law*.

Why is it, then, that there is such a strong correlation between this narrow domain and this unusual scheme? The scheme set motivated in (Katzav and Reed, 2004a) clearly identifies problems with schemes that are built around argument forms, and argues instead for schemes built, at least initially, on intrinsic semantics. In other words, following Kienpointner (1992), it is the semantics of the warrant by which an argument can be classified. The domain of these arguments is one in which in addition to more traditional semantic argument forms, there is also another that is quite common – namely reference to divine law. It is no surprise, therefore, that a schemeset built on semantic grounds should uniquely identify a domain which has at its disposal a semantic inferential structure that is (virtually) unique.

The discussion so far has explored relationships between scheme usage and the domain of argumentation. There are other variables that can be explored, and perhaps one of the most interesting is to ask if there are cultural differences: with examples from various domains drawn from India, Japan, South Africa, UK, Australia, and the US, are there identifiable similarities between arguments from geographic regions or culturally similar environments, and similarly, are there identifiable differences between different such regions or environments?

Probably the most striking difference is that amongst the Indian texts, 40% use *Argument from Singular Cause*. Though not a particularly uncommon scheme (it occurs in 15% of the examples throughout the corpus), half of those occurrences are from Indian sources, despite the fact that less than one fifth of the corpus (18%) is drawn from India. The result is not confounded by domain – the Indian resources include both popular and parliamentary sources, and in any case, *Argument from Singular Cause* does not seem to be associated with domains identified in the corpus. (Interestingly, however, every single example from an Indian newspaper involved the scheme).

It is not at all clear why this should be. Perhaps as part of the discourse community or culture, this kind of causal argument is selected more often as a result of rhetorical or linguistic preference; perhaps *Argument from Singular Cause* is seen to be a more persuasive form, other things being equal. Perhaps the structure maps more closely on to Hindi or other popular languages (though the examples in the corpus are in original English – they are not translations). In any case, the finding is certainly intriguing and demands further investigation.

Finally, there is an equally peculiar, though less marked difference between the transatlantic subsets. These two are the largest in the corpus, with 33 examples drawn from the UK and 39 from the US. From early work in argumentation schemes, the difference between the direction of the inference over a causal relationhas been recognised explicitly (Hastings, 1963). That is, Argument From Cause to Effect has been clearly distinguished from Argument from Effect to Cause in almost every work on scheme usage that identifies causality at all. The same distinction is also made in the (Katzav and Reed, 2004b) taxonomy, though the exact specification differs somewhat. What is surprising is that the different geographical subsets seem to demonstrate noticeably different preferences between the two directions. So, for example, where the UK has over 12% of examples using *Argument **from** Singular Cause* and only 3% *Argument **to** Singular Cause*; the US has 8% *Argument **from** Singular Cause* and 13% *Argument **to** Singular Cause*. The following table summarises the oddity:

Country	TO cause	FROM cause
UK	3%	12%
US	13%	8%
Australia	0%	25%
India	10%	40%
Japan	0%	0%
South Africa	14%	0%

Though the data points for Australia, Japan, and South Africa are very few (8, 6 and 7 respectively), what is surprising, particularly amongst the others, is that the TO/FROM-cause bias is large, and different in different subsets. Again, this finding poses an interesting research question in first, further substantiation and then second, justified explanation of the phenomenon.

Conclusions

Clearly, this preliminary investigation is not supported by statistical analysis – on datasets of this size, any firm conclusions would be dubious at best. But the aims did not include presentation of a fait accompli in this way. Rather, this investigation serves to identify priorities as the work progresses.

Specifically, and with respect to the three aims laid out in the first section, this exploration has delivered several successes. First, it clearly demonstrates that an argument corpus such as that being built at Dundee can support interesting observations. Most of the observations here need more data to be substantiated with statistical significance. But all are sufficient to pique curiosity and to pose interesting and challenging questions of theory and practice in argument use and its relationship to context. The construction of argument corpora for extended analysis can thus play an important role in studying the expression of solo and inter-personal reasoning.

Secondly, this exploration has identified a small set of issues that can form priorities for further study. In particular: (i) the frequency of normative arguments in all debate arenas; (ii) the distribution of the sign of normative (and

non-normative) arguments; (iii) the role of schemes with strong argumentational characters, such as *Argument from Implication* in the (Katzav and Reed, 2004b) taxonomy in extracts from the popular media, and newspaper editorials in particular; (iv) the relationship between clue word usage and scheme selection; (v) the relationship between cultural or discourse community and bias in usage of schemes involving cause. As the dataset expands, the same exploratory techniques piloted here can be used to refine the research agenda.

Thirdly, as research in philosophy, communication studies, and artificial intelligence starts to push forward theories of argumentation schemes, it will become necessary to formulate mechanisms for assessing the efficacy of schemesets, and their classification systems. At least some of those mechanisms might be expected to be data driven, in that a set's success at handling real world argument is one measure of its efficacy. So, in comparing (Walton, 1996), (Kienpointner, 1992) and (Katzav and Reed, 2004b), for example, it may be useful to examine how well they characterise argumentation in different domains, particularly specialised domains such as law.

In conclusion, the world's first corpus of analysed natural argument is starting to show early signs of its potential utility. As the dataset grows, it will become possible to explore with ever finer-grained detail patterns of usage and organisation of arguments in real world settings, and thereby provide a significant empirical resource that can contribute to further theoretical development on both the philosophical and computational sides of argumentation theory.

Acknowledgements

The author would like to thank The Leverhulme Trust in the UK for its support of this work under the grant, “Argumentation Schemes in Natural and Artificial Communication”, and to Joel Katzav, Louise McIver, and Fabrizio Mancagno at the University of Dundee, all of whom contributed to the development of the corpus.

References

- Araucaria (2004) Available online at http://araucaria.computing.dundee.ac.uk/
- Hastings, A. (1963) A Reformulation of the Modes of Reasoning in Argumentation, Ph.D. Dissertation, Northwestern University.
- Katzav, J., Reed, C. & Rowe, G.W.A. (2003) “An Argument Research Corpus”, Practical Appl.s of Ling. Corpora 2003, Lodz.
- Katzav, J. & Reed, C.A. (2004a) “On Argumentation Schemes and the Natural Classification of Arguments”, *Argumentation* **18 (2)**: 239-259.
- Katzav, J. & Reed, C.A. (2004b) “A Classification System for Arguments”, Division of Applied Computing, University of Dundee Technical Report, Available from http://www.computing.dundee.ac.uk/staff/creed/
- Kienpointner, M. (1992) “How to Classify Arguments” in van Eemeren F.H., Grootendorst, R., Blair, J.A., Willard, C.A. (eds) *Argumentation Illuminated* pp 178-187, Amsterdam University Press.
- Kirschner, P.A., Buckingham Shum, S.J. And Carr, C.S. (2003) *Visualizing Argumentation*, Springer.
- Knott, A. (1996) A Data Driven Methodolgy for Motivating a Set of Coherence Relations, Ph.D. Dissertation, U of Edinburgh.
- McGuire, W.J. (1974) “The Nature of Attitudes and Attitude Change” in *Handbook of Social Psychology* pp136-314
- Reed, C. & Norman, T.J. (2003) *Argumentation Machines*, Kluwer.
- Reed, C. & Rowe, G.W.A. (2005) “Araucaria: Software Tools for Argument Analysis, Diagramming and Representation”, *International Journal Artificial Intelligence Tools* **14 (3-4)** .
- Snoeck Henkemanns, A.F. (2003) “Indicators of Analogy Argumentation”, *Proceedings of the Fifth Conference of the International Society for the Study of Argumentation*, pp969-973, Sicsat.
- Walton, D.N. (1997) *Argumentation Schemes for Presumptive Reasoning*, LEA.
- Wason, P. (1966) “Reasoning” in *New Horizons in Psychology*, Penguin.
- Wilson, B. A. (1986) *The Anatomy of Argument*, Revised Edition, University Press of America.
- van Eemeren, F.H. and Grootendorst, R. (1992) *Argumentation, Communication and Fallacies*, LEA.
- van Eemeren, F.H., Grootendorst, R. and Snoeck Henkemanns, F. (1996) *Fundamentals of Argumentation Theory*, LEA.