Pushing Toulmin Too Far:
Learning From an Argument Representation Scheme

Susan E. Newman[1] and Catherine C. Marshall

Xerox Palo Alto Research Center
3333 Coyote Hill Road,
Palo Alto, CA 94034

Abstract

Many researchers have proposed representational schemes to capture complex reasoned discourses. In this paper, we use our experiences with argument representation to examine some of the issues affecting the design of these representational schemes. Our discussions focus on how well a particular scheme, Toulmin structures, maps into the domain of argumentative discourse and captures and highlights various phenomena we consider central to argumentation. We then use this analysis to explore several complementary representational schemes. Finally, we discuss some relatively unexplored factors that influence the usability of these schemes.

1. Introduction

Human-computer interaction researchers have become increasingly interested in embedding task-specific representational schemes in applications designed to support idea-processing work such authoring, analysis, and design. In particular, a number of researchers are working to develop computational support for the construction, evaluation, and retrieval of complex reasoned discourses (Streitz et al., 1989; Conklin and Begeman, 1988; Neuwirth and Kaufer, 1989; MacLean, et al., 1989; Smolensky, et al., 1988; Storrs, 1989; Lowe, 1986; Russell et al., 1990; Fischer et al., 1989; and Lee, 1990a, 1990b). All of these systems incorporate representational frameworks for structuring the discourse in standardized, computationally manageable forms -- that is, all incorporate argument representation schemes. The work has produced interesting proposals for argumentation environments and related computational functionality, as well as some much-needed analysis of the nature and role of argumentation in particular work settings. However, there has been relatively little analysis of argument representation schemes per se, or reflection on trade-offs involved in representation design.

Over the last four years we have been engaged in developing and evaluating argument representation schemes for use in fields such as policy, law, design, and science. Our purpose here is both to report on these efforts and to use our experiences to articulate some of the issues that surround the design of representational schemes for use in idea-processing work. We address these

---

1. also at University of California at Berkeley, Education in Math, Science, and Technology

issues within the frame of reference provided by the domain of argumentation. However, we believe that many of them are relevant to development of representation-based tools in other domains.

To frame subsequent discussion, we note that computer-mediated representational schemes participate in three distinct (though in practice interrelated) relations: (1) a representational relation to the represented domain -- in this case, argumentative discourse; (2) a relation to the cognitive, perceptual, and activity systems of the users, in the context of specific argumentation tasks such as analysis or formulation; and (3) a relation to the system's computational capabilities, such as search and retrieval, automatic inferencing, constraint management, and so on. Each of these relations establishes a "design dimension" along which to evaluate representational schemes: (1) concerns the scheme's *representational fit*; (2) concerns its *usability*; and (3) concerns its *computational power*. We focus in this paper primarily on issues of representational fit, reserving a few words about usability for the conclusion.

The paper is organized as follows: In Section 1.1, we explain the goals and assumptions behind our efforts to develop argument representation schemes for use in idea-processing tools. In particular our thinking has been influenced by a general analysis of sources of structure in argumentative discourse. In Section 1.1.2 and 1.1.3, we present this analysis followed by a corresponding taxonomy of argument representation schemes that we have discovered in the literature and through our own work. This analysis provides a very general delineation of the space of argument representation schemes, setting the stage for the evaluation of particular alternatives and discussion of how different schemes overlap or may be combined. In Section 1.2, we briefly discuss our research strategy, focusing in particular on what we hoped to gain from an analysis-based approach to developing and evaluating argument representation schemes. With that background established, we turn to discussion of specific argument representation schemes and approaches.

Section 2 develops a space of representational requirements and possibilities for capturing argument structure. We motivate the space by analysis of the characteristics and limitations of a particular representational scheme, Toulmin structures, as applied to real, complex arguments from the legal domain. The section is organized as follows: We discuss Toulmin's theory and how we have applied it to the oral and written forms of a specific legal argument, *People v. Carney*, in Sections 2.1 and 2.2. Attempting to encode the *Carney* argument leads to a set of "first-order" extensions to the basic scheme. These extensions, presented in Section 2.3, are of two types: one shows how the categories provided by Toulmin's theory map onto argumentative discourse; the second allows the representation of different kinds of logical connectivity within the Toulmin

framework. In Section 2.4 we evaluate the representational fit of the extended Toulmin scheme, as applied to the *Carney* arguments. Specifically we consider the characteristics of coverage, perspicuity, encodability, and comprehensibility. Discussion of these characteristics helps motivate proposals for "second-order" extensions, discussed in Section 2.5. These are classes of complementary representational schemes that highlight different aspects of an argument's reasoning structure, provide organizational infrastructure or problem decomposition, and/or enable comparison between different lines of argument.

While this paper focuses on issues of representational fit, the other "design dimensions" mentioned above are clearly also important. Our experiences using a variety of argument representation schemes, in particular our struggles in "fitting" representation schemes to arguments together with our observations of difficulties encountered by other researchers, suggest particular perspectives on questions of usability. In Section 3, we briefly outline these perspectives and pose some specific questions that we think need to be addressed in future research. The trade-offs between representations that users find appropriate, and representations that facilitate computational analysis has long been recognized as a tension in the development of semi-formal representations. We leave discussion of this tension and proposals for addressing it for later work.

1.1. Goals and Assumptions. Representational approaches to supporting argumentation rest on a few key assumptions. First is the idea that argumentation, like other types of discourse, is characterized by structures that can be described independently of specific content. We believe that this assumption is borne out by our and others' success in encoding arguments in generalized representational schemes. In the next section, we discuss our ideas about the types of structure underlying argumentative discourse and show how they define a taxonomy of argument representation schemes.

Secondly, it seems clear that even though surface features of the discourse (logical connectives, discourse markers) may provide clues for recognizing these structures, the structures are largely implicit. Consider, for example, how you distinguish hypotheses or claims from evidence or how you identify assumptions in a standard scholarly text. Such structure markers as do appear are generally ambiguous in scope; as readers we often depend upon our full understanding of the text's meaning to determine what elements are linked by markers such as "because" or "however". In oral discourse our ability to follow an argument's structure depends much on shared knowledge and the surrounding discourse context.

Finally, current discourse processing theories lead us to expect that the implicitness of argument structures contributes to the difficulty of comprehending and producing complex arguments.

Discourse comprehension theorists (e.g., van Dijk and Kintsch, 1983) maintain that much of the work in processing texts or oral discourse has to do with recognizing and using various kinds of discourse structure in support of building up models both of the discourse and of the situation that the discourse is about (see Newman, 1987 for a specific proposal for how knowledge of certain kinds of argument structure guides comprehension of argumentative text). Conversely, there is some experimental evidence that arguments presented in a structured form are better understood and remembered than those presented in a standard linear format (Schum and Martin, 1985). On the production side, writing researchers (Scardamalia & Bereiter, 1987; Hayes and Flower, 1980) have pointed to the difficulty of coordinating local and global structures in the writing process -- being sure that larger argument structures are complete and well-ordered while struggling to maintain coherence at the local level. Similarly, proponents of representation-based meeting tools have pointed to the difficulty of keeping track of the structure of a complex argument as it develops. (See Stefik, et al., 1987, on Colab and Argnoter.)

These assumptions have led us, among others, to explore the possibility that representational tools that reify implicit structures of arguments might aid users in constructing and comprehending them. However, it seems useful to first consider the types of structure that we might represent. The following section provides an analysis of types of structure implicit in argumentative discourse and a corresponding taxonomy of argument representation schemes.

1.1.2. A framework for characterizing argument structure. According to our analysis, we can distinguish an argument's immediate presentational form, or "surface structure," from its implicit "deep" structure.[1] The latter can be analyzed in terms of (1) functional structure; (2) ideational or content structure; and (3) reasoning and/or presentational strategies that account for the "logical" organization of the deep structure or for its presentation in linear form.

The functional structure is the organization of the arguments elements in terms of functional or relationally defined roles; thus a given element is understood as "evidence" in so far as it "supports" a "claim." While it is possible to define many kinds of functional roles for elements of a discourse, we have focused on those that help define its character as an argument. Thus, characterization of

1. We borrow Chomsky's distinction between surface and deep structures to emphasize the point that important regularities in an argument are implicit, although they may be signalled by surface features of a discourse. We do not thereby imply that these "deep" structures inhere either in the discourse or in the head of the argument processor, although they may be (at least partially and fleetingly) constructed by the reader/hearer/producer in the process of comprehending or producing the discourse. We only mean to say that the discourse can be analyzed in these terms, assuming that a large number of discourse conventions and background knowledge are shared by the author and the analyst. The issue of how easily and under what conditions a comparable on-line or post-hoc analysis can be carried out by users armed with an argument representation scheme is a key question for the usability of representation-based idea-processing tools.

arguments in terms of functional structure yields an argument model.

The ideational or content structure derives from the causal or other (e.g., classificatory or taxonomic) dependencies that characterize an author's or community's knowledge of the subject matter. This knowledge may be structured in terms ranging from very general, widely applicable models -- e.g., a specification of the characteristics of a good system design or the structure of rational action -- to very specific situation models that codify the reasoner's expectations in a particular case.

The presentational structure is the final design of the sentences and paragraphs (or utterances), ranging from the details of the use of lexical and grammatical resources to the ordering of the sentences within paragraphs and paragraphs within sections, and so on. For on-line presentation of structured discourses, this presentational structure can also be conceived in terms of how it orders and represents the content or functional structures.

In addition to these structural characterizations, we can describe the strategic organization of an argument. We can arrive a such a description through analysis of presentational strategies, which organize the structure for presentation, and reasoning strategies, which determine the progress of an arguments reasoning. Both reasoning and presentational strategies can be defined over functional or content categories.

1.2.1. A taxonomy of argument representation schemes. Analysis of arguments in these terms yields three general types of schemas or models that are candidates for reification in representational schemes:

(1) <u>schemas based on functional categories and relations, or argument models</u>. Most argument representation schemes we have reviewed are based on functional schemas. A number of researchers are exploring the usefulness of Toulmin structures (Storrs, 1989; Taylor et al., 1989; Streitz, 1989; and Lowe, 1985). van Lehn's (1985) hypothesis-evidence-argument structures are similar in structure to Toulmin's claim-data-warrant structures, but incorporate expectations about the particular types of claims (hypotheses) that are being supported. Smolensky, Bell, Fox, King, and Lewis (1988) are also developing a detailed set of functional representation schemes based on both rhetorical and logical analysis of an arguments functional roles.

Both Toulmin structures and hypothesis-evidence structures are based on fairly detailed models of argumentative reasoning. Conklin and Begeman's (1988) adaptation of IBIS structures (Rittel and Kunz, 1970) and MacLean, et al.'s (1989) Questions, Positions, and Arguments seem to be based on models of argumentation as an activity in service of problem solving, involving both

decomposition of the problem into issues and positions, and opportunities for supporting or contesting positions. The structure of argumentative reasoning per se is not addressed. Lee's (1990a, 1990b) Decision Representation Language combines facilities for representing the structure of the "decision problem" with a simplified model of argumentative reasoning, consisting of claims, qualifications, and support-denies relations.

(2) schemas based on content categories and relations, or subject matter models. We know of only a few cases in which researchers have developed argument representation schemes around models of the subject matter. Marshall (1986) has developed a representational framework for foreign policy decision-making and explanation that organizes reasoning around categories derived from a normative model of rational action (Allison, 1971). She also has experimented with supporting reasoning from taxonomic models of complex historical events, such as revolution. In addition, some researchers are experimenting with representing subject matter knowledge in traditional AI knowledge representation frameworks so that this knowledge can be accessed and used in building structured arguments using functional schemes (see for example Fischer et al. 1989 and Taylor, et al., 1989).

(3) hybrid schemas based on analysis of reasoning or rhetorical goals and strategies, formulated over content or functional categories. Again we know of relatively few examples of this approach to representing argument structure. Argument Scripts (Newman, 1986) were an attempt to specify micro-strategies for developing and presenting an argument as a possible basis for an argument outlining tool. Argument Scripts are defined over both content and functional categories. In addition, in the course of our use of Toulmin structures, we experimented informally with representing the reasoning strategies employed in the *Carney* oral argument, as a way of showing how the reasoning developed and was resolved.

1.3. A Note on Research Strategy. Given that these different kinds of representational schemes are possible, we need a better understanding of how they function in representing arguments, in particular how well particular schemes fit the domain of argumentative discourse. To address this question, we have in general pursued an analysis-based approach. This approach involves adopting or developing an argument or subject matter model; using it as a basis for a representational scheme; and applying the scheme to realistic arguments to test its representational adequacy. We will describe how this process worked with the development and evaluation of Toulmin structures, a detailed functional scheme oriented toward argumentative reasoning, in subsequent sections. The key point here is the focus on characteristics of particular schemes in relative isolation from both tool use and argumentation as an activity. Although we did develop some scheme-specific tools and representational embodiments in NoteCards (in particular,

Marshall, 1987, 1989; Newman, 1986), most of our effort has gone into analyzing and extending particular schemes as needed to capture arguments, seeing how different schemes fit together, and experimenting with the complex networks that result from the analysis.

We see both advantages and disadvantages to this sort of decomposition of research on representational tools. Obviously, because it concentrates on the relation between representational schemes and argument products, this strategy tells us little about how use of standardized representational schemes is incorporated into the diverse activities involved in argumentation. On the other hand, the analytic approach does help separate problems that might arise from the idiosyncrasies of a particular interface or from the normal management of argumentation processes from problems with the representational scheme per se. As a result, it has the potential to yield a more principled, better rationalized, and better understood representational system. In particular, it contributes to the articulation of encoding rules for relating domain information to the categories of the representational scheme. In addition, as we have employed it, an analytic approach promotes the development of a more general schematization of argument structure; as one struggles with problems of coercing existing arguments into a single representational framework, one is encouraged to notice and articulate aspects of the argument that do not fit, which in turn leads to recognizing new structural dimensions and representational possibilities.

In addition, although we have not addressed the question of representation and tool use head-on, focusing on the representational fit between schemes and arguments has contributed to our understanding of relatively unappreciated aspects of usability, i.e., those that are tied to the scheme's representational characteristics. We discuss some of these issues in the concluding section.

2. Argument Representation Schemes: Development, Evaluation and Analysis

2.1. Toulmin's Theory of Argument. Our most extensive exploration of argument structure and its representation has been based on the framework developed by Stephen Toulmin in *The Uses of Argument* (1958). Toulmin's goal in *Uses* was to understand why formal logics developed in philosophy provided so little insight into human reason and to develop an alternative notion of logical form based on the analysis of reasoning practice. Taking the view that argumentation is a primary site of practical human reasoning, he proposed a scheme for analyzing the logical microstructure of everyday arguments, contrasting this scheme and the ideas behind it with philosophy's analysis of the syllogism.

2.1.1. The Logical Microstructure of Arguments. Toulmin proposed that the logical micro-structure of human argumentation and reasoning consists, implicitly or explicitly, of six categories.

According to Toulmin, a *datum* is some fact or observation about the situation under discussion. The datum functions as the ground or basis for making a *claim*, i.e., some further, potentially controversial, observation, prediction, or characterization. The relation between the two is characterized by a rule of inference, a *warrant*, that serves to link the information set forth in the datum and claim. Data-claim-warrant (D-C-W) structures constitute the inferential core of the argument.

Thus far, the system is not unlike the familiar syllogism, with a rule of inference linking the terms of two propositions. However, it is different in two general ways: First the terms linked by the warrant need not be analytically related to the predicates expressed in a datum and claim. That is, taking the predicates of D and C as instantiating the categories or descriptions represented in W may be a substantial leap in itself, depending on the knowledge and beliefs of the interlocutors. Second, Toulmin recognized that our ability to predict, decide, and so on is not generally based on infallible knowledge leading deductively to certain conclusions. To capture different aspects of the contingent nature of our reasoning, he added three other argument elements.

First, he noted that arguments often contain *qualifiers*, phrases or terms that signal our epistemological stance toward our claims or register our degree of certainty. Second, he noted that "real" or substantive arguments, unlike the analytic arguments of standard logic, allow for exceptions. *Rebuttals* are Toulmin's category for explicitly registering the conditions under which the claim can be taken as true by oneself or an interlocutor. As Toulmin realized in naming this category, these conditions provide a direct path for development of counterarguments.

Third and most important, he realized that warrants, i.e., our knowledge, beliefs, and assumptions about relationships among "facts" in the domain at hand, come from somewhere -- our past experience, case law, the dominant theory of the field. The *backing* of an argument is some knowledge structure from which we derive the warrant as "a practical moral" for reasoning about the domain and which serves to justify the warrant. Note that the sources of warrants vary depending upon the domain. For example, one common type of backing in law is the set of relevant previous cases or pieces of legislation. In science, one may back arguments by reference to theory, by reference to constraints on experimental practice, and so on.

Figure 1 lays out the diagrammatic structure that Toulmin developed to symbolize the relationship among these argument categories and instantiates it with an example from Toulmin. The words in italics are simply logical connectors that allow one to string together the elements of the argument into a comprehensible (if long) sentence that preserves the logical sense of the argument.

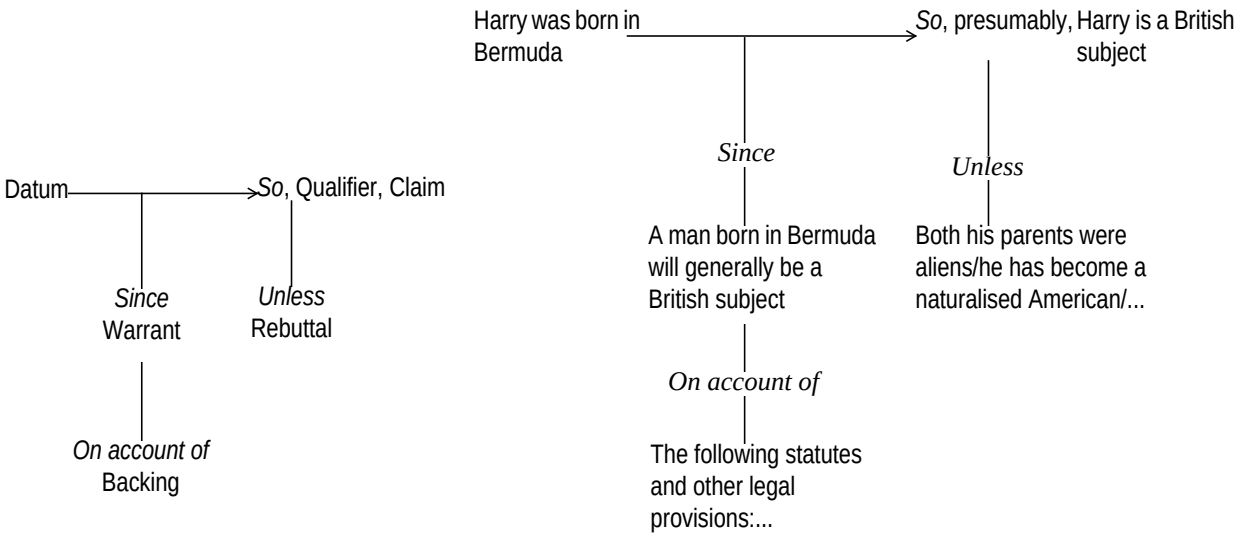

Figure 1. Toulmin structure and its instantiation with an example

Toulmin's work has been widely influential in the various fields that touch on argumentation. His initial contribution was to the philosophy of logic and epistemology, where his analysis of argument enabled him to address issues such as the nature of domain dependencies and invariances in argument and reasoning, the relation between deductive and inductive reasoning, and the relation between analytic and substantive reasoning. Toulmin's work has also influenced the field of rhetoric, where theorists have attempted to show how a classical taxonomy of argument types could be expressed within his framework (Brockreide and Ehringer, 1960). In addition, Toulmin's scheme has been used in cognitive science and education as an analysis tool for characterizing differences in expert-novice reasoning (Voss, 1983) and for use in evaluating students' argumentative writing (Hillocks, 1987). And finally, Toulmin himself has developed a textbook which uses his framework to teach reasoning skills to young adults through the process of producing and evaluating arguments (Toulmin et al., 1984).

As a starting point for argument and representation analysis, Toulmin structures have a number of advantages. Toulmin structures provide an intuitively plausible set of categories and relations for representing the logical structure of arguments organized in a distinctive graphical layout. Perhaps for these reasons, the scheme has appealed to a number of researchers interested in argument representation tools. For example, David Lowe adapted it for SYNVIEW, a community knowledge structuring tool (1985, 1986); Taylor, et al. (1989) are using Toulmin structures as the representational substrate for constructing detailed argumentation to support recommendations for changes and initiatives in social service policy, Streitz and colleagues (1989) propose it as

representational basis for the activity of structuring the logic of an argument independently of its rhetorical organization.

In spite of its influence, however, little has been done to assess the representational or theoretical adequacy of Toulmin's framework or to understand how argument structure as described by Toulmin is actually encoded in argumentative discourses. As a purported general theory of an argument's logical structure, it provides an interesting representational test case. How well does it work in representing real arguments? What are its limitations? What is entailed in encoding arguments in Toulmin structures? Does it provide a useful set of distinctions for understanding argumentation more generally? Our work with Toulmin aimed at a detailed understanding of these issues.

2.2. When the Warrant Hits the Road. To get at these issues, we undertook several related argument analyses in the Toulmin framework. We decided to focus on legal arguments, partly because legal argumentation was Toulmin's own exemplar of a highly developed system of practical reasoning and partly because of characteristics intrinsic to legal reasoning. In particular, argumentation in law is both crucially important to decision making and community memory, and has been shown to encompass case-, rule-, and model-based reasoning. Since we were specifically interested in understanding how Toulmin structures handle argumentative reasoning, it seemed appropriate to analyze a domain rich in reasoning paradigms (Skalak, 1989; Ashley and Rissland, 1985). In addition, legal decision making has characteristics of both design problem solving, insofar as it entails constructing policy rules for remote implementation, and analysis, insofar as it involves developing or applying descriptive models to situational facts.

***2.2.1. The case* People v. Carney.** We applied the scheme to arguments from two cases, *Tennessee v. Garner* (see Marshall [1989] for a discussion of the Toulmin encoding of this case and of issues surrounding the on-line presentation of hypertext structures) and *People v. Carney*. *Garner* is a U.S. Supreme Court case concerned with the constitutionality of using deadly force to stop a fleeing felon. *Carney* is a California Supreme Court case testing the constitutionality of the warrantless search of the defendant's mobile home. The main issue is whether the search could be justified by the so-called "automobile exception" to the protection against unreasonable search and seizure established by both the U.S. and California Constitutions. Although both cases contributed to our understanding and refinement of Toulmin structures, our examples will be drawn primarily from the latter.

The arguments in *People v. Carney* divide along two basic lines of thought. One says that because a motor home has an inherent quality of mobility, it is subject to the automobile exception. The

focus on a vehicle's mobility is justified in previous rulings by the need to prevent rapid removal of evidence from a state's legal jurisdiction; preservation of evidence in turn serves to promote effective law enforcement. The second line of thought says that because a motor home can function as a home, it has a reasonable expectation of privacy and thus cannot be searched without prior judicial approval. The protection of citizens' privacy in the home is viewed as a fundamental social value and is guaranteed by a long line of interpretations of the Fourth Amendment.

Because a mobile home can function both as a vehicle and as a home, a central issue concerns how it should be classified for the purpose of deciding the case. Thus many arguments in the case are concerned with developing and justifying one method of classification and refuting another. Classification arguments are important in law (and many policy and scientific arguments as well) because decision rules are inevitably defined over "open-textured" (Rissland, 1985) terms; much of the argumentation task involves relating particular situations to unavoidably vague generalizations. The winning classification method and its justification has the status of a new legal rule (or warrant-backing) that affects subsequent cases; as such it must both anticipate how future cases will be affected and avoid violating precedents established by previous cases, particularly cases ruled by a higher court. Avoiding violation of precedents may be accomplished either by using the same rule to decide the current case (with consequent argumentation aimed at relating the rule's open-textured terms to the current situation) or by defining a new rule which would lead to the same decision, but on new grounds. The first strategy entails articulating the conditions and rules governing previous decisions, particularly boundary cases, and relating those conditions to the current case. The second strategy entails a sort of paradigm shift, where the grounds for and implications of previous decisions are reconceptualized.

2.2.2. Written v. oral argument. We used Toulmin structures to encode written and oral arguments from the *Carney* case, both to provide a broader test of Toulmin and to increase our understanding of the representational differences between the two discourse types.

In general, representational differences derive from differences in the two discourse situations and hence in their argument goals and methods. In particular, the oral argument is concerned with *arriving at* an acceptable decision rule (warrant) that will meet constraints on legal rulings. This goal is pursued through a highly collaborative question-and-answer process in which warrants are proposed and subjected to hypothetical test cases leading to refinement or refutation. The written argument, on the other hand, aims at presenting a compelling reasoned argument that will serve both to justify the ruling and, by laying out the court's reasoning, to establish the conditions for future interpretation of the decisive warrant. One challenge to Toulmin structures was whether they could capture the logical structures of both types of argumentative discourse. Another issue

was whether different aspects of oral and written arguments would be lost or hard to capture in Toulmin. Some of these differences will emerge in subsequent discussion.

2.2.3. Goals and methods of analysis. In applying Toulmin's scheme to the *Carney* arguments, we adopted three encoding goals. First, to explore the scheme's representational limits, we tried to capture as many of the relationships between microarguments as possible within the Toulmin framework. Second, we tried to make semantic relations among argument elements transparent -- that is we tried to minimize the inferential leaps required to understand, for example, how warrants are related to data and claims, or rebuttals to particular arguments. Although this latter practice involved reifying much implicit structure -- arguably more than most users would wish or need to represent -- we expected it to help us understand more about how argumentative reasoning works, as well as about how it is related to linear discourse. Third, we tried to articulate problems and constraints in encoding, along with a characterization of the sorts of statements filling the Toulmin slots. This we hoped would help us understand what arguments are made of. Together these efforts led us to develop a set of first-order representational extensions, both to accommodate Toulmin's categories to the range of argumentation we encountered and to capture the argument's logical connectivity. We discuss these extensions in the next section. In addition, careful analysis here helped us in the evaluation of Toulmin's encodability.

Once the arguments were encoded in the extended Toulmin representational system, we performed two sorts of analyses. One strategy, applied to both written and oral arguments, was to encode as much of the argument as possible in extended Toulmin structures and then to analyze the resulting network for intelligibility. In particular we were concerned about the availability to prospective users of important reasoning strategies or patterns, relationships among lines of argument, and unfolding of the argument in time.

The second, complementary analysis, based on encoding of the written text, involved a detailed comparison of the Toulmin-based representation with the structuring resources and methods of linear discourse. The strategy here was to render the text into simple declarative statements; encode as many statements as possible into Toulmin structures; and analyze the elements of the text that could not be represented in the Toulmin framework. Our reasoning was that while some of this work might be peculiar to linear text (for example, some textual information is intended to help the reader construct the argument structure from the linear stream), other aspects might provide us with clues about the sorts of information that are "washed out" of Toulmin structures but that may be needed to render the argument intelligible. In this case we were particularly interested in how the argument's logical structure was organized for presentation and in how the text manages the relation of the reader to the argument.

While we do not report on these analyses in detail here, we use the results to inform a general evaluation of Toulmin structures in terms of its representational fit to the domain of argumentative discourse. To organize the evaluation, we use four general criteria: coverage, perspicuity, encodability, and comprehensibility. We discuss each of these criteria and their application to Toulmin structures in Section 2.4. These criteria proved useful in understanding the limitations of Toulmin structures and help motivate proposals for complementary representational schemes and approaches. These schemes may be viewed either as second-order extensions or, in some cases, as alternatives to Toulmin. We discuss them in Section 2.5.

2.3. First-order extensions. In order to capture the *Carney* arguments in Toulmin structures, we developed two kinds of representational extensions.

2.3.1. Extensions to facilitate encoding of elements. The first arose from the difficulty of relating the elements as defined by Toulmin to the represented domain. These difficulties led to mutation of the elements' internal (category) structure and development of elaborated encoding rules, as well as to representational innovations. Extensions were required for four of the five major elements.[1]

Data: For Toulmin, data are prototypically singular facts about individuals (for example, Harry was born in Bermuda or Swen is a Swede). Not surprisingly, such a definition proved wholly inadequate. We required four basic extensions:

• The first minimum extension (made by Toulmin himself in his recent textbook on reasoning, 1984) allowed *generalizations or statements about classes of individuals* to function as data, as in "The primary function of automobiles is transportation." In addition to simple generalizations, it was also common to find a general statement functioning as a datum for a claim, backed up by examples or instances. While it is possible to treat the examples in turn as data supporting the generalization, supported by a warrant such as "Given x examples we can draw a general conclusion," this seems awkward and somehow misses the flavor of examples, which often function as much to explicate the generalization as to support it. Also, in some cases, such as an existence proof, the number of examples does not matter at all -- you only need one. Thus we found it desirable at times to represent examples as part of a single complex datum consisting of a generalization and instances. Figure 7 below provides an example of such a complex datum.

---

1. We generally avoided use of qualifiers since, as defined by Toulmin, these consist only of highly ambiguous and indefinite discourse cues. To the extent that qualifiers might be understood as specific conditions on the truth conditions of a claim, they are subsumed by rebuttals.

• Only infrequently did we find statements adduced as evidence for a claim that would fit an objectivist's definition of "fact." Early on it became necessary to include as data *statements of belief under a basically social model of the domain*. Compare, for example, a statement to the effect that "the primary function of motor home X is transportation" to the statement that "motor home X has wheels." Note that while the former statement is disputable, it cannot be proved or disproved by reference to any particular method. What determines the believability of the datum are shared social understandings and the goal of the argument (in this case to emphasize the similarity between motor homes and automobiles).

• In addition, these descriptive and evaluative data were often not simple categorical statements, but *statements of relative belief*. That is, predicates were applied to some entities in comparison to other entities or in comparison to other predicates. For example, one argument in the written opinion is based on a relative evaluation of the neutrality and detachment of magistrates in comparison to the police, who are seen as involved and competitive and thus *more likely* to exceed the bounds of executive authority.

• Quite often, a claim was not backed up by a single statement but required *several statements operating together*. To handle these cases we created *conjunctive data*. Figure 2 below provides an example of conjunctive data in service of a claim. Conjunctive data are employed when multiple descriptions or characterizations of a situation must be accepted simultaneously in order to support a given claim. As we shall see, conjunctive data are not sufficient to account for all cases of multiple sources of evidence supporting a claim. (Cf. J. Lee's (1990a) provision of grouping structures which allows the user to collect argument objects and specify the constraints on their relations -- for example, mutually exclusive, conjunctive, and disjunctive)

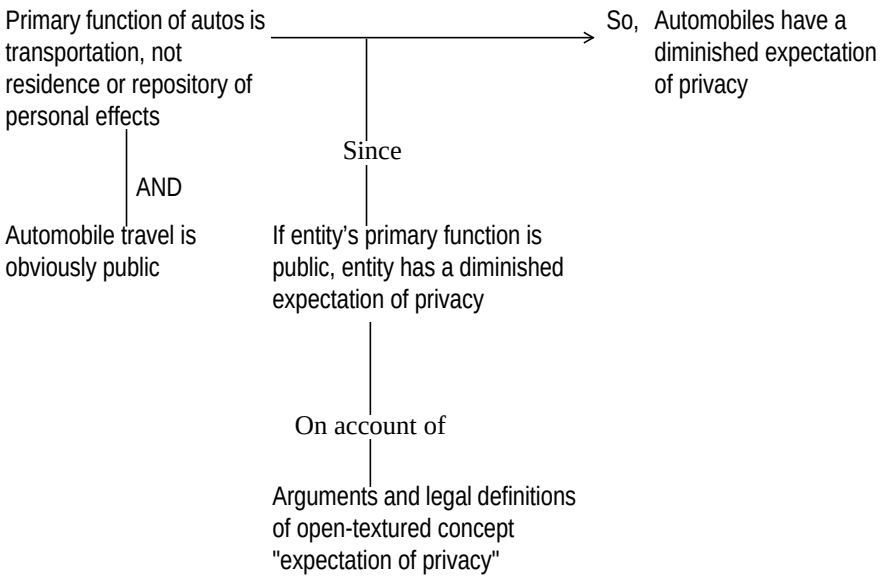

Figure 2. Example of conjunctive data

Warrants: According to Toulmin, warrants are more general than data and claims; that is, the predicates linked in a warrant subsume the predicates asserted in its datum and claim. However, warrants are often implicit; as a result, one has a great deal of freedom in deciding what generalizations to make. The resulting encoding problem is particularly difficult where there is a wide inferential gap between an explicitly stated datum and the claim it supports. Such moves can be characterized either by a warrant statement whose relation to data and claim is itself highly implicit and knowledge dependent, or by generating intermediate low-level inferences to reduce the gap. We have opted for the latter practice. This option increase the overhead of encoding, but also increases the scheme's intelligibility and analytic power.

Although expressing an argument's intermediate inferences narrows the choice of warrants, the problem of deciding what generalization to make remains. For example, one can express the warrant governing the search of cars as "mobile *entities* can be searched without a warrant" or as "mobile *vehicles* can be searched without a warrant." Since a key issue in legal (and presumably other sorts of reasoning as well) is what sorts of things are included in the categories one reasons over, decisions about category description in the warrant are consequential for future reasoning and decision making. Explicit expression of the inference rules underlying an argument is a major source of the Toulmin scheme's representational power. However, it seems likely that ongoing negotiation, interpretation, and modification of such rules will occur in the course of argument and analysis.

Backings: In Toulmin's examples in *The Uses of Argument*, backings resist questioning; they are

unassailable "facts," like the provisions of a statute or something that has been experienced directly by the arguer. However, arguments rarely rest on so firm a foundation. We found it necessary to elaborate our understanding of backings in several ways in order to develop a consistent encoding of the *Carney* arguments.

• *Citations*: As expected, citing other cases is common practice in the written opinion; often (but not always) these can be seen to function as backing for an argument. In such cases, we have encoded the case cite in the backing slot of the appropriate microargument. However, for a consistent understanding of backing, it is important to realize that the "findings" of a previous case (or, analogously, of a piece of research in scholarly or scientific argument) are rarely either unquestionable or unambiguously applicable to a new case. Instead, citation backings should be understood as a placeholder for a line of argument establishing the warrant and/or its applicability to the current case.

• *Backing arguments*: Occasionally in the written opinion, some of the underlying argument is actually made explicit; the warrant for one argument then becomes a claim in a "backing argument." (See the example in the discussion of argument hierarchies below.)

• *Implicit backings*: Except for case cites in the written argument, however, explicit statements of backing are rare. In general, legal argumentation takes place among participants with a good deal of shared background knowledge, which renders explication of the basis of an argument unnecessary. Where backings cannot be resolved into a line of argument or summarized by a case cite, we have used the backing as a location for classifying the type or general source of the warrant. In this role they provide a resource for analyzing the knowledge types underlying a complex argument. A good many of the micro-arguments in *Carney* ground out in commonsense knowledge and practices. Heuristics for classifying social objects like cars and motor homes (by function, by attribute) are particularly important.

Rebuttals: Rebuttals are the most problematic and undeveloped aspect of Toulmin structures. As we discussed, Toulmin conceptualized rebuttals as a location for enumerating the conditions on an argument's validity. However, because of our goal of precisely specifying the interconnections among arguments, this general characterization proved insufficient; it does not inform us about either the particular element(s) under attack or the nature of the challenge. As a guide to modifications of the Toulmin representation, we developed an analysis of types and functions of rebuttals, as they appear in the *Carney* arguments.

• The first is perhaps the prototype that Toulmin had in mind when he described rebuttals. This type of rebuttal encodes the *conditions under which a given warrant no longer holds*, thus defining

possible exceptions to the rule. This definition works best in well-defined argument domains where enumeration of qualifying conditions is possible. Statutes and judicial rulings often attempt to specify the limiting conditions under which a warrant is applicable. In the *Carney* argument, for example, the justices describe the "controlling principles" of the "automobile exception" to the warrant requirement, as follows: "Officers are empowered to search an automobile as long as it can be demonstrated that (1) exigent circumstances rendered the obtaining of a warrant an impossible or impracticable alternative and (2) probable cause existed for the search." Figure 3 represents these conditions as rebuttals in a Toulmin structure. Note that Toulmin depicts these rebuttals as being attached to the qualifier and thus to the relation between the datum and the claim as mediated by the warrant.

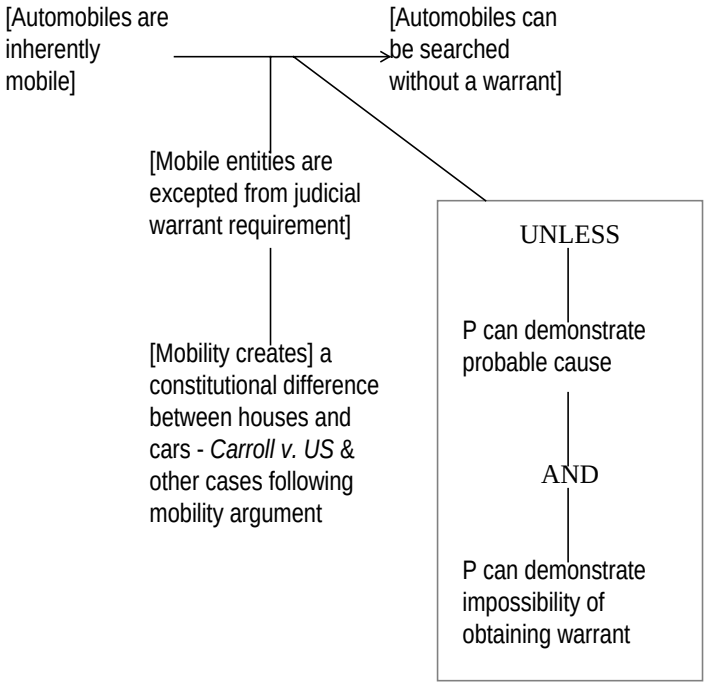

Figure 3. Rebuttals as enumerated exceptions[1]

However, objections to arguments are not restricted to the enumeration of definitive conditions-of-applicability for the warrant, or to claims that those conditions are instantiated in the current case. Rather, challenges take a number of different forms with correspondingly different representational requirements. One set of rebuttal types are defined relative to individual argument constituents.

• *Denial of datum*: A datum (for example, "The primary function of a mobile home is

1. Note when dealing with simple qualifying rebuttals, it is sometimes convenient to use *as long as* rather than express the negation.

transportation") can be denied by negation, proposing an alternative characterization ("The primary function of a mobile home is to provide a residence") and/or by adducing evidence to the contrary.

• *Counterclaim*: Similarly, an argument may also be rebutted by a counterclaim (with or without evidence). We observed two sorts of implications of counterclaims for an argument. In the first case, the counterclaim stands as an exception or limiting case to the rule of inference expressed in the warrant. For example, given the warrant "mobile entities can be searched without a judicial warrant," one might infer (as did the prosecuting attorney) that suitcases found in a moving car can be searched without a warrant. In fact, however, such searches are for the most part disallowed, even though in general the warrant stands. In the second case, the negation of the claim implies the rejection of the warrant that produced it. For example, in the oral argument, the fact that trailers cannot be searched without a judicial warrant negates the proposed warrant that if a vehicle has wheels, it is subject to the automobile exception (and thus can be searched without prior judicial approval). Which of these implications applies probably depends on conditions such as how long-standing the warrant is, what the goals of the argumentation activity are, and so on.

• *Questioning of the warrant*: If too many exceptions to the warrant arise, it may be attacked directly. For example, the written opinion argues that because there are so many exceptions to the warrant that excludes mobile entities from the Fourth Amendment requirements, mobility is not a good basis for a decision rule. See the example in Figure 7 below. One may also question the applicability of a warrant to a particular case by distinguishing the current case from the case or cases in which the rule arose. Such argument aims at *establishing* (rather than enumerating) boundary conditions for the use of a decision rule. Much of this sort of reasoning involves operationalizing the open-textured concepts in terms of which many legal rules are defined (See Rissland, 1985, for a discussion of the role of hypotheticals in testing the limits of open-textured terms in law).

• *Attack on backing*: Finally, one may attack the model underlying an argument by attacking the backing. For example, in our work on *Tennessee v. Garner*, an inference is made on the basis of a commonsense model of criminal behavior, that the threat of deadly force acts as a deterrent to escape. A justice's rebuttal questions the validity of this model, which is used as a justification for the rule that the police may use deadly force against fleeing felons.

In normal discourse, many elements of an argument, in particular, intermediate data and claims, warrants, and backings, are implicit. As a result, to identify which element of a given Toulmin structure is under attack requires the elaboration of much implicit structure. Our efforts to figure

out the rebuttal relations implied by the posing of hypotheticals or adducing of counterclaims from previous cases in the oral argument provide a good example of the difficulty. Consider the following: At one point in the discourse, the prosecuting attorney proposed to operationalize the open-textured term "mobility" in the existing decision rule ("Mobile entities can be searched without a warrant") by specifying the attributes that determine mobility: "If a vehicle has wheels -- that provides a bright line." The rejoinder "But what about trailers?" constitutes the rebuttal, but it is difficult to understand (and therefore to represent) in Toulmin terms just how the rebuttal functions.

In order to carry out this and other rebuttal analyses, we found it necessary to distinguish between: (1) the explicit statement in the discourse that signals a rebuttal; (2) the implied rebuttal statement; (3) the implied argument or counter-argument set up by the rebuttal statement; and (4) the rebuttal strategy and associated goal. In this case, the rebuttal-signalling statement simply raises the case of trailers for consideration. By assimilating trailers to the current discourse context, we get the implied argument: "(D) Trailers have wheels, therefore (C) trailers can be searched without a warrant, since (W) vehicles with wheels can be searched without a warrant." But it is known from previous cases that trailers cannot be searched without a warrant. Thus the *implied* rebuttal statement "Trailers cannot be searched without a warrant" denies the *implied* claim "Trailers can be searched without a warrant," which is, under a Toulmin analysis at least, automatically "produced" through the introduction of the hypothetical case. Note further that if we stopped here with the analysis we would still have missed the *point* of the rebuttal, or the rebuttal strategy, which is not to decide whether trailers can be searched without a warrant, but to deny or undermine the proposed warrant "If a vehicle has wheels, it can be searched without a warrant."

In the previous examples it is possible (if sometimes difficult) to relate the rebuttal to particular elements of the Toulmin structure. However, rebuttals may address issues that arise from the implicit context in which the argument takes place. In addition, the rebuttals discussed so far question the validity of particular statements or of relations between statements. But objections may be lodged on grounds other than validity. We have identified two sorts of rebuttals whose focus is outside the argument microstructure and that concern themselves with issues other than truth or falsity.

• In legal argument, and we suspect in design and policy arguments as well, warrants are proposed as decision rules intended to function effectively in a remote context of action. For example, legal rules formulated or refined in court may be intended to guide police behavior in response to situations that arise in the course of duty. *But a rule for remote action can be faulty not because of a lack of validity, but because it can't be implemented reliably.* Thus one kind of rebuttal

questions a warrant by attacking a datum on the grounds that it can't be known in the context of practical police activity. For example, a proposal to use the length of time a vehicle has been parked to decide if it is being used as a home is rebutted on the grounds that there is no practical way for the police to make the necessary discrimination. In some sense this challenge questions the appropriateness or possibility of using a proposition such as "Vehicle X has been parked for 9 months" as a datum, at least with respect to this particular warrant in the context of its use as a guide for future action.

• Finally and most importantly, rebuttals are posed because a particular argument or line of reasoning fails to meet the constraints imposed by a *conflicting set of goals and values*. The relative importance ascribed to preserving evidence and promoting effective law enforcement versus protecting privacy and maintaining limits on executive authority determines the relative weight ascribed to the mobile home's mobility and similarity to an automobile versus its function as a home. We can illustrate this point with an excerpt from the oral argument. Just previously, the prosecuting attorney has (effectively) proposed a warrant for deciding the *Carney* case: "Entities capable of quick movement should be excepted from the judicial warrant requirement." One of the justices proposes a hypothetical to test that warrant's boundary conditions:

> Q: What if the vehicle is in one of these mobile home parks and hooked up to water and electricity but still has its wheels on?
>
> A: If it still has its wheels and it still has its engines, it is capable of movement and it is capable of movement very quickly.

The prosecuting attorney's response does not directly address either the hypothetical datum proposed by the justice ("Mobile home X is parked and hooked up to water and electricity") or the implicit line of argument that it implies (that motor homes hooked up to utilities are obviously being used as homes and therefore are included in the Fourth Amendment protections). Rather, it asserts the irrelevance of those attributes to the underlying concerns about police effectiveness that motivate the "auto exception." Note that conflicting goals and values are generally accompanied by fundamentally different models of how things ought to be done, what aspects of the status quo are candidates for change, etc. In essence, conflicting goals and values lead not only to conflicting "design criteria" for decision-making (and thus to different warrants), but also to different ways of characterizing situations (and thus to different data as well).

2.3.2. Extensions to capture argument connectivity. Because Toulmin was concerned with philosophical (and later pedagogical) issues surrounding the logical microstructure of arguments, his system does not provide ready-made ways of capturing an argument's connectivity (although

note that in *Reasoning* (1984) he briefly discusses the linking of arguments through data; see argument chains below). We have developed four extensions for linking micro-arguments to capture extended reasoning. Several of these extensions are related to mutations of the internal structure of specific elements discussed above.

Argument chains: As we discussed in the section on rebuttals, data can be challenged. This often leads to producing a second datum in support of the first and to a recursion of argument structure, with a consequent transformation of the role of the first datum -- from datum to claim -- with respect to the new argument. Argument chains are also required to build up any substantial line of argument. Figure 4 displays interconnection of arguments by chaining.

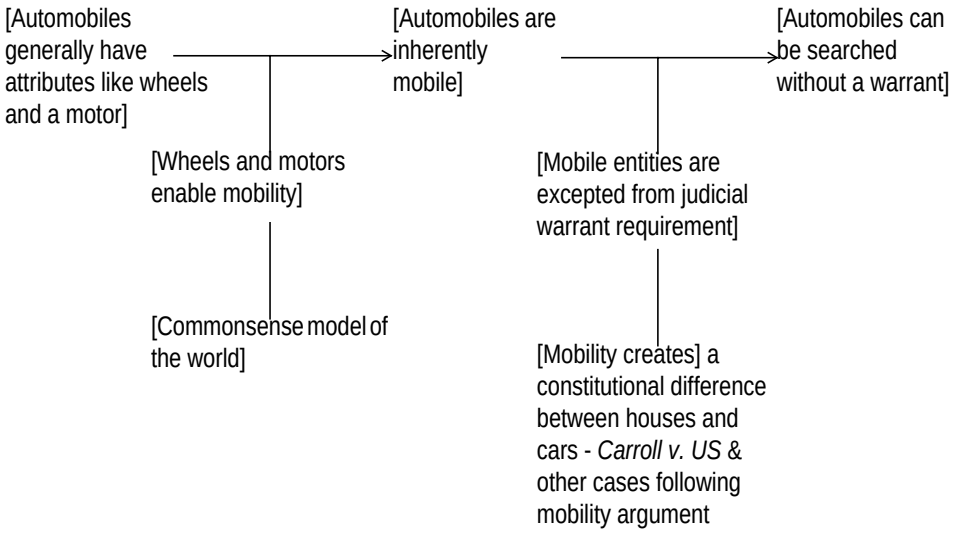

Figure 4. Argument chains

Argument hierarchies: As we suggested earlier, backing may be seen as a place-holder for the shared commonsense and domain-specific models, or simply shared practices, that give an argument's warrant its credibility. Quite often, because it is shared, this background information need not (sometimes cannot) be articulated; arguments "ground out" in social consensus. However, either for rhetorical reasons or if the warrant is questioned, it may become necessary to articulate the backing. In most cases, this means further argumentation and thus recursion of the basic argument structure. Since the role of the backing is to establish the credibility of the warrant, this recursion results in a transformation of the warrant's functional role with respect to the new argument -- it becomes a claim. Figure 5 shows a schematic representation instantiated with an example from the written argument.

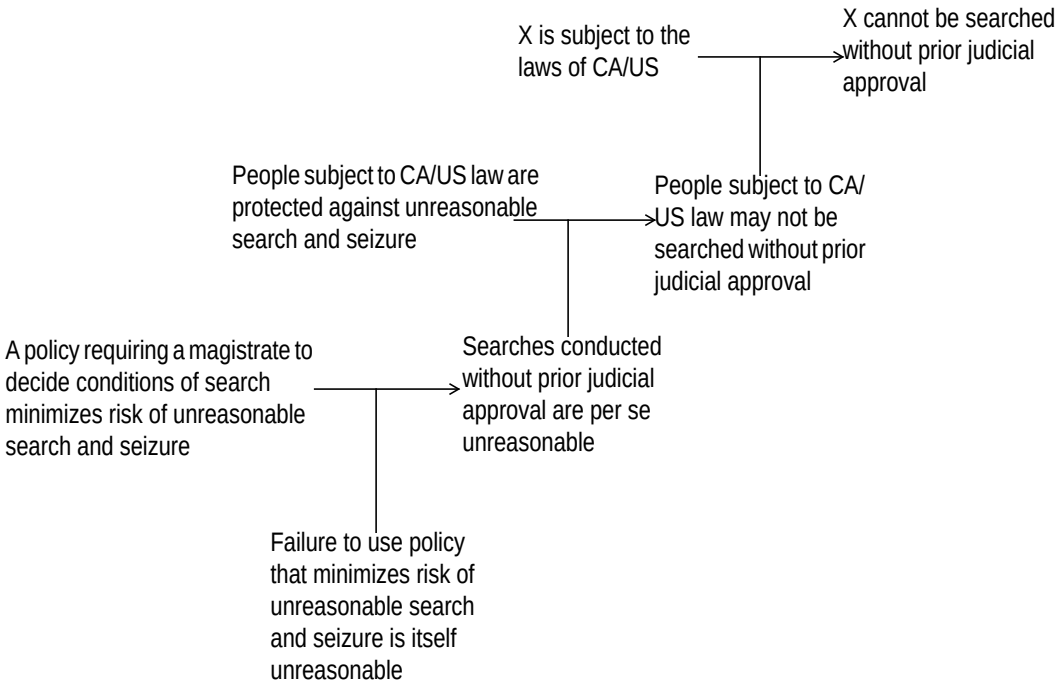

Figure 5. An example of a hierarchical argument structure

Confluence arguments: In addition to conjunctive data, where two or more conditions must hold in support of a claim, we found it necessary to represent what we call confluence arguments. In such arguments, which are particularly common in the written argument, independent lines of argument or data are brought together to support a single claim. We developed two different representational possibilities: The first decomposes the truth conditions under which a claim holds by specifying separate warrants connecting each datum to the claim. The second employs a single generic summary warrant that pronounces the claim true "on the weight of the evidence." Figure 6 demonstrates these two possibilities.

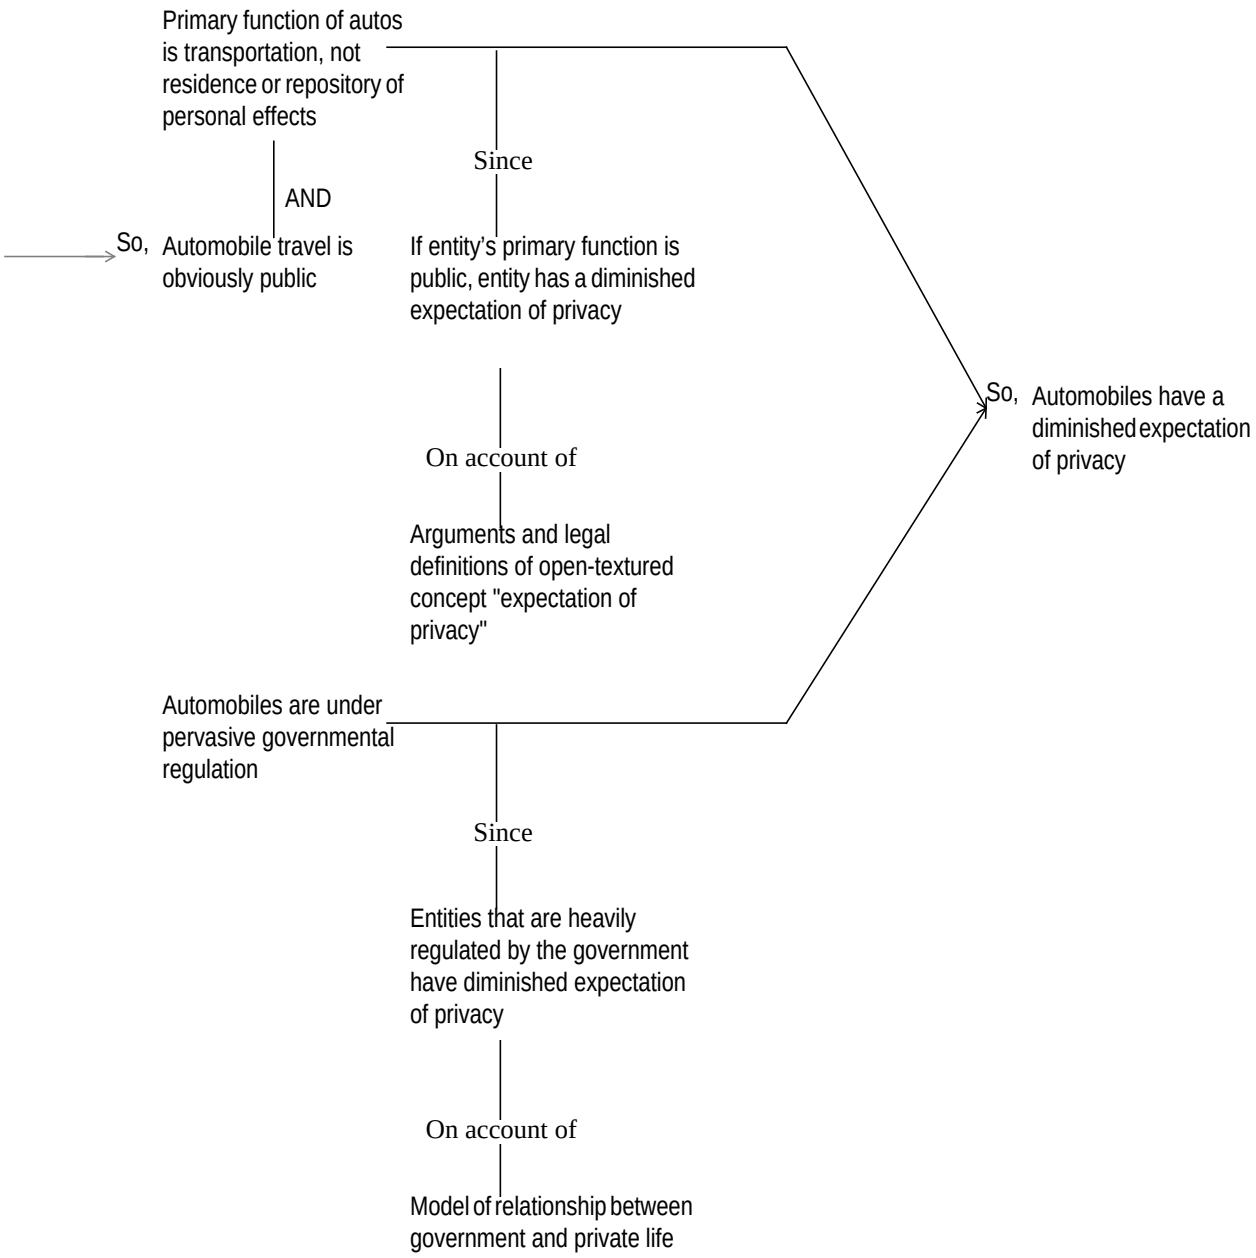

Figure 6. Example of a claim supported by two separate argument structures

It seems likely that evidence in confluence arguments is often brought together in a situation-specific, ad hoc way. However, it is worth noting the possibility that analysis of the *types* of data and warrants coordinated in a confluence argument may allow us to uncover recurrent argument strategies organized by underlying domain models or argument-structuring schemas. Although we have not evaluated the generality of the reasoning patterns found in *Carney* (for example, classification by function, by actual use, by particular attributes), we have, in other work, discovered what seem to be very general strategies for producing policy arguments (see Newman,

1986, for a discussion of Argument Scripts). To argue against a given policy, for example, one generally attempts to show that the policy leads to undesirable consequences, that it is ineffective in achieving its goals, that it is based on an incorrect formulation of the problem, and that there are alternatives. We will discuss later how general argument strategies such as these, or the model of the domain that underlies them, might provide the basis for more domain-specific or semantic ways of representing argument structure.

• Connection by rebuttal: Just as rebuttals are problematic and undeveloped in Toulmin structures, so are its facilities for representing the relation between conflicting positions. We believe that the problem of expressing these conflictive interrelations is a general one for argument representation schemes[1] and arises, at least in part, from an inadequate understanding of the sources, and expression, of conflict in argumentation. Analyses of rebuttal types such as the one we conducted around the Toulmin framework are an important step toward understanding the requirements and possibilities for representing conflict.

In the previous section, we discussed how rebuttals can arise as local challenges to particular argument constituents, as well as more global challenges motivated by implicit goals and alternative models or world views. Global challenges do not lend themselves to representation through local connectivity. Indeed the inability to capture interactions between globally opposed lines of arguments is one of the limitations of Toulmin structures that motivate the second-order extensions and alternative representational schemes that we describe in Section 2.5. Representing localized rebuttal relations, on the other hand, seems fairly straightforward: a rebuttal attaches to the particular element under attack. Homer-Dixon and Karapin (1989) propose several extensions of Toulmin that accomplish this sort of connectivity; according to their proposal, rebuttals can point to any other element (indicated graphically by arrows).

With a few exceptions, rebuttals in the *Carney* arguments are not best understood as simple counter-statements dangling from a given micro-argument. In most cases, a rebuttal sets up a counter-argument -- an opposing line of reasoning (although, as we saw previously, it may also function to apply the current line of reasoning to a known counter-example). Thus the rebuttal provides yet another location for extending and developing the argument, with the rebuttal statement undergoing a transformation of function with respect to the counter-argument. How the rebuttal statement functions vis a vis the elaborated argument structure depends upon the nature of the challenge. In the example below (Figure 7), a rebuttal attacks the warrant "mobile entities can

1. For example, IBIS structures support the representation of "pro arguments" and "con arguments" via supports and objects-to links to Positions; however, there is no explicit way to attack specific parts of a pro or con argument.

be searched without a judicial warrant," attempting to show it is not a good decision rule. The rebuttal statement "mobility alone is an inadequate justification for warrantless search" becomes a claim supported by evidence about the resolution of past cases.

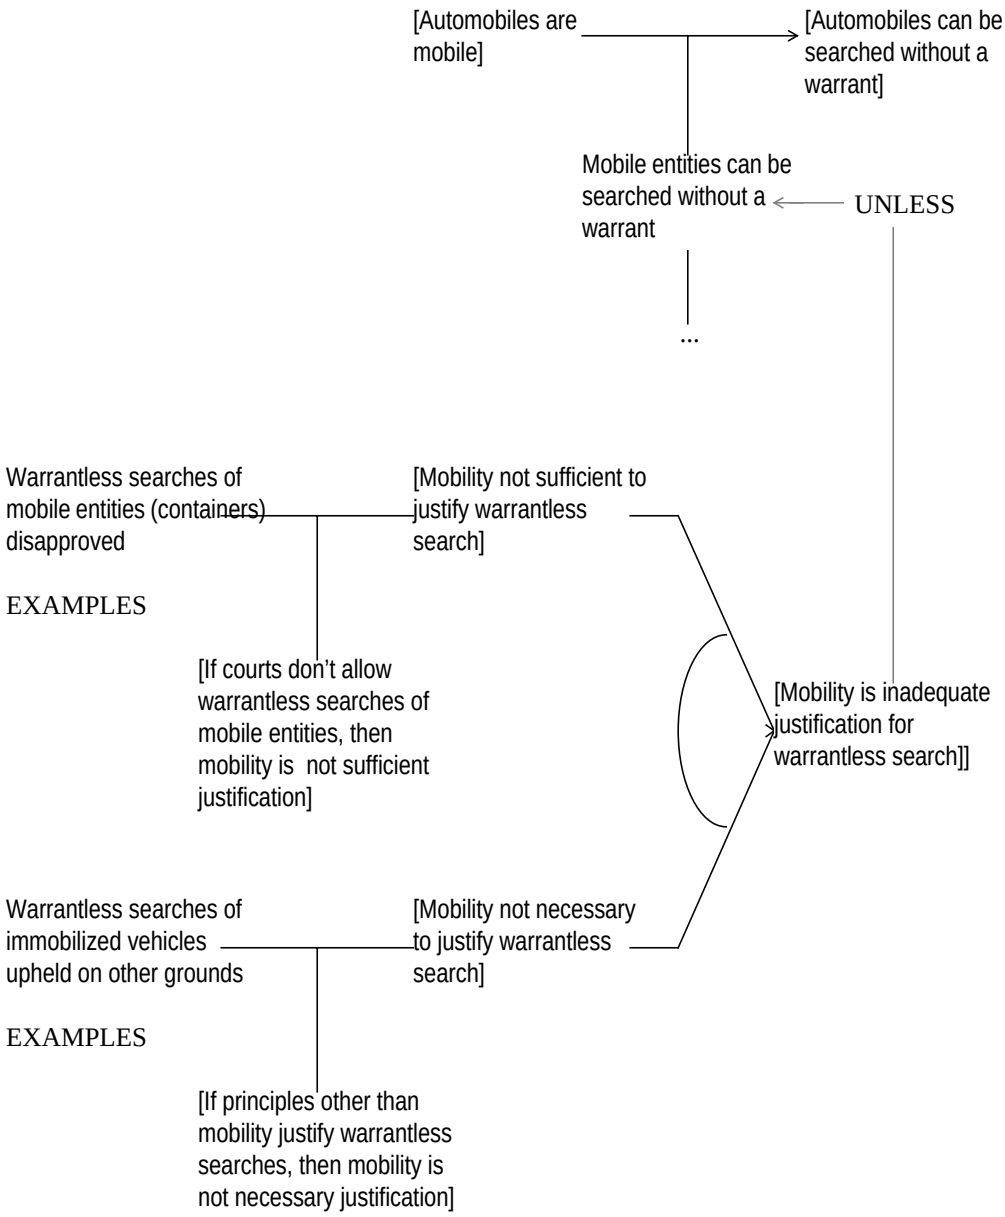

Figure 7. Connection by rebuttal

2.4. Evaluation of Toulmin structures. In addition to proposing ways that Toulmin structure might be modified to capture extended argumentative reasoning, the preceding discussion has helped us to understand how Toulmin's categories and relations can be mapped into argumentative

discourse and hence something about the elements from which arguments are composed. Before discussing further extensions to Toulmin structures, let us step back and evaluate the system in more general terms. We propose two ways of evaluating Toulmin structures -- first, as an analytic scheme and tool for research into argumentation; and second, as an argument representation scheme for use in constructing, analyzing, evaluating and comprehending structured reasoned discourses.

With respect to the first, we believe that Toulmin's scheme has a number of interesting properties as an analytic scheme and research tool. First it provides a structured system for categorizing the inference rules (warrants) and knowledge sources (backings) underlying particular arguments in a domain. Toulmin analyses carried out over a number of arguments in a field could help us understand what knowledge sources are important in that domain and how they are woven together in reasoning. Analysis of inference rules might enable us to build useful taxonomies of argument types and a better understanding of how different kinds of arguments (classificatory, predictive, etc.) are related in different domains and in different kinds of reasoning tasks. Since different argument types may have different informational and representational needs, such an analysis could help guide the design of flexible idea-processing environments. In addition, analysis of warrant and data types can identify recurring patterns of reasoning organized by conceptual models and accompanying argument strategies. Such recurrences may provide the basis for content- or strategy-based representational schemes.

Secondly, even in the short run, Toulmin's schema and theory provides an accessible analytic language and useful set of distinctions for describing certain aspects of argumentative practice. Looked at as answers to the types of questions that can come up with respect to any basic argument move (i.e., adducing evidence to support a claim), the elements of Toulmin's scheme provide the minimum categories we need to understand how arguments proceed -- specifically, what it is possible to question or dispute, and how. (As we suggested in our discussion of rebuttals, adequate treatment of rebuttals also requires "contextual" elements such as the framing goals and values that structure a particular line of reasoning, since these may be called into question as well.) Without understanding that these elements are always "present" in an argument, it may be difficult to understand how comprehension among argument participants breaks down. For this reason, the elaborated Toulmin model of argument that we present here may provide a useful framework for understanding problems that arise in the use of other argument representation schemes.

With respect to the second question -- the value of Toulmin structures as an argument representation scheme -- we propose an evaluation along four dimensions related to its representational fit to the domain of argumentative discourse: coverage, perspicuity, encodability,

and comprehensibility. We arrived at these dimensions from examining how, on one hand, the representation captures and highlights key aspects of the domain (coverage and perspicuity), and on the other hand, how the domain is mapped into the representation and subsequently reconstructed from that mapping (encodability and comprehensibility).

We wish to note here that, although we emphasize the relation between the representation and the represented domain without reference to a context of use, it is important to realize that assessment of representational fit is at bottom an interpretative enterprise. That is, representational fit is evaluated by a human understander, who must decide questions with respect to both the representational scheme and the represented domain. With respect to the representational scheme, s/he must arrive at a stable, and hopefully appropriate, interpretation of the scheme's categories with respect to the target domain, as we saw in the previous discussion of Toulmin elements. And of course with respect to the represented domain of argumentative discourse, s/he must both arrive at some interpretation of the discourse and decide questions such as what the "important" aspects of its structure are.

2.4.1. Coverage. Coverage addresses whether the phenomena we recognize as central to argumentative reasoning are captured by the representation. To evaluate the coverage of Toulmin structures, we need to ask several questions. First, are there entire categories of arguments that somehow elude representation as Toulmin structures? Second, do the resulting structures capture all of the important features of the reasoning?

In answer to the first issue, Toulmin structures are general enough to capture the basic inferential structure of most clearly argumentative discourse. Thus we found very few types of arguments which could not be coerced into Toulmin structures. However, to understand Toulmin's limitations, it is instructive to look at the case of analogies. Although they can be expressed as Toulmin structures (see Figure 8), the representation does not capture central aspects of that type of reasoning. In particular, it is interesting to explore the kind of analogies that are tightly connected to the argument's rhetorical force. Since their effectiveness does not simply arise from the cogency of their underlying inferential structure but rather from the creation of associations between ideas which leads to the transfer of commitment or feeling from one to another, it is important to capture not only the principles motivating the comparison, but also the values associated with the objects used in the analogy. The principles will allow us to pick out the crucial determinants of similarity that make the analogy possible, and the values will help us understand the transfer of commitment.

For example, in the oral argument, the defense attorney appeals to the hypothetical example of a

covered wagon -- would a covered wagon (which has wheels and is mobile) be included in the automobile exception and thus excepted from Fourth Amendment protection. To capture the meaning and force of this analogy, the representation must express the fact that the privacy and mobility characteristics that a covered wagon shares with a motor home are the crucial determinants of similarity, and that the set of values evoked by the symbol of the covered wagon is meant to be transferred to the motor home. Toulmin structures have no convenient way of expressing either of these things. We will discuss some other problems with representing analogy in the following section on perspicuity.

The second question of whether all the important features of the argument's reasoning are captured can most easily be addressed by looking at the results of an encoding performed using the first-order extensions and comparing it to the oral and written discourses. Once we see the argument's reasoning not as a search for timeless truths through inference and the counterposing of objective propositions by disinterested parties, but as skilled wrangling by interested participants guided by a melange of overlapping and conflicting goals and assumptions, it is easy to see that the network of interconnected logical structures is missing several important aspects of an argument's reasoning.

First, Toulmin structures provide no way of expressing ownership. Yet knowing who "owns" a given argument is necessary to having a clear understanding of the implications of particular argument moves. For example, in order to understand the force of certain kinds of rebuttals (see previous discussion of types of rebuttals), it is necessary to know whether the owner of the original claim is merely placing qualifications on the warrant, or whether an adversary is attacking it.

A related problem is the lack of facilities for adopting or presenting different perspectives on the argument and the issues it addresses, depending on the point of view of the producer or comprehender of the argument. In contrast, the written argument demonstrates many resources for capturing the relation between the author and different parts of the argument. One example of such a resources is the ability of the author to adopt an external "meta" view of an argument, by describing or summarizing it, or an internal one, by making it.

Furthermore, the network does not capture how an argument develops through time; the structure that has been built is essentially static, while in linear text as well as in oral discourse, the argument progresses. Thus, in the progression of both the oral and the written argument, beliefs and inferential pathways may be developed and summarized as consensus is reached or discarded as they are refuted. As a result, some parts of the argument become less important than others as the argumentative situation changes, but in the existing Toulmin scheme, these have the same

representational status as currently important lines of thought. (See also Conklin and Begeman [1988] for additional discussion of issues surrounding the need for information decay in static representations).

These problems are related to the homogenization of the argument, with respect to point of view and time, that the representation imposes. This homogenization prevents it from capturing shifts in the significance of the reasoning.

2.4.2. Perspicuity. To evaluate perspicuity we ask how well the representational scheme delineates and highlights salient structural characteristics of the domain. Given that Toulmin structures are intended to express reasoning, it is important to identify whether the structure of some types of reasoning is obscured by the way arguments are expressed.

Although we found that some types of reasoning - in particular, detailed inferential chains - are well delineated by Toulmin structures, others such as analogy and parallelism, comparative reasoning, constraint-based reasoning, and reasoning from models, become difficult to perceive when expressed as Toulmin structures. Once again, we will look at analogical reasoning to demonstrate how perspicuity is lost. In Figure 8, we compare three possible ways to express analogy in Toulmin structures; the illustration includes both generic forms and their instantiation with examples from our analysis. The first form (developed by Brockreide and Ehringer, 1960) uses a standard warrant/backing pair for analogical reasoning, "X and Y are essentially similar, and here's how." This form relegates the enumeration of the dimensions of similarity of the individual data to the backing and omits the reason why the dimensions are important for the comparison. The second form applies the same warrant, "Temporary residences are afforded full protection," to both data to support analogous claims, but leaves the actual analogy-making implicit. Finally the third uses an intermediate step of posing category membership. However, none of the three highlight all aspects of the structure of the analogy. In particular the view of analogy as a mapping between domains along particular dimensions is not supported by either the graphical layout or the way Toulmin's categories carve up the domain of argument.

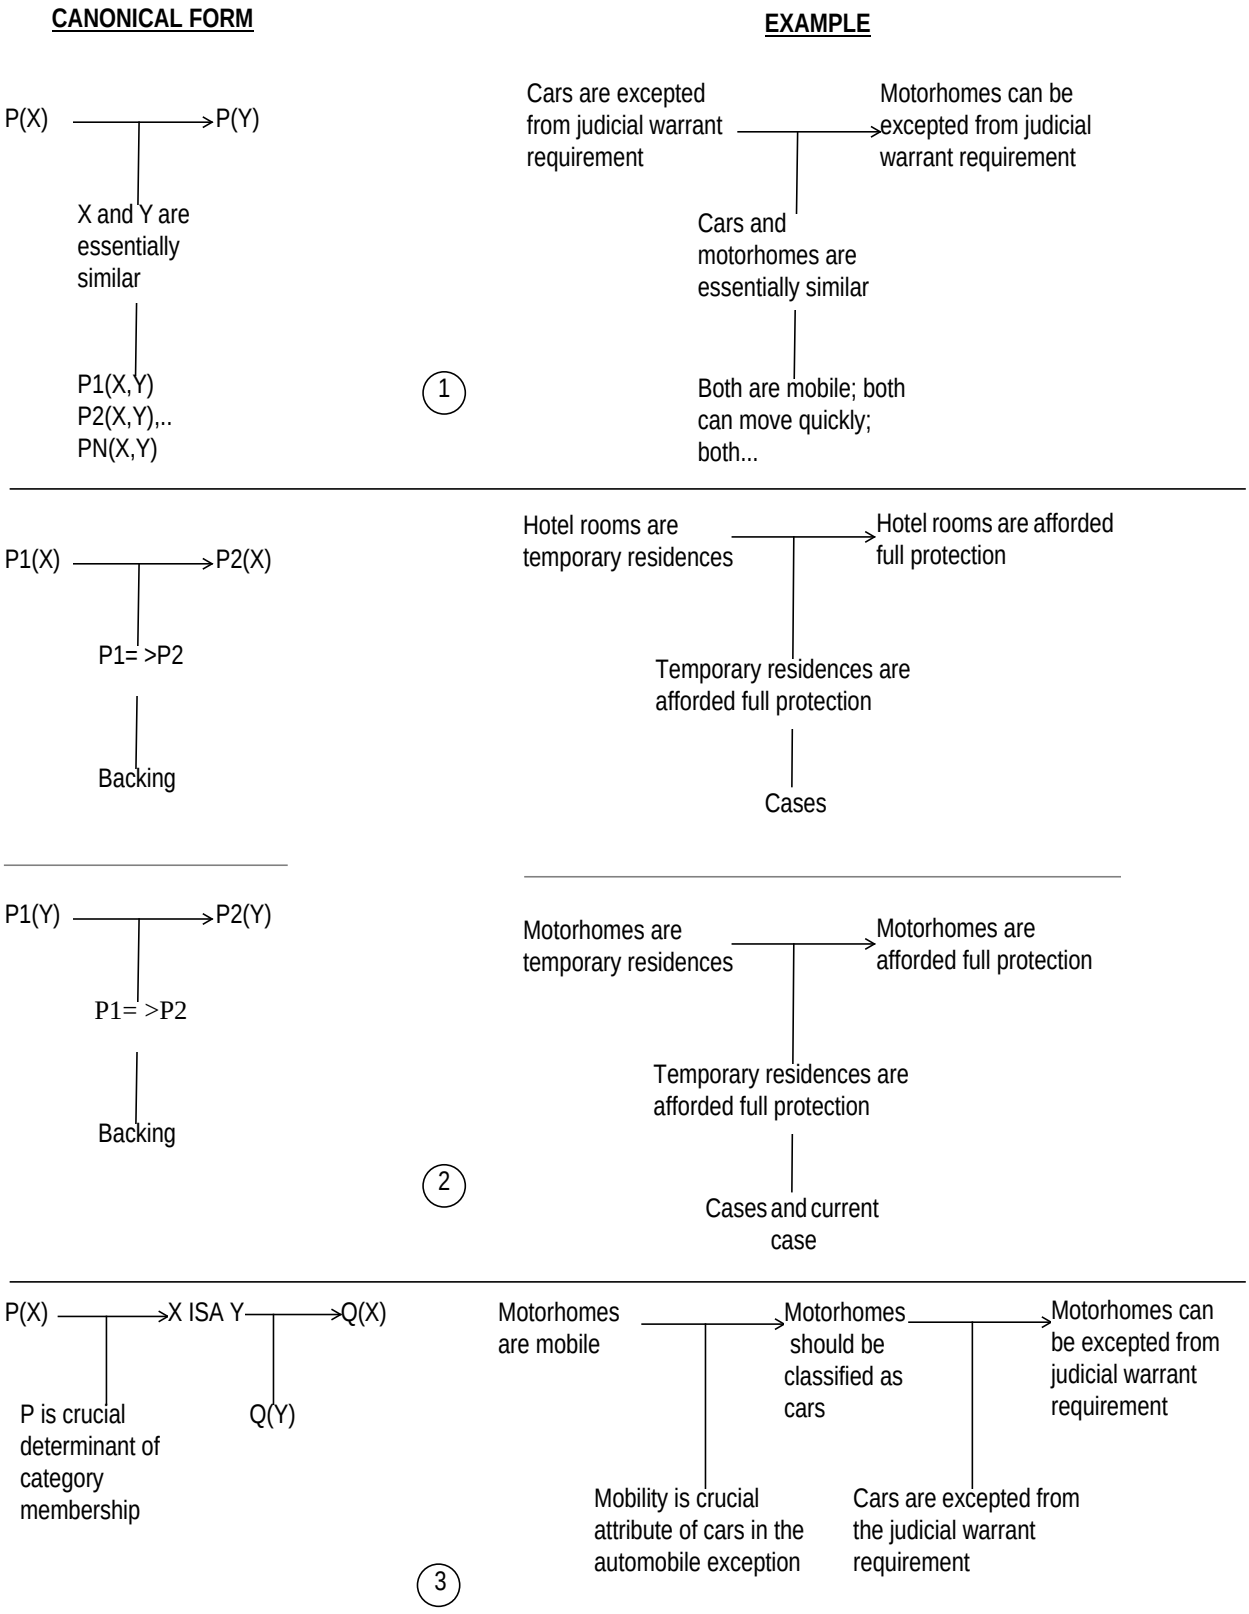

Figure 8. Comparison of three forms (canonical and examples) for expressing analogies in Toulmin structures

In the case of reasoning from models, we find that the knowledge embodied by the model is distributed throughout the network of Toulmin structures, and is therefore not accessible as an integrated structure. For example, in *Carney*, some of the reasoning relies on a classificatory model of the attributes of vehicles and residences (see Figure 10 below); instead of promoting the coherent articulation of this model, the Toulmin encoding spreads it over many structures. In addition, model-based reasoning involves the mapping of particular situations into general categories. However, Toulmin provides no perspicuous support for relating the general and the particular, or indeed, as we saw in the case of analogy, for mappings of any sort.

2.4.3. Encodability. Encodability concerns how readily and consistently elements in argumentative discourse can, in practice, be mapped into the categories and relations provided by the target representation, in this case Toulmin structures.

In using Toulmin structures to analyze the *Carney* arguments, we found two general types of encoding problems. One was concerned with implicitness - the fact that elements necessary to the Toulmin scheme frequently do not show up in the discourse - and the other with classificatory ambiguity. With respect to implicitness, consider again Figure 7 showing rebuttal relations between arguments. In this case, only the two elements encoded as data, "Warrantless searches of mobile entities (containers) disapproved" and "Warrantless searches of immobilized vehicles upheld on other grounds," along with examples from past cases, are provided by the source text. As we have discussed, the encoder must infer much implicit structure in order to hook these propositions and the examples that explicate them into the existing argument structure.

As a rough index of what aspects of argument, viewed under the Toulmin model, are implicit, we counted the number of elements of each type that were explicitly expressed in the written argument. Not surprisingly, data and claims were explicit most often. On our analysis, 22 out 27 data were explicit, and 12 out of 19 claims were explicit. Of those that were implicit, most could be inferred from the known focus of the discourse and so were only nominally implicit; a few were necessary for intermediate inferences. These trends were reversed for warrants and backings. Only 4 out of 23 warrants and 5 out of 23 backings were explicitly expressed. The four explicit warrants, interestingly, were crucial to the court's deciding of the case. Similarly some of the five explicit backings constituted the justificatory structure behind the dominant warrant: "In most cases private property may not be searched without prior approval of a judicial officer." Some rebuttals (3 out of 8) were implicit in the sense that they were only expressed by a statement negating them.

The second encoding problem, classificatory ambiguity, can be illustrated with the example text

from above. In order to encode this text, we had to decide whether the examples from past cases should be encoded as data to support a generalized formulation or, as in other encodings, as case-based backings from which some implicit warrant should instead be drawn. Resolving the uncertainty depended on resolving the argument structure of the immediately surrounding discourse; that is, decisions about particular encodings provide constraints on other nearby encodings.

Resolving questions about implicit elements and ambiguities led, over time, to the development of a set of partially articulated encoding conventions that helped us coordinate our analyses. This experience leads us to believe that the ability to resolve ambiguities or to find the "best" registration with respect to the representation requires a theory and practice of encoding that must be developed and negotiated within a community.

Other researchers have found that using Toulmin structures for argument analysis involves tackling difficult encoding problems (see for example (Hair and Lewis 1990)). However, we should not assume that other "better" schemes will not encounter similar problems. Specifically we expect that encoding problems associated with implicitness and ambiguity will arise whenever people attempt to adapt their diversely and implicitly structured understandings of their domain and task to a standardized representational scheme. Examination of encodability issues should thus be a part of our evaluation of any proposed representation scheme.

2.4.4. Comprehensibility. The evaluation of comprehensibility can be thought of as a kind of second-order evaluation of the representation's coverage. In other words, even though Toulmin structures were never intended to capture aspects of argument structure above the level of detailed reasoning, we can use the Toulmin analysis of the argument in order to throw into relief the aspects of the structure that are not included in the resulting network. We used two strategies to evaluate comprehensibility. First, we examined the encoded arguments and assessed whether the results are comprehensible, i.e., whether the argument is organized by Toulmin structures in such a way that it can be understood or re-interpreted. Second, we analyzed the non-Toulminizable aspects of the written text to understand the resources employed by linear text to render an argument comprehensible.

Our analyses here agree with that of other argumentation researchers, who note that Toulmin structures exhibit several important deficiencies that prevent a reader from understanding and using the network (see for example (Lowe 1986), (Lee 1990a, 1990b), (Homer-Dixon and Karapin 1989), or (Marshall 1989)). These deficiencies include lack of a topic or other sorts of infrastructure for organizing the detailed argumentation, inability to express argument summaries

and groupings, inability to compare positions (including points of agreement), lack of information about argument resolution, and inability to display the structuring role of values and goals.

These deficiencies should not be too surprising, since Toulmin addresses the structure of detailed argumentation. Many other representational needs arise when we consider the organization of the argument for activities such as competitive evaluation (where, for example, comparison between positions with respect to goals or other evaluative criteria become important) or problem decomposition (where, for example, topic infrastructure is helpful). Representations covering these needs frequently provide the infrastructure necessary for comprehensibility. Many developers of argumentation-related tools have compensated for the representational deficiencies of Toulmin structures by coming up with systems of representations; we describe our own ideas for these kinds of augmentations in the next section covering second-order extensions and alternatives.

2.5. Second order extensions and alternatives. The first-order extensions to Toulmin structures define some of the ways in which argument micro-structures can be constructed and composed to form larger networks expressing lines of reasoning. Many of these extensions resulted from grappling with problems of encodability. However, other problems remain, especially problems of coverage and comprehensibility; in particular, we have been concerned with capturing the way lines of reasoning are organized, compared, and evaluated. We have also attempted to address the perspicuity problem that we saw manifested by the distribution of model-derived knowledge. We have found several general types of second-order extensions that perform these functions: issue or problem decomposition structures, descriptive or classificatory models, and frameworks that enable comparison or competitive argumentation.

2.5.1. Problem decomposition representations. In our experiences developing and understanding argument networks, problem decomposition is an important organizing mechanism for large and complex spaces. In order to understand the relation between lines of argument encoded as Toulmin structures, we require some decomposition of a problem or separation of issues that group arguments with respect to particular topic they are addressing; the IBIS representation performs this kind of decomposition, as does Lee's (1989) notion of goals and subgoals in DRL. Van Lehn (1985) also used an issue decomposition representation to structure subsequent argumentation.

Figure 9 shows a top-level issue decomposition for the detailed argumentation that takes place in *People v. Carney* (see also Rissland, 1989 for a more detailed, and slightly different decomposition of the issues or "dimensions" in a framework oriented toward the use and generation of hypotheticals in the *Carney* oral argument). This diagram can be seen in IBIS terms as an issue

node expressing a central question in the case -- How we should determine whether a vehicle is excepted from 4th amendment protection -- and the three positions that respond to it -- by its inherent mobility, by its function, and by its expectation of privacy. We have already discussed how we have captured some of the detailed arguments that support and object to the three positions in Toulmin structures.

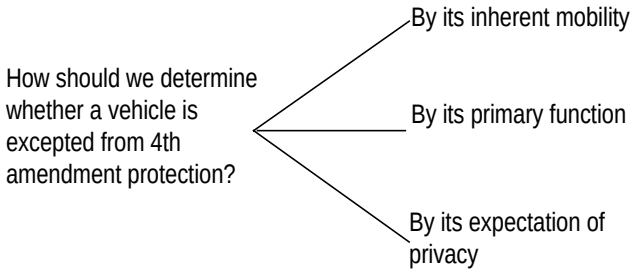

Figure 9. Top-level issue decomposition organizing *People v. Carney*

In addition to issue structures, we can use descriptive models to decompose or structure detailed analysis or argumentation. Such models can be used to extend Toulmin's notion of backing; they provide the detailed set of relationships from which a rule is derived or over which reasoning is performed. For example, we can collect all of the assumptions we have made about the characteristics of vehicles and homes and structure them as a classificatory model. We can then reason about whether a given entity should be classified as a vehicle or a home simply by comparing its attributes to the categories of the model. A portion of this model is shown in Figure 10. In addition to supporting reasoning by mapping, this type of model functions to gather knowledge that has been distributed throughout the network of Toulmin structures; once the knowledge is gathered and structured, its structure becomes perspicuous in the system of argument representations.

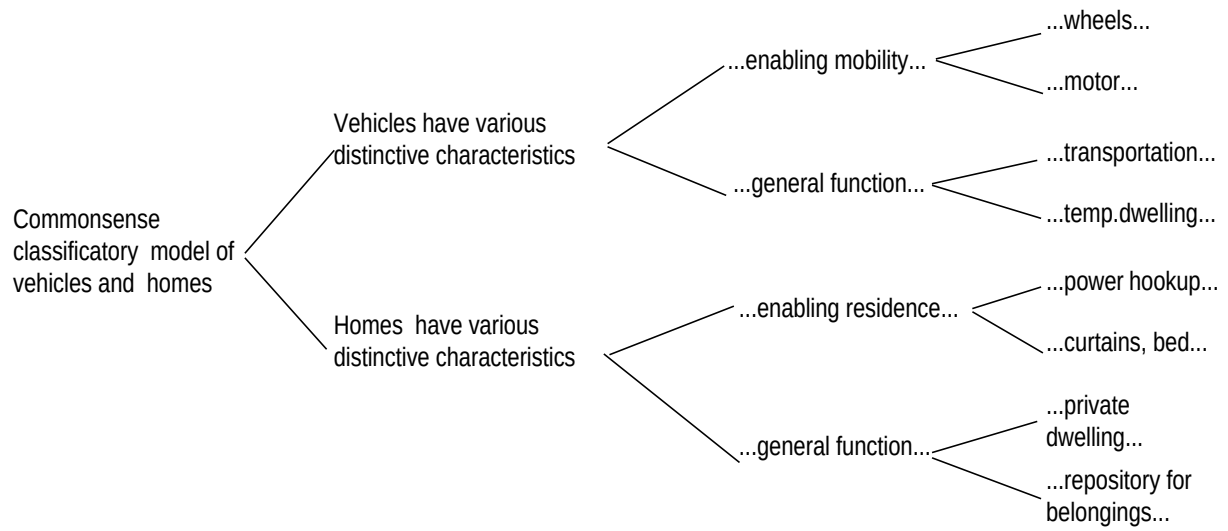

Figure 10. A portion of a classificatory model for vehicles and homes

2.5.2. Comparison and competition. Another organizing function for second order extensions to Toulmin is to allow comparison and competitive evaluation of lines of argument with respect to goals, values, theories, principles, and other structuring criteria for judging their relative strength and merit. Often we see this type of extension realized graphically as a matrix, where the structuring elements form one dimension of the matrix, the competing options (e.g., different proposed decision rules) the other dimension. The detailed argumentation generally lies beneath the cells. See vanLehn (1985), Marshall (1987), Lee (1990a, 1990b) and MacLean et al. (1989) for examples of how matrices can be used to structure argumentation.

In our example of the *Carney* case, ways of defining open-textured concepts are evaluated with respect to the various judicial principles that guide the case. Figure 11 shows an example of a matrix used to present this kind of competing argumentation. The purpose of the matrix is to explore whether mobility can be defined in terms of a vehicle's attributes, such as having wheels, having a motor and wheels, being able to move very quickly, or being able to move at all. These options are weighed against the judicial principles at work in a search-and-seizure case; on one hand, the justices would like to protect privacy, and on the other, they recognize the law enforcement establishment's desire to preserve evidence. They arrive at a third principle, the idea that the operators of a vehicle, in certain circumstances, will have a diminished expectation of privacy. The values in the cells of the matrix reflect the evaluation of the options for defining mobility with respect to the judicial goals. A "+" value means that Toulmin structures exist that test the option with an example; a "-" value signals the existence of a counterexample or effective rebuttal. The "NA" value means that no evaluation has taken place, and neither examples nor counterexamples have been offered.

	Judicial principles		
Mobility definitions (in service of automobile exception)	Protection of privacy	Preservation of evidence	Diminished expectation of privacy in vehicle
Entities with wheels are mobile	-	+	NA
Self-propelling entities are mobile	+	-	NA
Entities that can be moved quickly are mobile	-	+	NA
Movable entities are (potentially) mobile	-	+	NA

Figure 11. An options/goals matrix for competing arguments in *People v. Carney*

One problematic aspect of using matrices to show how arguments compete in a semi-structured representational system is related to the use of the value "NA" in the evaluative cells of the matrix. The need for this value results from the fact that the matrix form often requires irrelevant judgements to be made. In our example, the principle of "diminished expectation of privacy in a vehicle" became important only after all the mobility criteria had been explored and successfully refuted. Hence, a principle that is crucial for evaluating some of the decision rules simply does not apply to others.

3.0. Usability Issues for Argument Representation Schemes

We began this paper by elaborating three "design dimensions" for evaluating computer-mediated representational schemes: representational fit, usability, and computational power. We have devoted most of our attention to elaborating aspects of the relation between argument representation schemes -- Toulmin structures in particular -- and the domain of argumentative discourse. However, a scheme's representational characteristics are not unrelated to either its usability or its computational power. Our experience using Toulmin structures and other argument representation schemes to analyze realistic arguments has given us a new perspective on questions of usability and computational power. In this concluding section, we discuss this perspective as it relates to usability, leaving discussion of computational power for another paper.

3.1. Fitting Representational Schemes to Existing Cognitive, Perceptual, and Activity Structures. The most obvious approach to assessing the usability of representational schemes -- at least prior to widespread field testing of representational tools -- is to evaluate how well the scheme fits with existing structures. Traditionally, human-computer interaction researchers have

addressed such questions from a cognitive perspective, proposing criteria intended to evaluate a system's "cognitive compatibility" (see, for example, Neuwirth and Kaufer's 1989 discussion of the cognitive determinants of a "good" representation). In line with this way of thinking, it seems reasonable to suppose that research into prospective users' intuitive (articulate) models of argument structure might lead to representational schemes with a better fit to users' cognitive and perceptual structures. In particular, it would be interesting to compare the relative importance of content, functional, and strategic models of argument structure.

More recently, researchers have begun to pay attention to what we call here "activity structures" -- that is, the detailed ways that people structure the sorts of tasks that our systems are intended to support. Of particular interest are the nature and role of external resources in organizing the task and the importance of social negotiation of all aspects of the task (see Tatar, Foster, and Bobrow, forthcoming, for an interesting discussion of how research into the structure of conversational activity helped pinpoint problems and suggest solutions in the design of a tool intended to provide computational support for design meetings). In contrast, reliance on purely cognitive task analyses is likely to be inadequate or misleading, for several reasons: (1) such analyses tend to focus on individual cognition, thus entirely missing the role of social negotiation; (2) traditionally, cognitive analyses do not systematically distinguish between internal and external sources of structure -- for example, between internal and external memories. In order to build representational tools that fit existing activity structures, we may need to expand our understanding of the task to include aspects of the mundane activity by which the task is carried out -- including in particular the ways people use existing representational tools and media, such as word processors, whiteboards, paper-and-pencil, and post-its, as well as media that impose more structure, such as outlining tools, forms, and domain-specific representational languages or notational systems.

3.2. The Relation between Usability and Learning, Adaptation, and Change. Defining usability as the fit between representational schemes and the cognitive, perceptual and activity structures of users provides a useful heuristic for the design of usable representational schemes and associated tools. However, such a definition may also be misleading. Specifically , our emphasis as a field on finding the "right" or "better" representational forms may interfere with recognizing the inevitability of learning, adaptation, and change in the use of representational schemes and tools. We propose that, to a greater degree than has been recognized, usability is a function of the flexibility and mutability of the overall user-tool-task system.

In particular we see a pivotal role for learning in the acceptance and productive use of representational tools. Our experiences using argument representation schemes suggest that two general "directions" of learning are important: (1) adaptation of the cognitive, perceptual, and

activity systems of users to the representational scheme and tool; and (2) adaptation/mutation of representational media to satisfy the evolving requirements of the activity. To conclude, we discuss a few different aspects of learning and change as they relate specifically to the use of representational schemes.

3.2.1. Adaptation in the direction of the representational scheme: cognitive and task evolution. An important and underappreciated aspect of the adoption and use of representational tools concerns the need to learn to encode ideas, discourse, events, etc. in the representational framework. In the discussion of Toulmin structures, we described encoding difficulties that arise from implicitness and ambiguity in the relation between the scheme and the represented domain. While other representational schemes may not suffer to the same degree from these particular encodability problems, we would argue that encoding difficulties are inherent in the enterprise of mapping a limited, regularized schema to the contingent, multiply structured, dynamically reconfigurable "stuff" of a domain. At a minimum users must learn to ignore or transform accustomed ways of structuring their thoughts and discourse in favor of relating them to the categories and relations imposed by the representational schema. Segmentation of ideas into the different slots of the schema is likely to provoke difficulties, as are decisions about what to do with material or understandings that falls outside the scheme's coverage or perspicuity limitations.

Moreover, the resolution of these difficulties cannot be seen simply as an individual cognitive issue if representational schemes are to serve to structure cooperative work. As we worked with Toulmin structures, for example, we found it necessary to negotiate a number of shared encoding conventions in order to coordinate our analyses. This was not always easy, since we had somewhat different goals for the analysis, which in turn influenced our ideas about useful encoding practice. We doubt that we are unique in this regard. The encoding practices involved in using standardized representational schemes must inevitably be developed in a social context characterized by multiple agendas and different initial understandings of the domain.

Furthermore, as designers we need to consider, and anticipate responses to, the possibility that more radical changes will result from using standardized representational schemes to structure idea-processing work. For example, the social and cognitive work involved in categorizing one's actions and/or the content of one's ideas may contribute to task fragmentation, or at least change the processes by which users form an integrated understanding of the domain -- or even the nature of understanding itself (see Russell et al., 1990, for evidence of these kinds of changes from their experiences with IDE). Such changes affect not only individuals' relation to their idea-processing work, but also potentially the distribution of task knowledge and responsibility.

In sum, as many designers have pointed out, we can anticipate a number of advantages from the widespread use of computer-mediated representational schemes and associated tools, such as being able to store, access, and manipulate structured representations of the domain. Explicit representation and manipulation of implicit structures may even contribute to interesting restructurings of users' knowledge and activities, discovery of useful new generalities, and streamlining of processes. However, the human activities involved in organizing, maintaining, and communicating information in structured forms are different and will require the development of new social as well as cognitive structures. It is perhaps for these sorts of reasons that Conklin and Yakemovic (in press) propose that in designing representational tools for capturing design rationale, we should be looking to replace old activity structures with new ones. Such a view requires a different look at issues surrounding adaptation and change; specifically proponents need to understand how resources and constraints affect the reorganization of activity structures within organizational settings. That is we may need to think of learning not just at the individual cognitive or small-group level, but also at the organizational level.

3.2.2. Adaptation in the direction of the evolving task and cognitive requirements: representational mutation. The previous discussion highlights some of the ways in which activities involved in using representational tools may lead to changes in the cognitive, social and even institutional organization of idea-processing work. In our experience, pressures for adaptation and change move in the other direction as well; that is, as the users' understanding of the task evolves over time, so do their ideas about how to represent relevant knowledge and task structure. For example, as we progressed in our understanding of both the domain of argumentation and of our analytic task, we felt the need to collapse some distinctions provided by the Toulmin scheme and to invent others. The first- and second-order distinctions discussed in Section 2 provide some evidence of these changes.

Many researchers already recognize the evolution of representational commitment in the course of so-called ill-structured idea-processing tasks, with resulting proposals for moving from informal to formal and from general to specific representational schemes as the structure of the task develops. We believe that in addition to shifting between pre-determined representational schemes, users may well need the ability to modify the categories and relations of existing schemes, as well as the rules for composing them. If our experiences are typical, the capacity for representational mutation, reconfiguration, and retrofitting of existing structured databases is crucial to the widespread acceptance and use of representational tools.

Acknowledgments

We wish to thank the members of ARG (Argumentation Reading Group) - Danny Bobrow, Frank Halasz, Bill Janssen, Jintae Lee, Russ Rogers, Dan Russell, and Norbert Streitz - for their energetic criticisms of an earlier draft. We express special thanks to Frank Halasz for his on-going support and help in developing our ideas.

References

(Allison, 1971) Allison, G., *Essence of Decision: Explaining the Cuban Missile Crisis*. Boston, MA: Little Brown and Company, 1971.

(Ashley and Rissland, 1985) Ashley, K. and Rissland, E. "Toward Modelling Legal Argument." CPTM #7, University of Massachusetts, 1985.

(Brockreide and Ehringer, 1960) Brockreide, W. E. and Ehringer, D., "Toulmin on Argument: An Interpretation and Application," *Quarterly Journal of Speech*, XLVI (February 1960), pp. 44-53. Reprinted in R.L. Johannese, (Ed.) *Contemporary Theories of Rhetoric: Selected Readings*. New York: Harper & Row, 1971.

(Conklin and Yakemovic, in press) Conklin, E. J. and Yakemovic, K. B., "A Process-Oriented Paradigm for Design Rationale," *Human-Computer Interaction Journal*, Vol. 6, No.3-4.

(Conklin and Begeman, 1988) Conklin, E. J. and Begeman, M.L., "gIBIS: A Hypertext Tool for Exploratory Policy Discussion," MCC Technical Report Number STP-082-88, Austin, Texas, 1988.

(Cutting et al., 1989) Cutting, D., Halvorsen, P.K., Pedersen, J., and Withgott, M., "Information Theater versus Information Refinery," Xerox Palo Alto Research Center Technical Report, December, 1989.

(Fischer et al., 1989) Fischer, G., McCall, R., and Morch, A., "JANUS: Integrating Hypertext with a Knowledge-based Design Environment," *Proceedings of Hypertext '89*, Pittsburgh, PA, November 5-8, pp. 105-117.

(Hair and Lewis, 1990) Hair, C. and Lewis, C., "Are Argument Representation Schemes Useful?" University of Colorado at Boulder, Department of Computer Science Report Number CU-CS-475-90, June 1990.

(Hayes and Flower, 1980) Hayes, J.R., and Flower, L.S., "Identifying the Organization of Writing Process," in L.W. Gregg and E.R. Steinberg (Eds.) *Cognitive Processes in Writing*. Hillsdale, NJ: Lawrence Erlbaum, 1980, pp. 3-30.

(Hillocks, 1987) Hillocks, G. Jr., "Analyzing the Structure of Written Arguments," Paper presented at the Annual Meeting of AERA, Washington, D.C., April 1987.

(Homer-Dixon and Karapin, 1989) Homer-Dixon, T.F. and Karapin, R.S., "Graphical Argument Analysis: A New Approach to Understanding Arguments, Applied to a Debate about the Window of Vulnerability," *International Studies Quarterly* (1989) **33**, 389-410.

(Kintsch and van Dijk, 1985) Kintsch, W. & van Dijk, T.A. "Toward a Model of Text Comprehension and Production." *Psychological Review*, January 1985, 92:1, 109-129.

(Lee, 1990a) Lee, J. "SIBYL: A Qualitative Decision Management System," to appear in Winston, P. and S. Shellard (Eds.) *Artificial Intelligence at MIT: Expanding Frontiers*, Chapter 5, The MIT Press: Cambridge, MA, 1990.

(Lee, 1990b) Lee, J. "SIBYL: A Tool for Sharing Knowledge in Group Decision Making," to appear in *Proceedings of CSCW '90*, Los Angeles, CA, October 7-10, 1990.

(Levi, 1949) Levi, Edward H., *An Introduction to Legal Reasoning*, The University of Chicago Press, 1949.

(Lowe, 1985) Lowe, D. G., "Co-operative structuring of information: the representation of reasoning and debate," *Journal of Man-Machine Studies* (1985) **23**, 97-111.

(Lowe, 1986) Lowe, D. G., "SYNVIEW: The Design of a System for Cooperative Structuring of Information," *Proceedings of CSCW '86*, Austin TX, December 3-5, 1986.

(MacLean et al., 1989) MacLean, A., Young, R., and Moran, T., "Design Rationale: the Argument behind the Artifact," In *Proceedings of CHI '89*, (Austin, TX, April 30-May 4, 1989), ACM, New York, pp. 247-252.

(Marshall, 1986) Marshall, C.C. "Representation of a Rational Actor Model in NoteCards," Xerox Special Information Systems Technical Report, August, 1986.

(Marshall, 1987) Marshall, C. C., "Exploring Representation Problems Using Hypertext," *Proceedings of Hypertext '87*, Chapel Hill, North Carolina.

(Marshall, 1989) Marshall, C. C., "Representing the Structure of a Legal Argument," *Proceedings of the 2nd International Conference on AI and Law*, Vancouver, BC, 1989, pp. 121-127.

(Neuwirth and Kaufer, 1989) Neuwirth, C., and Kaufer, D., "The Role of External Representations in the Writing Process: Implications for the Design of Hypertext-based Writing Tools," *Proceedings of Hypertext '89*, Pittsburgh, PA, November 5-8, pp. 319-341.

(Newman, 1986) Newman, S.E., "Argument Scripts: Schemata for Representing (Roughly) the Discourse Structure of an Argument," Xerox Palo Alto Research Center Technical Report, October, 1986.

(Newman, 1987) Newman, S.E., "Argument Scripts: Functional Schemata and the Comprehension of Written Argument." Xerox Palo Alto Research Center Technical Report, 1987.

(Rissland, 1985) Rissland, E.L. "Argument Moves and Hypotheticals." In C. Walker, (Ed.), *Computing Power and Legal Reasoning*, West Publishing Co., St. Paul, MN, 1985.

(Rissland, 1989) Rissland, E.L., "Dimension-Based Analysis of Hypotheticals from Supreme Court Oral Argument," *Proceedings of the 2nd International Conference on AI and Law*, Vancouver, BC, June 13-16, 1989, pp. 111-120.

(Rissland et al., 1984) Rissland, E.L., Valcarce, E.M., and Ashley, K.D., "Explaining and Arguing with Examples," In *Proceedings AAAI-84*, Austin, TX, August, 1984, pp. 288-294.

(Rittel and Kunz, 1970) Rittel, H. and Kunz, W., "Issues as elements of information systems." Working paper #131. Institut fur Grundlagen der Planung I.A. University of Stuttgart.

(Russell et al., 1990) Russell, D. M., Burton, R.R., Jordan, D.S., Jensen, A.M., Rogers, R.A., Cohen, J., "Creating Instruction with IDE: Tools for Instructional Designers," Xerox PARC SSL report P88-00076, (August, 1988), in: *Intelligent Tutoring Systems*, (May, 1990), Reprinted in: Yazdani, M. and Lawler, R., (eds.) *AI and Education*, Learned Information, (1990)

(Scardamalia and Bereiter, 1987) Scardamalia, M. and Bereiter, C., "Knowledge telling and knowledge transforming in written composition." In S. Rosenberg (Ed.), *Advances in Applied Psycholinguistics: Vol. 2. Reading, writing, and language learning*, Cambridge: Cambridge University Press, 1987.

(Schum and Martin, 1985) Schum D. and Martin, A, "Formal and Empirical Research on Cascaded Inference in Jurisprudence." *Law and Society Review*, Vol 17, No. 1, 1982, pp. 105-152.

(Skalak, 1989) Skalak, D., "Taking Advantage of Models for Legal Classification, "*Proceedings of the 2nd International Conference on AI and Law*, Vancouver, BC, June 13-16, 1989, pp. 234-241.

(Smolensky et al., 1988) Smolensky, P., Fox, B., King, R., and Lewis, C., "Computer-Aided Reasoned Discourse or, How to Argue with a Computer" In R. Guindon (Ed.), *Cognitive Science and its application for human-computer interaction.* Norwood, NJ: Ablex, 1988, pp. 109-162.

(Stefik et al., 1987) Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S., and Suchman, L., "Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings," Communications of the ACM, Vol. 30, No. 1., January, 1987, pp. 32-47.

(Storrs, 1989) Storrs, G., "Group Working in the DHSS Large Demonstrator Project," Technical Report, Logica Cambridge Ltd., Cambridge, England, 1989

(Streitz et al., 1989) Streitz, N.A., Hannemann, J., and Thuring, M., "From Ideas and Arguments to Hyperdocuments: Travelling through Activity Spaces," *Hypertext '89 Proceedings,* Pittsburgh, PA November, 1989.

(Tatar et al., forthcoming) Tatar, D., Foster, G., and Bobrow, D. "Conversation in and around Cognoter", to be published in *International Journal of Man-Machine Studies.*

(Taylor et al., 1989) Taylor, A., Knox, M., and Weaver, J., "Policy Argumentation Support," Technical Report, Policy and Legal Informatics Group, University of Lancaster, Lancaster, UK, 1989.

(Toulmin, 1958) Toulmin, S., *The Uses of Argument,* Cambridge University Press, Cambridge, 1958.

(Toulmin et al., 1984) Toulmin, S., Rieke, R., Janik, A., *An Introduction to Reasoning*. MacMillan: New York, 1984.

(vanDijk and Kintsch, 1983) *Strategies of Discourse Comprehension,* New York, NY: Academic Press, 1983.

(vanLehn, 1985) vanLehn, K., "Theory reform caused by an argumentation tool," Xerox Palo Alto Research Center Technical Report, ISL-11, July, 1985.

(Voss et al., 1983) Voss, J.F., Tyler, S.W., and Yengo, L.A., "Individual Differences in the Solving of Social Science Problems." from *Individual Differences in Cognition,* Volume 1, pp. 205-233.