Translating Toulmin Diagrams: Theory Neutrality in Argument Representation

Chris Reed & Glenn Rowe

Division of Applied Computing

University of Dundee

Dundee DD1 4HN

Scotland, UK

[chris | growe]@computing.dundee.ac.uk

Abstract. The Toulmin diagram layout is very familiar and widely used, particularly in the teaching of critical thinking skills. The conventional box-and-arrow diagram is equally familiar and widespread. Translation between the two throws up a number of interesting challenges. Some of these challenges (such as the relationship between Toulmin warrants and their counterparts in traditional diagrams) represent slightly different ways of looking at old and deep theoretical questions. Others (such as how to allow Toulmin diagrams to be recursive) are diagrammatic versions of questions that have already been addressed in artificial intelligence models of argument. But there are further questions (such as the relationships between refutations, rebuttals and undercutters, and the roles of multiple warrants) that are posed as a specific result of examining the diagram inter-translation problem. These three classes of problems are discussed. To the first class are addressed solutions based on engineering pragmatism; to the second class, are addressed solutions drawn from the appropriate literature; and to the third class, fuller exploration is offered justifying the approaches taken in developing solutions that offer both pragmatic utility and theoretical interest. Finally, these solutions are explored briefly in the context of the *Araucaria* system, showing the ways in which analysts can tackle arguments either using one diagrammatic style or another, or even a combination of the two.

Keywords: Argument analysis; Argument diagramming; Argument software; Computational Models of Argument; Toulmin diagrams

1. Introduction

The analysis of arguments is often hard, not only for students, but for experts too. As a result, a variety of tools and techniques have emerged from the theory of argumentation and the theory of argument/informal logic/critical thinking pedagogy that aim to help in the task of analysis. One of the most common and intuitive of these tools is diagramming, by which the abstract form of an argument can be identified and seen at a glance, and according to which it is then possible to analyse more closely the relationships between an argument's parts. The utility of argument diagramming is seen in its almost universal adoption in the teaching of critical thinking and argumentation skills, as well as its deployment in various practical tools employed where complex argumentation is part of a profession (most notably in legal domains). There are a wide range of diagramming techniques, some very general, some tailored to particular domains. But there are two that are perhaps most well known through the various pedagogic and professional applications of argumentation theory.

The first technique is the conventional “box-and-arrow” approach of identifying atomic components of an argument, and then indicating links between them with arrows. One of the first proponents of the approach in a pedagogic context was Beardsley (1950), and little has changed since then. In addition to identifying relationships of support between atoms in an argument, the scheme has become refined to also identify four distinct ways in which compounds can be formed: as serial argument (in which one statement supports another, which in turn supports a third); convergent argument (in which two or more statements independently support a third); linked argument (in which two or more statements jointly support a third) and divergent argument (in which two or more statements are supported by a third). Complex argumentation can be constructed through arbitrarily complex combinations of these forms. As it is so familiar to many, and by analogy to the terminology used in fallacy theory, we here refer to this “box-and-arrow” approach as the *standard treatment* of diagrammatic argument analysis.

But almost contemporaneously with the development of the standard treatment, a second approach, which sprang from quite different concerns, has developed into an equally successful, well known and widely used method for diagramming, *viz*. the Toulmin schema (Toulmin, 1958). Rather than viewing arguments as essentially just more or less complex binary relationships of support, Toulmin sees arguments as six-part complexes, comprising the familiar Data, Warrant, Claim, Backing, Rebuttal, Qualifier. Though the starting point was jurisprudential, the resulting theory and its subsequent application are very general, and a Toulmin-style approach, replete with appropriate diagrams is commonplace in current undergraduate curricula.

An important observation is that both the standard treatment and the Toulmin schema are, of course, much more than just ways of drawing pictures. There is more to the standard treatment than p2 of (Freeman, 1991), in just the same way that there is more to Toulmin than p104 of (Toulmin, 1958). Both systems embody many theoretical assumptions and conclusions, and work as a way of packaging up substantial theories into practical tools that are simple and easy to understand – and produce analyses that are the products of those background theories. The motivation that drives the remainder of this paper is a challenge initiated through the apparently innocuous desire to allow diagrams in one form to be translated into another. The challenge lies in the fact that for a variety of reasons, such translation demands coherence and interaction between the two approaches at a deep, theoretical level. For if an analysis embodies the theory by which it has been constructed, then transmutation of that analysis into one constructed in according to a different theory demands some sort of tying together of the two theories.

2. The Computational Angle

Recently, there have been an increasing number of applications of argumentation in computational sciences, including computational linguistics, artificial intelligence and distributed computing (Reed and Norman, 2003). These different applications have used different accounts of argument: some, such as (Bench-Capon, 2003) have been rhetorical in style, using models from Perelman and Olbrechts-Tyteca (1969); others (perhaps rather surprisingly) have involved predicate calculus interpretations of Toulmin, e.g. (Fox & Das, 1996). Many have developed intuitive models that are closest to the standard treatment: (Dung, 1995) offers a particularly influential example.

One area in which computing has contributed back to the theory and particularly pedagogic practice of argument is in the design and implementation of software tools to support diagramming. There are many dozens of such systems each with different strengths and weaknesses: Kirschner *et al*. (2003) offer a review of many of them. Interestingly, most of these systems adopt the standard treatment, rather than a Toulmin structure, as the basis for diagramming (though there are exceptions: (Chryssafidou & Sharples, 2002) is a Toulmin-based diagramming tool, and a number of those described by Kirschner *et al*. use models from decision making such as IBIS (Rittel & Webber, 1973)).

With this exciting and rapid proliferation in computational applications and diagramming tools, has come an interest in standardisation and interoperability. To our knowledge, the only application to date that aims to tackle argumentation representation within different theoretical frameworks is *Araucaria* (Reed & Rowe, 2001; 2004). *Araucaria* is a software tool designed to support the manual analysis of argument, and the subsequent storage of that argument for a variety of computational and pedagogic purposes. Though the software has enjoyed relatively widespread use in universities and colleges, it was until recently severely hampered, simply because many instructors value the clarity and other pedagogic virtues of the Toulmin model. From an entirely pragmatic perspective, then, the development of a Toulmin model in *Araucaria* has practical classroom utility.

Working specifically in the context of *Araucaria*, the aim is to develop both the theory and implementation (in software) of diagramming tools that are “theory neutral”. Specifically, therefore, there are several objectives:

(i) It should be possible for diagrams to be constructed in more than one theoretical framework

(ii) Resulting analyses should be stored in a common format or interlingua

(iii) Partly as a result of (ii), it must be possible to convert, in software, from a diagram in one theoretical framework to an equivalent in another

(iv) Conversion according to (iii) must be consistent and deterministic, and should not require additional input from the user

(v) Analysts working solely within one theoretical framework should not be impacted at all by features, contrivances or oddities from other theoretical frameworks

Perhaps surprisingly, simple image processing software, that is not specialised to argument at all, fairs well against these objectives, with the exception of (iii) and (iv). It is these that make the task one that is specialised to argumentation. Together, then, these five objectives for software development frame a research project that tackles both theoretical and practical strands of work. The two strands are intimately connected, with practical, classroom needs dictating the remit of the more abstract work, and the equivalences developed theoretically tested in practice in software with real examples.

3. Translation

3.1. Atoms

A variety of initial translation “building blocks” are relatively unproblematic. The first concerns atomicity. The notion of what constitutes an argument or an atomic component of an argument (Katzav and Reed, 2004; Parsons, 1996; Wreen, 1998) is highly contentious. There is an enormous range of opinions on what might reasonably stand as a premise or a conclusion. Yet the theoretical framework in which such analysis might take place is for the most part independent of such differences: whether one analyses according to the standard treatment, to the Toulmin schema, to the pragma-dialetical model, or to some other system has relatively little impact on the size and character of the stretches of lexical material from which the argument are constructed. In textbook presentations, there are canonical examples (e.g. of warrants, or conclusions, etc.) that simplify the picture and sharpen contrasts between theoretical apparatus, but when analysing real world examples of argument, textual complexity brings atomic concepts of one theory into direct correspondence with those of another. Our conclusion, then, is that atoms in one theory are not substantially different from another: a Toulmin datum might be associated with a standard treatment premise, for example.

3.2 Warrants

The second building block concerns the simplest structural translation: from a Data-Warrant-Claim (henceforth, DWC) complex of a Toulmin analysis, to a linked argument in the standard treatment. Linkage expresses some need for both components to be present (explicitly or implicitly) in order for the argument to go through. Sometimes, moving from one linked premise to another functions as a way of delivering or manifesting relevance between premise and conclusion. Freeman (1991) offers a dialectical analysis, such that a linked premise is elicited by asking, “Why is that [first premise] relevant [to that conclusion]?” Freeman's discussion is tabled in the context of the Toulmin warrant, and although he identifies many problems with explications of the latter, we are not here trying to critique either of theories involved. Instead, the aim is to adopt them “warts and all”, and provide mechanisms for those that adhere to one or the other (or both) to work within their frameworks. There seems, in this context, to be good reason therefore for identifying the DWC complex with a linked argument in which a single conclusion is supported by two linked premise. In combination with the direct mapping of argument atoms, the following is therefore a reasonable intertranslation[1]:

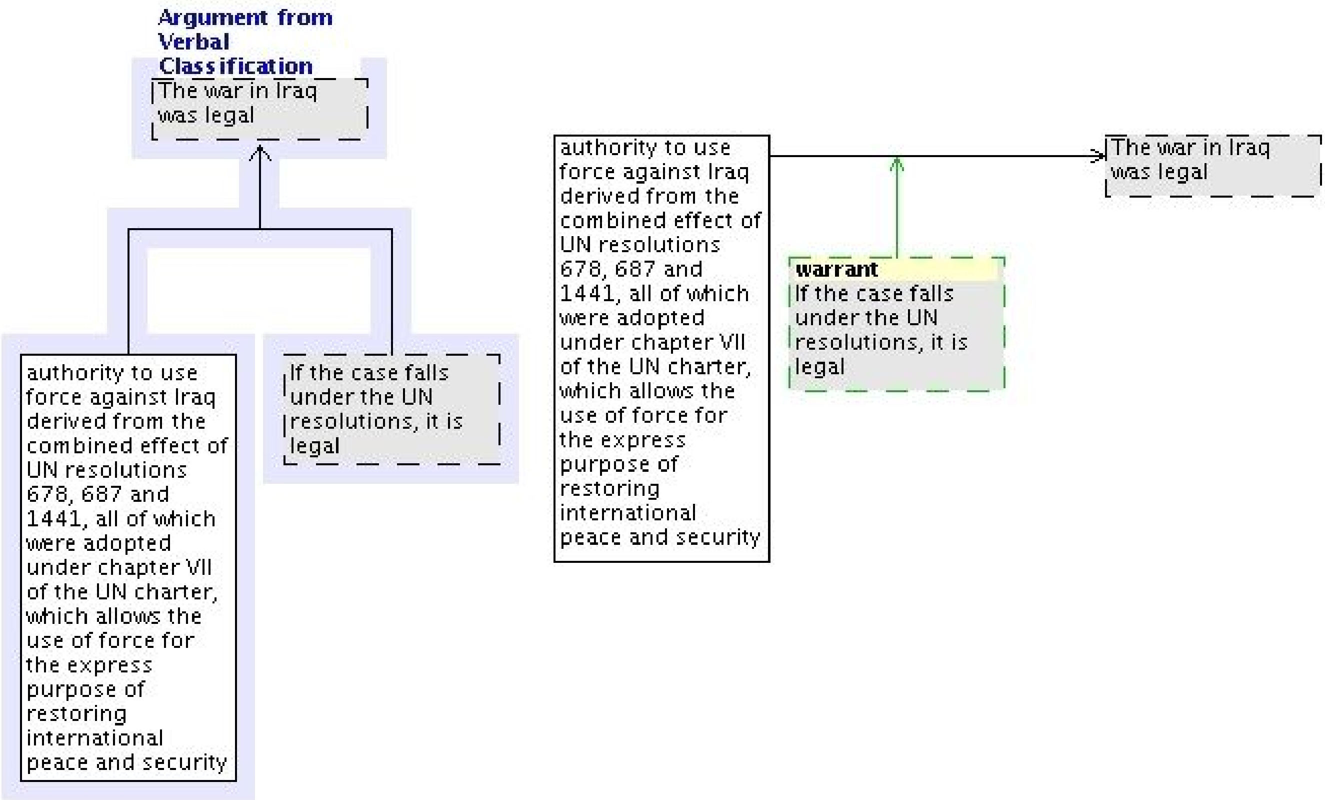

Figure 1. A linked argument as a single DWC complex

(The example is from Hansard, and is taken from the *AraucariaDB* online corpus). It is important not to read too much in to Figure 1. Specifically, it is not making ontological claims about the interpretation of one language for expressing argument *by* another (though it is providing interpretation of one language *in* another). For some authors, presumably including Toulmin, the warrant is most certainly not a premise (Hitchcock, 2003): the Toulminian framework is simply and deeply richer than that. Yet in the standard treatment, there are no other ontological categories. With a Toulminian analysis of some argument that yields a DWC complex, how might a standard treatment analyst go about performing the same analysis? Without a concept of warrant, it seems reasonable that that analyst might view two components of the argument as linked premises – and one of those would happen to correspond to what the Toulminian identified as a warrant. In this way, there is no implicit claim that either analysis is right, or more right, or more basic, but merely that the analysis conducted in one framework might be rendered in such a way as to make sense to an analyst from another framework. This is what Figure 1 is depicting.

Figure 1 also highlights a closely related point of correspondence in the translation. In the standard analysis, the example has been marked as an instance of a particular argumentation scheme. Toulmin analyses do not usually employ such devices. However, as both Walton (1997) and Katzav & Reed (2004) have pointed out, argumentation schemes can be seen as analytic devices for handling warrants. So perhaps Toulmin warrants are best translated as instances of particular schemes. There is unfortunately a practical problem with this solution. Where Toulmin analyses identify a warrant explicitly, and can therefore mark any number of different warrants as such, models of argumentation schemes are in their infancy and are consequently limited. Though efforts at taxonomic work have begun (*vide* Walton and Katzav & Reed, *ibid*), coherent and wide-ranging systems of argumentation schemes are simply not available. Worse, it is far from clear that a single, uncontentious, field-independent system might ever be produced. If there is no way of guaranteeing the schemes that might be available, it is difficult to use those schemes in translation. There is, however, a way forward. With recent proposals for seeing argumentation schemes as (at least in part) "shorthand" for a number of characteristic implicit premises or sets of implicit premises (Bex *et al*., 2003), the explicit warrant of a Toulmin analysis might be seen as one of the linked premises associated with a scheme in a standard analysis. This leaves scope for the analyst to identify a scheme from those that are available, but does not necessitate such identification in the automated translation process. We return to the issue again in the context of rebutting, below.

One challenge remains. In the standard treatment, a linked argument can have any number of premises; a Toulmin analysis on the other hand typically has a single datum and a single associated warrant. How then, can many-premised, linked arguments be faithfully represented in Toulmin schema? One possibility is simply to ignore linked premises beyond the first two – i.e. a Toulmin analysis recognises exactly one D, W and C in each DWC complex. This is unattractive because it fails to preserve information between frameworks. Perhaps, then, an additional premise in a linked argument might be seen as a fulfilling one of the other roles in the Toulmin model. Unfortunately, there are no other roles that could be filled in a consistent way: the relationship between backing and warrant is most closely similar to the relationship of support in the standard treatment – and not the relationship holding between “sibling” premises. The only alternative left open is to broaden what Toulmin diagrams can handle, either by allowing more than one datum in an argument, or allowing more than one warrant. Permitting more than one datum in a single DWC complex requires extension to the Toulmin model that is unnecessary (multiple data, each with corresponding warrants, all supporting the same claim is a different problem, tackled below). A single datum seems to offer a single basis from which to build an argument to support a claim. The final option – to permit multiple warrants – is a little strange, but not downright offensive to the Toulminian theoretical framework, particularly given more recent exegesis: 'The question [for a warrant] is not “How do you get there?” but “How might you get there?” And then: “Is one of the ways you might get there a reliable route?” ' (Hitchcock, 2003: §4) So, perhaps the best (default) Toulminian interpretation of a standard treatment analysis involving more than two linked premises is of an argument with more than one warrant. Though taking liberties with the Toulmin picture, this meets objectives (iii) and (iv) from the introduction, and most importantly, means that as described in objective (v), analysts working in either tradition needn't worry about the foibles of the other (just because Toulmin diagrams can be constructed in which more than one warrant supports the move from datum to claim does not mean that such analyses will be at all common for those working in the Toulmin framework).

3.3 Complex Arguments

A similar approach is required with another general problem. The standard treatment allows the construction of analyses of arbitrary complexity and depth. In this respect it is like most methods for analysing, synthesising and representing argument, including Wigmore's (1931) method of analysing legal argument, Pollock's (1995) argumentation based reasoning system, and so on. Toulmin was unconcerned with such larger scale structures, and focused therefore upon the simple, individual argument with its six components. The simplest solution to this problem is to see each of the components as points for expansion. That is, if a given argument, A, for some claim comprises a datum, a warrant, a backing, a rebuttal and a qualifier, then each of those five components might stand as a claim in some other argument, B. In this way, Toulmin arguments can be glued together. From a computational point of view, this is extremely attractive, as it allows a simple recursive definition that supports Toulmin diagrams of arbitrary complexity. Partly as a result of this attraction, and partly because the solution seems for the most part to remain true to the spirit of the original work, this recursive definition has been adopted in most computational uses of the Toulmin model, such as Fox*et al.*'s (1996) implementation of automated Toulmin-esque reasoning, and Bench-Capon's (1998) dialectical system called simply, the Toulmin Dialogue Game. The same approach is followed by *Araucaria*. Figure 2a shows an example taken from a standard critical thinking textbook that is there used to illustrate serial argumentation (Hoaglund, 1999: 122). The guidance in the textbook would yield the three unshaded boxes in the standard treatment diagram. In anticipation of the translation, two further premises have been inserted by the analyst to form two linked arguments. (Of course, the analysis is arguable: the aim here is simply to ensure a faithful interpretation).

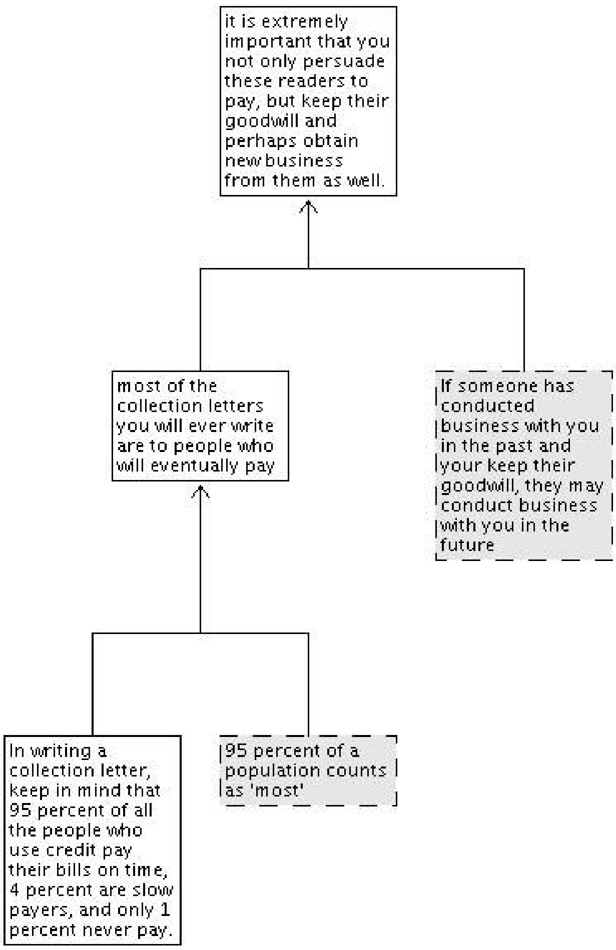

Figure 2a. Serial (linked) arguments in the standard treatment

Its conversion to a Toulmin diagram is then quite straightforward, with the datum of the righthand argument forming the claim of the lefthand:

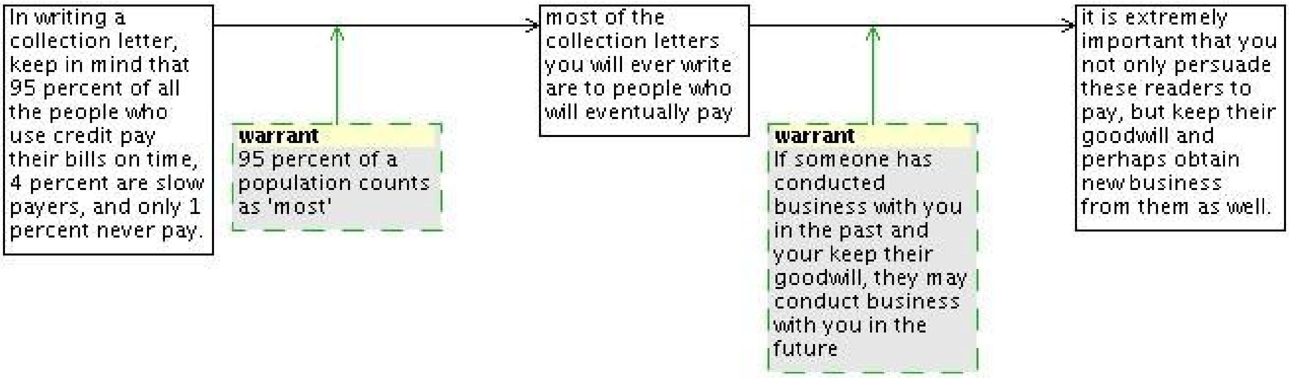

Figure 2b. Serial arguments in Toulmin form

In a similar fashion, computational approaches usually allow multiple instantiations of Toulmin schema to support a single claim, as in Figure 3, which shows a large example from the *AraucariaDB* online corpus. Each horizontal arrow represents a single DWC complex, sometimes with warrants present, sometimes not, and with reconstructed components in shaded boxes.

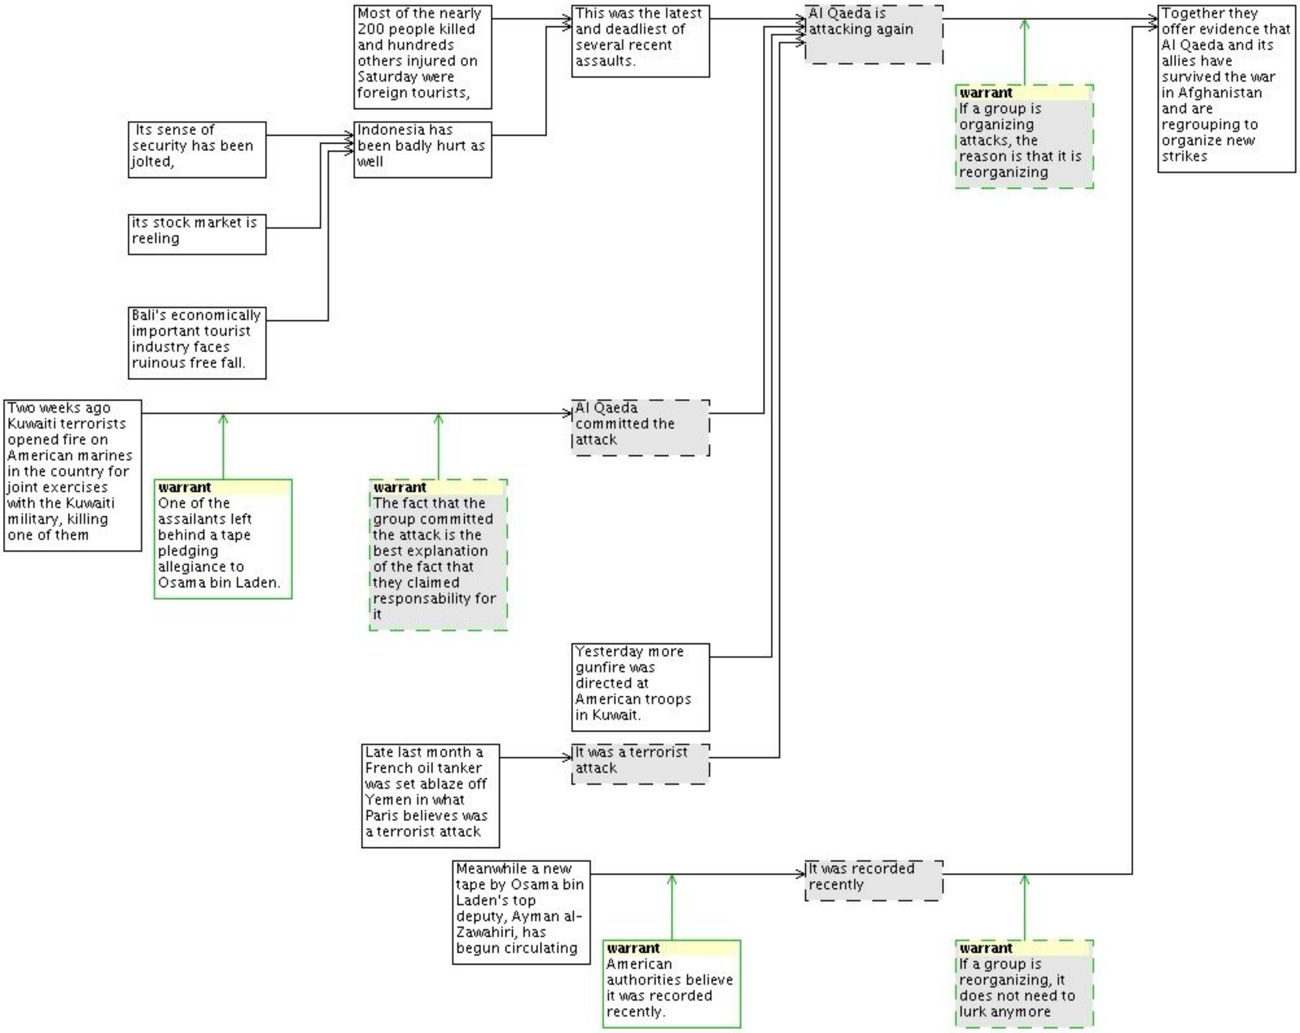

Figure 3. Large Toulmin analysis

3.4 Qualifiers

In fact, for various reasons, all Toulmin-based computational systems adapt the original model somewhat to their own purposes. One of the effects of such adaptations is that not all of the five components that form an argument for a claim are able to stand as claims in other arguments. In *Araucaria*, the only exception is the qualifier. In the standard treatment, qualifiers (called "evaluations" in the standard treatment tools in *Araucaria*) are rarely included in analyses and diagrams: when they are, they are identified with support relations. Thus, qualifiers work as a modality or modifier, expressing the degree of support captured by an arrow between two argument components. This role is very similar to that carried out by the qualifiers in the Toulmin approach, so a mechanism for exchange suggests itself, as in Figure 4. Here, we use Toulmin's introductory example.

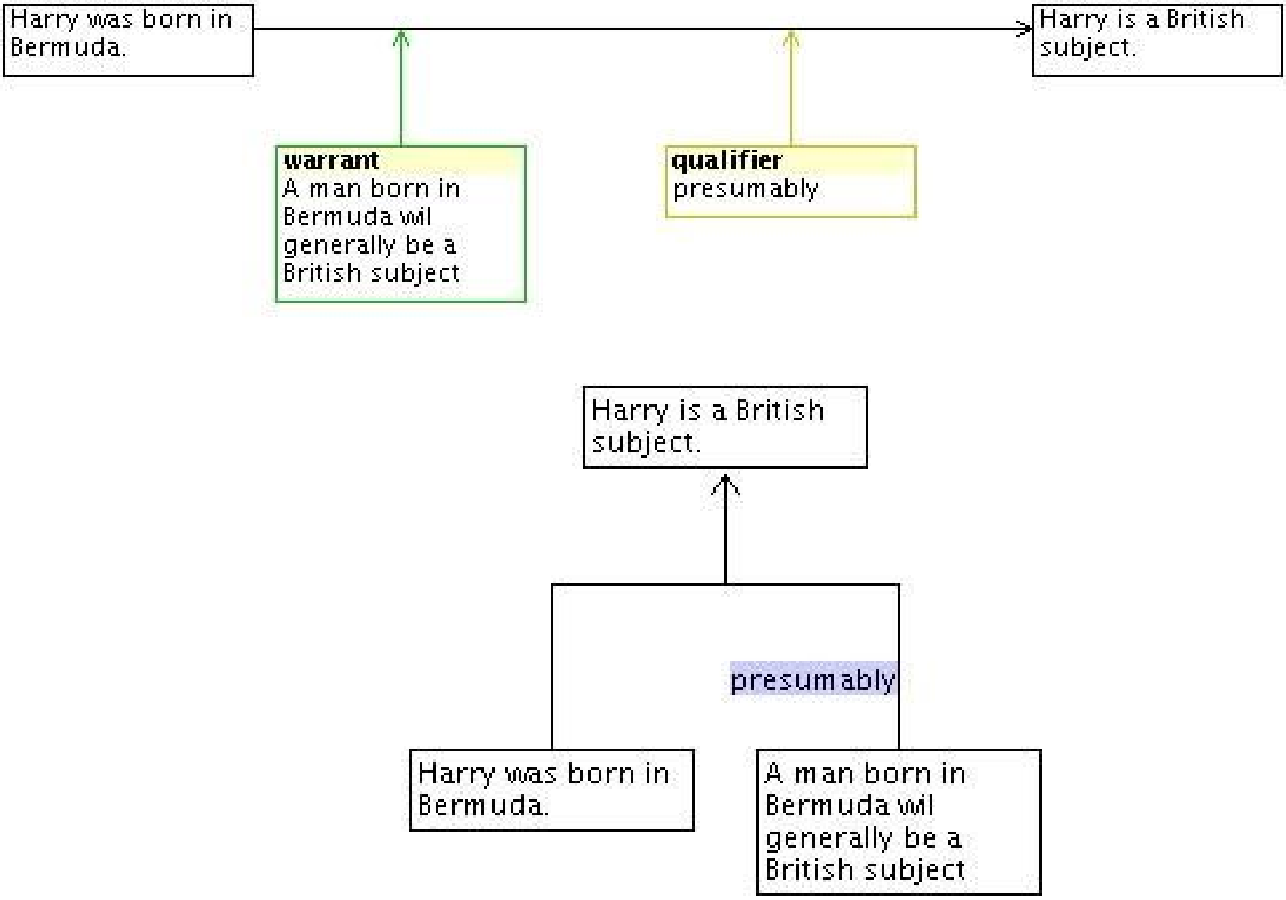

Figure 4. Mapping evaluations and qualifiers

If there are multiple evaluations in a standard analysis, it is the one attached to the warrant (i.e. the premise that stands in the Toulmin “warrant” role) that is identified as the qualifier. Most of Toulmin's examples suggest that the qualifier modifies the scope of the warrant, or is picking out the defeasibility, plausibility or presumptive nature of the warrant (so, in the example above, it is the “generally” of the warrant that seems to require qualification of the claim with “presumably”). Toulmin points this out explicitly in introducing qualifiers: “Warrants are of different kinds, and may confer different degrees of force on the conclusions they justify” (Toulmin, 1958: 100).

3.5 Backings

The penultimate translation issue, backing, is interesting in that a number of computational interpretations of Toulmin omit the category entirely. Backings are perhaps the clearest indication of the jurisprudential heritage of the Toulmin model, indicating links to legal precedents, case law and so on. Though these are adapted in some work (so, for example, Fox & Das (1996), use backings to indicate links into the medical literature in their Toulmin model of oncological reasoning), the challenge is often that the relationship between backing and warrant is identical, ontologically and formally, to the relationship between a warrant in one DWC complex that is standing as a claim in another, and the datum in that other complex. Referring to Figure 5, if we permit the recursivity conditions that allow arguments of arbitrary complexity, as discussed above, then we must permit diagrams such as the one on the left. But in that case, it is difficult to see how it differs in any important way from the analysis on the right.

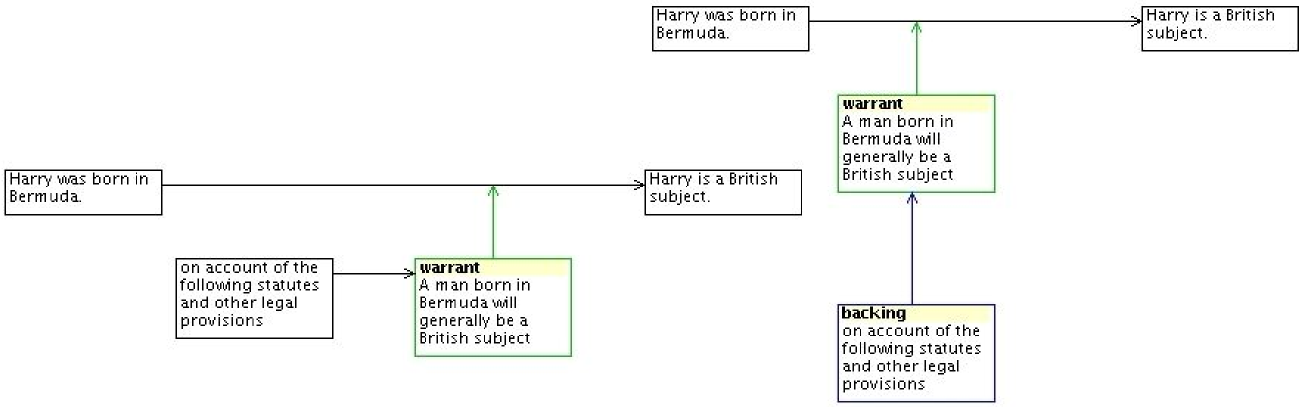

Figure 5. Two ways of supporting a warrant.

The argument from a Toulmin perspective would undoubtedly turn upon the nature of the material that constitutes the "on account of" component in Figure 5. And that determination may be context-dependent: in some circumstances one may want to ask of the leap from backing to warrant, "How did you get there?" - and if so, perhaps the lefthand approach is analytically clearer. Alternatively, if one is focusing upon the data-claim link, and the backing is merely an anchor in the literature of the field, then the relationship between the backing and the warrant may be self-evident and the right hand analysis be more appropriate. Similarly, if the literature (case law, medical scholarship, etc.) is itself inconsistent and contestable, then perhaps there will be call for analysis between those components – in that case, the recursive structure of either the left hand picture, or even the more complex structure of Figure 6, is more appropriate.

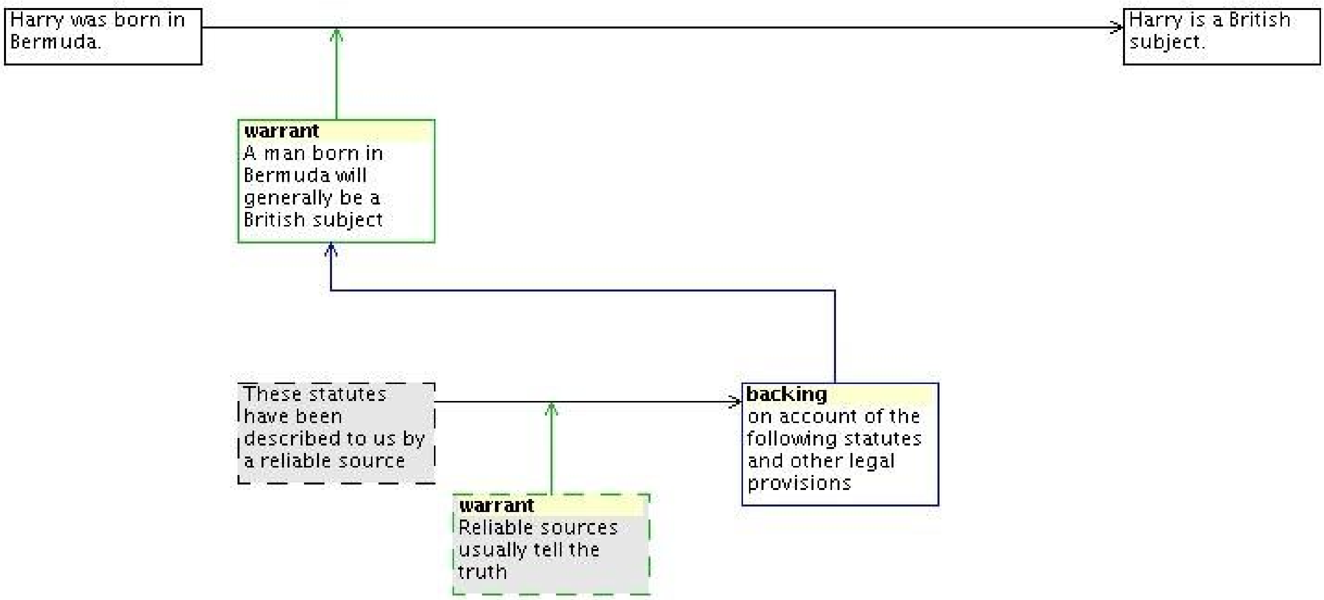

Figure 6. Supporting a backing.

The conclusion is that there are good reasons why an analyst might use any one of these structures in a given case, so it is neither the job of theoretical structure nor software tools to proscribe any of them. Each is permissible in *Araucaria*. The default is the conventional DWBC structure, but changing to either alternative is a simple matter. The decision as to which of the two diagrams in Figure 5 (or alternatively the one in Figure 6) should apply in a given situation is a decision that only has meaning within the Toulmin framework. Without the ontological difference (and rather slippery ontological difference at that) between backing and data-supporting-a-warrant, the standard treatment cannot and should not distinguish between the two approaches. As a result, the two analyses in Figure 5 should be rendered identically under the standard treatment. In order to meet the objectives listed in the introduction however, it is important that (a) any analytical decision taken under the Toulmin view is implicitly preserved even under the standard treatment, and (b) there is a deterministic way of identifying an appropriate Toulmin interpretation of the Toulmin-ambiguous standard treatment analysis. The solution to (a) lies in a general approach which is also employed in the data/warrant distinction, where there is an analogous problem of under-specification in one theoretical framework with respect to the distinctions that can be made in another. Each theoretical framework has various roles that it identifies for the atoms of argument. Those roles can be characterised by restrictions upon how they interact. A Toulmin backing, for example, cannot stand to support a qualifier. The identification of which components stand in which roles can only be carried out within the appropriate theoretical framework: identifying a component as a backing can only be carried out in the context of a theory, such as Toulmin's, that has backings. These role-assignments, by their very nature, cannot easily be represented in some lowest common denominator of an argument theory: any technique for translation must simply respect the differences in the target theories. But on the other hand, the mechanisms for translation must nevertheless be principled and well defined to avoid a combinatorial explosion in the effort of translating either large arguments, or arguments between large numbers of theoretical frameworks (that is, we do not want to have to build $n - 1$ new translation mechanisms on encountering the *n*th theory). The approach taken in *Araucaria* is to allow theory-specific roles to be identified, represented, and stored explicitly in the underlying language. So, in that language, a single component of an argument may simultaneously instantiate a backing role in the Toulmin theory, and a premise role in the standard treatment theory. Equally, if the analyst has not specified a role for a component in a given theory, then that role is simply undefined - or, more accurately is defined implicitly and by default through the semantics implemented in software. This is the solution to (b), in that a default translation is applied. Here, that default is to interpret support for warrants as a new DWC complex rather than as a backing – though in this case and in general, such defaults can be overridden by the analyst.

3.6 Rebuttals

The final component of the Toulmin picture is perhaps the single most troublesome – and most interesting from a theoretical point of view: rebuttals. Most standard treatment systems involve some mechanisms for identifying conflicts: propositional negations, counter-positions, incompatibilities, etc. For some reason, there does not seem to have emerged a consensus on how best to deal with the issue diagrammatically. This has transferred directly into software implementations of diagramming methods: Reason!Able, for example uses coloured arrows (van Gelder, 2003), Argue! has lines terminated in diamonds (Verheij, 2003) and so on. *Araucaria*'s solution is to use double-headed horizontal lines, and to restrict any given proposition to a single conflicting proposition (though that proposition in turn may have an additional conflicting proposition that is not the first, and so on). Whatever the exact mechanism for handling and representing these conflicts, the challenge is the same: is it possible to construe Toulmin rebuttals in terms of standard treatment refutations?

There seem to be (at least) four possible standard treatment interpretations of the Toulminian notion of rebuttal, summarised in Figure 7.

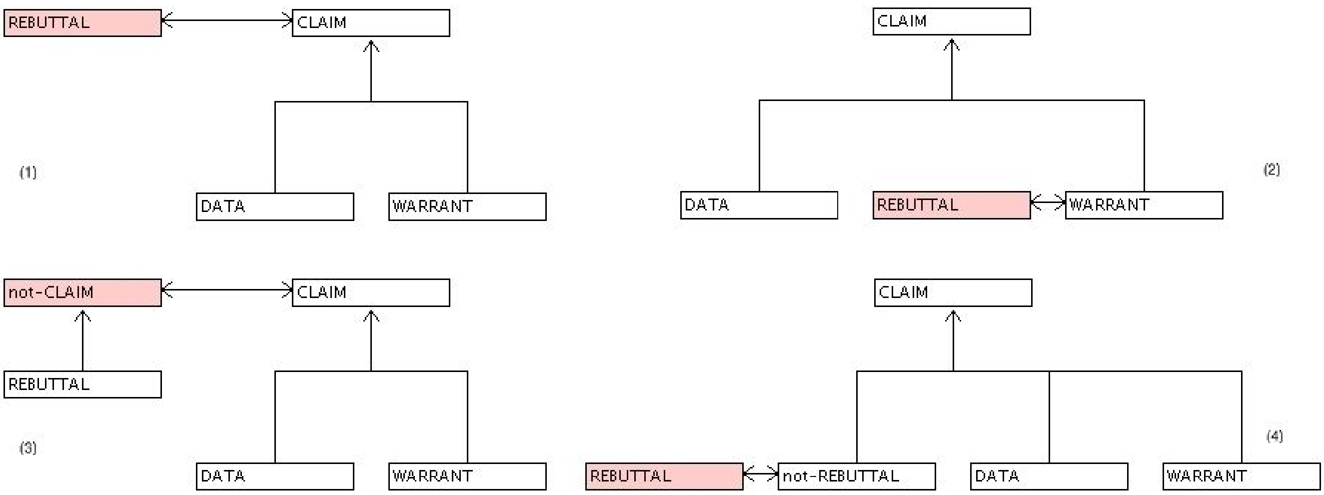

Figure 7. Four candidate standard treatment interpretations of Toulmin rebuttals

The first candidate is that a rebuttal refutes its claim (we use *rebuttal* to refer specifically to that Toulmin role, and *refutes* to refer specifically to the countering relationship expressed by a horizontal line in *Araucaria*'s implementation of the standard treatment). The single largest problem with this approach is that it seems to fail to capture accurately the function of the Toulmin rebuttal. Not only the examples in (Toulmin, 1958), but even the very diagrams that label rebuttals with "unless", suggest that rebuttals function not to refute the claim, but to capture exceptions, objections or ways in which the argument may not apply (and may perhaps not apply in the case at hand). In this way, rebuttals are functioning in a manner akin to undercutters in Pollock's (1995) terminology[2]. Undercutters take on the role of defeating an argument by attacking the inference, the way by which a conclusion was derived. Of course, in the Toulmin framework, the "way by which a conclusion was derived" is captured specifically by the warrant. Perhaps then, a second possible interpretation is more favourable: the rebuttal refutes the warrant. Again, though, this perverts the explication laid out by Toulmin. In the initial example, used in Figure 4, above, the warrant is "A man born in Bermuda will generally be a British subject". It is surely not the case that the rebuttal, "Both his [Harry's] parents were aliens" refutes this general statement. Even if the rebuttal is true in a specific circumstance, the general presumptive rule might nevertheless hold true. It might be argued that what the rebuttal does serve to do in this case is to lend implicit support to the conclusion that (in this case) Harry is not a British subject. This, then, offers a third possibility: that a rebuttal supports a refutation of the claim. The claim, C, has some counterposition which might be expressed loosely with the gloss, “it is not the case that C”. This component itself is then supported directly by the rebuttal. Though this seems to work in the Harry case, it captures our intuitions poorly since the rebuttal is now interpreted as being entirely distinct from the data and warrant – under this interpretation a rebuttal is interacting only with the claim, and not with the way in which the claim is being derived. Furthermore, if the relationship between rebuttal and Pollock-style undercutter is close, then Pollock's analysis is in direct conflict with this third option, for, crucially, undercutters do not offer support for any counter to the conclusion. Pollock offers the example shown in Figure 8:

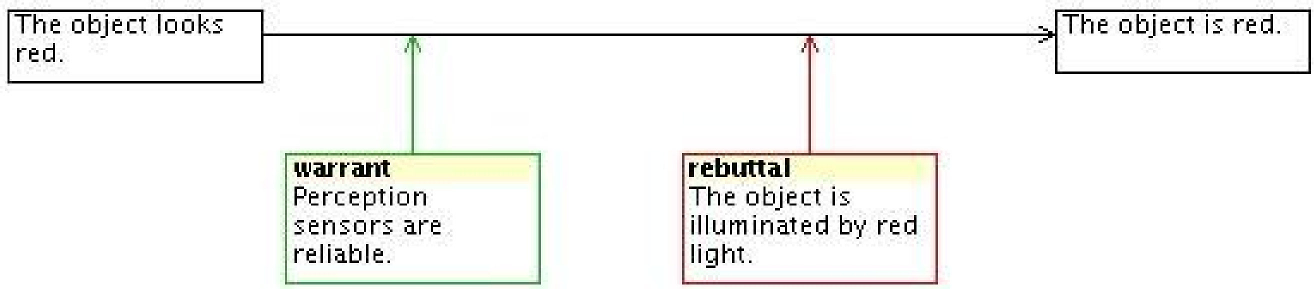

Figure 8. Pollock-style undercutters as Toulmin rebutters

Here, the fact that an object is illuminated by red light offers no support whatsoever for concluding that the object is not red. But it certainly casts doubt on the inference that it looking red suggests that it is, in fact, red.

Is there, therefore, a way of capturing this undercutting style of attack that seems so close to the Toulminian notion that a rebuttal serves to identify objections or exceptions to the way in which the conclusion has been reached using the warrant? There are two ways of achieving such a representation that are structurally identical, but semantically quite different. The first is to reify the inference. In this way, the DWC complex implicitly includes another component – represented, perhaps, by the horizontal line. The inference then runs, roughly, given the datum and the warrant, it is reasonable to conclude the claim. It is this implicit premise, that the rebuttal refutes. The approach has a direct counterpart in more traditional models of inference. A conventional approach to first order logic uses the principle of Modus Ponens to get from premises A and $(A \rightarrow B)$ to conclusion B. But it is just as reasonable to extract the leap of faith or "inference rule" and identify it explicitly, as a premise: A, $(A \rightarrow B)$, $(A \wedge (A \rightarrow B)) \rightarrow B$. The Carrollian regress looms instantly, and threatens the Toulmin model in an identical way if we go down this path. In addition to being a sly way of "deductivising" any non-deductive theoretical framework, a further problem is that it is far from clear that having the rebuttal refute this implicit premise is any better than having it refute the warrant. It may well be that the datum and warrant do still plausibly support the claim, even if the rebuttal holds.

The final alternative then, is to introduce an implicit premise, but have that premise represent nothing more than the counter of the rebuttal. This implicit premise might be seen (by the analyst) as an additional warrant. It could be that it is an attack on the entire inference scheme. It could be a specialisation of the warrant that is expressed. But perhaps the most common and accessible interpretation will be that this missing premise is some kind of implicit assumption. In this way, it is very similar to the implicit components expressed in argumentation schemes (Walton, 1997; Katzav and Reed, 2004). The approach taken in *Araucaria* (partly because it is designed also to handle such theoretical structures) is to use this scheme-like approach in implementation (cf. Bex *et al*. 2003). This approach naturally handles "conditions of exception or rebuttal" (Toulmin, 1958: 101) and "circumstances in which the general authority of the warrant [should] be put aside" (*ibid*.) as well as the full range of interpretations of rebutting used by analysts based on Toulmin's brief and ambiguous presentation. It also means that there is a clear relationship between components of argumentation schemes in the standard treatment and their (automatic) characterisation in Toulmin diagrams.

There remains a problem. The function of a rebuttal in a Toulmin diagram is, on our understanding of it, one of challenging an inference. The function of standard treatment refutation, at least as implemented in *Araucaria*, is one of representing some sort of dissonance between statements. These two theoretical frameworks thus manifest a fundamental difference in the way they handle inference: essentially, the former has a metaphysical basis that identifies multiple forms of inferencing, whilst the latter is cast in the deductivist mould. The only straightforward way in which translation between them might be accomplished is to reify the inference types of the former, so that they can be represented explicitly as statements in the latter. The problem then, is that it might be argued that the richer model is weakened by its translation to the more formal model. The first observation to make in response to such a challenge is that it is interesting and perhaps surprising that an apparently simple diagramming translation problem is intimately tied to the great deductivist debate that is still going strong (witness, e.g. (Groarke, 1999) and its responses). We do not here seek any kind of resolution of that debate, but rather seek to build a pluralistic approach that allows analysts and researchers to work within their many theoretical frameworks, allows work conducted in one to be re-used in another, and, perhaps, allows research exploring the differences between frameworks to have practical support.

4. A Common Interlingua

From a technical point of view, translation between two importantly different approaches requires an interlingua, or generic representation, that is very broad. Though clearly just two theories do not guarantee immediate generalisation to all structural theories of argument, the discussions in the previous section demonstrate that many key issues have been tackled. Initial work into the approaches of Pollock (1995) (which in many ways is similar to Toulmin) and Wigmore (1931) (which has been compared with both Toulmin and standard treatment theories in (Bex, *et al.*, 2003)) suggests that extensions to the interlingua itself are likely to be relatively minor.

Reed and Rowe (2004) describe AML, the Argument Markup Language in which analyses (such as those in the corpus) can be stored. AML encapsulates a relatively straightforward markup scheme, in which statements of an argument are individualised, and the relationships between them stored explicitly. AML, then, is the "deep structure" from which standard treatment, Toulmin, or, in the future, other diagrammatic or linguistic "surface structures" can be automatically generated. Thus, for each new theory there are two translation processes, one from AML to that theory, and one from that theory to AML. Each translation process must meet the objectives (iii) through (v) discussed above, but once in place, it is possible to translate from any one theory to any other in two steps, via AML. To ensure the precision of such translation, it is crucial that any theory-specific analysands are preserved through AML. Thus if two theories A and B are the only two that make use of some feature, X, it is important that X is preserved in AML, regardless of whether or not any other theory recognises X. Given that it is, in general, impossible to determine all the possible instances of such features a priori, AML instead offers a generic mechanism for representing theory-specific features. On each analysed statement, any number of *role* tags can added. Each *role* tag indicates a *class* and an *element*: the former for a theory identifier, the latter for a theory-specific feature. So, for example, where a Toulmin backing only makes sense in the context of a Toulmin analysis, the underlying AML stores the statement of the backing along with a role tag specifying the *class*: *toulmin element*: *backing*. Translation processes are then composed predominantly of a mapping to and from "vanilla", generic AML, with a small part of theory specific intertranslation. This approach avoids the combinatorial problems whilst still supporting theory-specific features, both in and between theories.

5. Conclusions

We have presented mechanisms for translation between the standard treatment of box-and-arrow diagrams and the Toulmin model of analysis. The translation presented is consistent, deterministic and requires no user intervention. Information loss during translation is limited to those features that are only expressible in one theory or the other; such information is preserved at a deep structure and is recoverable. Such intertranslation makes possible a single piece of software that can support teaching, diagramming, storage and manipulation of argument structures in the two frameworks. But more than that, it offers a mechanism for interchange and reuse between communities. As an example, *Araucaria* has been used to develop a corpus of natural argument, comprising over 500 analysed extracts from a wide variety of sources in several languages from around the world. The work was carried out as part of a project investigating argumentation schemes, and as a result adopted the standard treatment for its primary method of analysis. That corpus is currently being employed in a variety of research work, but is also available for teaching. With mechanisms for translation, a rich resource is now available not only to educators who use the argumentation scheme techniques, but also those who use the Toulmin model in the classroom. Future empirical work is required to explore how well automatically translated material fits the needs of professionals and educators.

This work has shown that theory neutrality in designing tools is a goal that can be met, and that solutions that do so are implementable in software that is still simple enough to use that students working within any of the supported theoretical frameworks find little challenge in working with the software. Inevitably, as with any translation, some of the subtlety or nuances that can be captured in one theory or language might be rendered imperfectly in another. But what has been demonstrated here is that such translation can be carried out with a sound theoretical basis to support both scholarly research into argument structure and practical application of such research in specific domains.

Acknowledgments

The authors would like to acknowledge the support of the Leverhulme Trust under grant F/00 143/C for a part of this work, and to thank both Doug Walton for ongoing discussions on closely related topics, and the anonymous reviewers who provided very detailed and helpful comments.

References

Beardsley, M.C. (1950) *Practical Logic*, Prentice Hall.

Bex, F., Prakken, H., Reed, C. & Walton, D. (2003) "Towards a formal account of Reasoning about Evidence: Argument Schemes and Generalisations", *Artificial Intelligence & Law*, 11 (2-3): 125-165.

Bench-Capon, T.J.M. (2003) "Persuasion in Practical Argument Using Value Based Argumentation Frameworks" *Journal of Logic and Computation* 13(3): 429-448.

Bench-Capon, T.J.M. (1998) "Specification and implementation of Toulmin dialogue game". In J. C. Hage *et al*., editor, *Legal Knowledge Based Systems*: 5--20. GNI, 1998.

Chryssafidou, E., Sharples, M. (2002) "Computer-supported planning of essay argument structure" in *Proceedings of the 5th International Conference of Argumentation*, Sic Sat, Amsterdam.

Dung, P.M. (1995) "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games" *Artificial Intelligence* 77(2): 321-357.

Fox, J. & Das, S. (1996) "A Unified Framework for Hypothetical and Practical Reasoning" in *Proceedings of the 1st International Conference on Formal and Applied Practical Reasoning (FAPR'96)* Springer.

Freeman, J.B. (1991) *Dialectics and the Macrostructure of Argument*, Foris.

Groarke, L. (1999) “Deductivism within pragma-dialectics” *Argumentation* 13: 1-16.

Hoaglund, J. (1999) *Critical Thinking* (3rd Edition), Vale Press.

Hitchcock, D. (2003) “Toulmin's warrants” in *Proceedings of the 5th International Conference of Argumentation*, Sic Sat, Amsterdam.

Katzav, J. and Reed, C. (2004) “On Argumentation Schemes and the Natural Classification of Arguments” *Argumentation* 18 (2): 239-259.

Kirschner, P.A., Buckingham-Shum, S.J. & Carr, C.S. (2003) *Visualizing Argument,* Springer.

Parsons, T. (1996) “What is Argument?” *Journal of Philosophy* 93: 164-85.

Perleman, C. and Olbrechts-Tyteca, L. (1969) *The New Rhetoric*, Notre Dame Press.

Pollock, J.L. (1995) *Cognitive Carpentry*, MIT Press.

Reed, C. & Norman, T.J. (2003) *Argumentation Machines*, Kluwer.

Reed, C. & Rowe, G.W.A. (2001) "Araucaria: Software for Puzzles in Argument Diagramming and XML" in *Department of Applied Computing, University of Dundee Technical Report*.

Reed, C. and Rowe, G.W.A. (2004, to appear) "Araucaria: Software for Argument Analysis, Diagramming and Representation", *International Journal of AI Tools*, 14 (3-4).

Rittel, H.W. & Webber, M.M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences* 4: 155-169.

Toulmin, S.E. (1958) *The Uses of Argument*, CUP.

Van Gelder. T.J. (2003) “Enhancing deliberation through computer-supported argument visualization” in (Kirschner *et al*., 2003)

Verheij, B. (2003) “Artificial Argument Assistants for Defeasible Argumentation” *Artificial Intelligence* 150(1-2): 291-324.

Walton, D.N. (1997) *Argumentation Schemes for Presumptive Reasoning* LEA.

Wigmore, J.H. (1931) *The Principles of Judicial Proof* (2nd Edition) Little, Brown & Co.

Wreen, M.J. (1998) “A few remarks on the individuation of arguments” in *Proceedings of the 5th International Conference of Argumentation,* Sic Sat, Amsterdam.

1 All the figures in this paper have been produced using *Araucaria* which is available for download from `http://araucaria.computing.dundee.ac.uk`

2 It is an unfortunate feature of terminology that Pollock contrasts these undercutters with direct counters that he calls rebutters. We re-emphasise that here we use the term rebuttal strictly in its Toulminian sense.