

A Memory Architecture for Case-Based Argumentation*

Eric Shafto
shafto@nwu.edu

Ray Bareiss
bareiss@ils.nwu.edu

Lawrence Birnbaum
birnbaum@ils.nwu.edu

The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201

Abstract

This paper describes a memory organization that supports intelligent memory-based argumentation. Our goal is to build a system that can argue opposite sides of an issue by retrieving stories that support or oppose it. Rather than attempting to determine how a story relates to a point on the fly, we explicitly represent the points that the stories support or oppose, as well as how they support or oppose those points. We have developed a hierarchy of story point types; associated with each type is a set of *rhetorical templates*, which describe the ways that a story could support or oppose a point of that type. Each template consists of a series of assertion types on which the argument depends. This enables the program to attack intelligently the foundations of the point it is trying to refute. Our approach is being developed within the context of the ILS Story Archive, a large multimedia case base which includes stories from a wide variety of domains.

Introduction

Argumentation is fundamentally a memory process. When we are presented with a point in an argument, we generally find that examples or facts that support or refute that point come to mind automatically and without conscious effort [McGuire *et al.*, 1981]. Clearly, complex inference is involved in determining whether some aspect of a given story can be used to strengthen or weaken a given argument. However, it is not at all clear that the bulk of this inference must take place in real time. For any issue about which a person

has an opinion, there are sure to be a large number of anecdotes, experiences, and memories of previous arguments that are permanently associated with that issue.¹ Our model is based on the idea that pre-existing links between points in an argument and stories that support or oppose those points play a key role in argumentation.

Our approach is intended to model a small but important part of the human ability to find stories that support or refute points. Obviously, the ability to hear a new story and determine that it supports a point, the ability to learn new points, and to re-evaluate a story with regard to other points is important to argumentation, and to understanding in general. These are interesting and difficult problems. However, we believe that a system that had to rely on arguing from first principles in this way would bog down immediately in an intractable memory search process. We wish to show here that much interesting behavior can be obtained through a system that relies on a corpus of pre-analyzed stories that are linked to the points they represent. The core of our theory thus lies in the taxonomy of points and link types embodied in the story point/rhetorical template hierarchy.

Our goal is to build a system that can argue for or against a point by recalling appropriate stories, and explaining how they support or oppose that point. This work is being carried out within the context of the ILS Story Archive Project. The Archive is a large multimedia case base of stories, including news footage, speeches, lectures, and interviews with experts. The overall goal of the system is to enable users to easily gain access to stories relevant to their interests, and in particular, to facilitate the retrieval of follow-up stories relevant to whatever story the user has just viewed. The underlying metaphor is that of a normal human conversation, in which the interaction flows naturally and coherently from one story to the next [Schank, 1977, Bareiss *et al.*, 1991, Ferguson *et al.*, in press]. At their best, arguments are a particularly informative kind of conversational

*This research was supported in part by the Defense Advanced Research Projects Agency, monitored by the Office of Naval Research under contracts N00014-90-J-4117 and N00014-91-J-4092, and by the Office of Naval Research under contract N00014-89-J-1987. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting. The Institute receives additional support from Ameritech, an Institute Partner, and from IBM.

¹We use the term "stories" to refer to all of these types of information in memory.

interaction. Our goal, then, is to extend the memory organization and retrieval strategies of the Archive to enable this kind of interaction.

Suppose that the user has just viewed the story "Brown on Violence," in which activist H. Rap Brown argues that violence is a justified means of achieving racial equality:

I say violence is necessary. Violence is part of America's culture. It is as American as cherry pie ... America taught the black man to be violent. We will use that violence, if necessary ... we will be free by any means necessary.²

Roughly speaking, Brown's main argument here is that the ends justify the means: Because racial equality is a desirable goal, violence is justified to the extent that it helps to achieve this goal.³ An argument of this sort can be opposed in at least four ways: By arguing that the means are unacceptable for some other reason, by arguing that they will not achieve the goal, by arguing that there are alternative means of achieving the goal, or by arguing that the goal itself is not desirable. Corresponding to each of these, the Archive contains at least one story, as follows:

- In "Birmingham Riots," Martin Luther King declares that violence is unacceptable because it is counter to Christian ethics.
- In "Race and Riots," Arthur Ashe argues that violence will not achieve racial equality, stating that it has never succeeded in the past.
- In "I Have a Dream," Martin Luther King claims that an alternative plan, non-violence, will achieve the goal of racial equality.
- Finally, in "Segregation Forever," George Wallace states his view that racial equality is not a proper goal for American society.

Our goal is to create a program that can oppose the points represented in a story by recalling appropriate stories, as illustrated above, and explaining their relevance. This requires, first, a representation of story points and how they can be supported or opposed; second, a retrieval method that can find stories that support or oppose a particular point; and, third, a method of generating bridging explanations between stories.

Representation

"The ends justify the means," as argued by H. Rap Brown, is exactly the sort of generic argument structure that our representation seeks to capture, in the form of *rhetorical templates* (cf. e.g. [Alvarado *et al.*, 1986, Birnbaum, 1982, Flowers *et al.*, 1982]). The two

²All of the examples in this paper are actual stories from the Archive.

³One subtlety missed by this representation is that violence needs some justification—that it is considered bad by default.

templates that make up Brown's argument are that the plan achieves a desirable goal, and that no other plan achieves that goal. Our representation of "Brown on Violence" includes instances of these templates, among others. It is of course apparent that there is more to Brown's argument than this, and some of it is quite subtle. However, such a representation of the point is sufficient to enable the retrieval of a number of relevant stories, as we have shown.

Much work has been done on the problem of representing and indexing stories for retrieval (see, e.g., [Schank and Osgood, 1990, Kass, 1990, Bareiss and Slator, 1992, Ferguson *et al.*, in press]). The earlier work on this problem generally tried to create detailed representations of the entities, actions, and relations that comprised a story. This approaches the ideal of a general-purpose representation — sufficient to infer the story's point, in whatever context one needs to understand the story. All too often, however, the points of stories are not readily emergent from this kind of representation, either to a human or a program. For example, it is difficult to determine computationally that some story is a case in which goal conflict was resolved by turn-taking, when the representation consists entirely of actions, agents, goals, plans, intentions, etc.

Our approach is to back off from this general-purpose ideal, and explicitly represent the points that the stories support or oppose (see also [Bareiss *et al.*, 1991]). Instead of creating detailed representations of the stories, we represent only the points of the stories, as perceived by human indexers. Furthermore, we do not build complex representations of the points themselves. Instead, we represent how the story supports or opposes the point. For example, one can argue that a person is good by arguing that they do good things, that they try to do good things, that they have good beliefs, that they possess a valuable skill, and so on.

For a point of a given type, there is a set of these rhetorical templates — ways a point of that type could be supported or opposed. We have developed a hierarchy of point types that is very well-behaved with regard to the rhetorical templates that are applicable to each type. That is, a template usable for some story-point type is usable by all the children of that type, and, in general, the same templates do not appear in different branches of the story-point hierarchy. A simplified portion of this hierarchy is shown in Figure 1.

The root of this hierarchy is *Any Assertion*. All story point types are specializations of this basic type. When an indexer categorizes a point as being of this very abstract type, all they are saying about it is that some statement was made. Nonetheless, there are a variety of weak, but commonly-used, templates that can be used to support or refute a point of this type. Assertion, authority, logical proof and appeals to common sense are all examples of rhetorical templates for supporting or opposing virtually any point. As we specialize the assertion, we gain more specific (and, in

general, more powerful) templates. For example, if the assertion is of the form *Something is good*, we can use the templates *Best available*, or *Fulfills function*.

However, when the *Something* in *Something is good* is specialized to be a plan, a resource, a goal, an agent, etc., the set of templates becomes extremely rich. For example, to support a point of type *Plan is good*, we can use any of the templates we have listed so far. In addition, we can use *Plan succeeds*, *Plan has positive side effects*, *Plan is low-risk*, *Agent has nothing to lose by executing plan*, and so on. This level of specificity seems to be a natural one for categorizing rhetorical templates. As we get more specific than this, we begin to talk about ways to support individual plans, or to show that certain people or types of people are good or bad. At that level, we would expect an explosion of categories, most of which would contain very little in the way of new, more-specific templates (see, e.g., [Rosch, 1988]).

Each instantiated rhetorical template is a sequence of assertions that are necessary to the story's ability to support or oppose the main point. They are represented in the same way as points, and this is what allows the program to undercut an argument. To rebut a point, the Archive can retrieve a story that opposes the main point directly, or it can retrieve stories that oppose the points that were used to support the main point. For example, as shown in Figure 1, the point type *Plan is good* can be supported by the template *Achieves good end*. This is elaborated as two statements (actually represented as clauses): *The plan meets some goal, and that goal is good*. As we saw in the Introduction, the program can oppose the *Plan is good* point by finding stories that oppose either of these two supporting points.

A partial representation in these terms for the H. Rap Brown story is shown in Figure 2. The point being indexed is that violence is a good plan for achieving racial equality. Since this is an assertion of the type, *Plan is good*, we must use one of the templates available to this type. There are two templates that are used to support this contention. The first is that violence achieves a good goal. The instantiated conjunct of clauses that supports that point can be seen in the figure. These clauses, as well as the point of the story itself, are used as queries to find opposing stories.

Retrieval

Given this representation, a very simple mechanism suffices for story retrieval. Any story point has pointers to all the stories that support or oppose it. When a story is needed to oppose a point, all stories listed as opposing it are retrieved.

In our first example, the probes *Violence is good*, *Violence achieves racial equality*, *No plan but violence achieves racial equality*, and *Racial equality is a good goal* are each used to find stories that oppose the idea




Figure 1: A portion of the story-point hierarchy. Pluses and minuses represent supporting and opposing templates, respectively. The italicized sentences are actually represented by a conjunctive list of assertion types, which are instantiated with fillers from the point.

“Brown on violence”

Speaker: H. Rap Brown

Supports: “Violence is a good plan”

Type: Plan is good.

Plan = violence

Agent = American-Blacks

Goal = Racial-Equality

Template: No alternative.

Relies on:

There is no plan that achieves racial equality except violence.

Template: Achieves good goal.

Relies on:

Violence achieves Racial equality.

Racial equality is a good goal.

“King: Birmingham Riots”

Speaker: Martin Luther King

Opposes: “Violence is a good plan”

Type: Plan is good

Plan = violence

Agent = American-Blacks

Goal = Racial-Equality

Template: Inconsistent with good theme.

Relies on:

Violence conflicts with Christianity.

Christianity is a good theme.

...

Figure 2: A sample representation. Each of the templates represents one index. There are other points that these stories support that are not shown here. The italicized sentences are actually represented by conjunctive lists of clauses.

that violence is a good plan⁴. Any story that opposes one of those points weakens Brown's argument in favor of violence.

However, this leaves us with a problem. Each of these retrieved stories may have many points. Sometimes, the point for which the story was retrieved is only a small part of the story as a whole. A user of the program might not understand why the story is being shown. This means that it is not sufficient to simply retrieve an appropriate story and present it to the user. Research on hypermedia presentations (*e.g.* [Oren *et al.*, 1990]) has demonstrated that users often have difficulty determining the relevance and intended meaning of a story. It is essential that any system that is intended to argue be able to explain why it is telling a given story — what its relevance is to the conversation.

Bridging

An important aspect of the work on the Story Archive is focused on the building of these “bridging” explanations. Viewing the Archive as a case-based reasoning program, bridging can be seen as a form of case adaptation. Instead of modifying the cases themselves⁵, we are modifying the user's perceptions and understanding. This modification takes the form of a short text explaining how the story fits in the given context before the story is shown.

From the viewpoint of the arguing system being described here, bridging is essential to the argument's coherence. It is part of the job of an arguer to explain how the evidence presented is connected to the point being argued. Otherwise, the story presentation will seem cryptic and incoherent. The retrieval shown in the Introduction is a good example of this. The first story, “Brown on violence,” makes many points. In addition to advocating violence to achieve civil rights, Brown comments on American culture, and threatens other black leaders. The first story retrieved as a counterargument, “Birmingham Riots,” is fairly long, and is not entirely about violence as a plan for achieving civil rights. It is also about courage, morality, withstanding violence, brotherhood and so on. The stories are clearly related, and a viewer might understand the transition, but the intended relationship between the stories is much clearer when the second story is preceded by an explanation such as the following:

⁴Since the part of the Archive that is most thoroughly indexed is the Civil Rights area, the majority of stories that are found will be from the domain of the US Civil Rights struggle in the 1960's. When the Archive has grown, we might expect stories about Gandhi, Sitting Bull and Spartacus to be retrieved as well: Whether the program should prefer more similar or more distant analogies is an ongoing research question.

⁵Due to the fact that all our materials are digitized video images, we have little ability to modify the cases.

In the story you have just viewed, H. Rap Brown makes the point that violence is a good plan because it achieves a good end, racial equality. In the story you are about to see, Martin Luther King makes the point that violence is a bad plan because violence is counter to Christianity, which is a good theme.

Before the story, "Segregation Forever," the bridge would be:

In the story you have just viewed, H. Rap Brown makes the point that violence is a good plan because it achieves a good end, racial equality. In the story you are about to see, George Wallace makes the point that racial equality is a bad goal, undermining H. Rap Brown's argument.

These bridges are assembled at run time, out of pieces of text associated with the point types, objects and templates. Taking the last bridge as an example, we have associated with the story point type *Plan is good* the string, "\$speaker makes the point that \$plan is a good plan to achieve \$goal, because \$instantiated-template." This is known as a half-bridge. Half-bridges are combined according to information associated with the type of opposition being used to oppose the point — in this case, undermining. The "boilerplate" for undermining is, "In the story you just viewed, \$half-bridge-1. In the story you are about to see, \$half-bridge-2, undermining \$speaker1's argument." We have not yet built a large set of bridging strategies, but the problem appears to be quite tractable.

Conclusion

Our story-point type hierarchy currently contains about 25 point types, although we expect this number to grow dramatically. Some of these types have as many as ten rhetorical templates, and many have only one. We are now in the process of indexing new stories on paper to determine which types and templates need to be added to our hierarchy. We have been encouraged by the fact that we have been able to index many new stories without needing to add many new items to our hierarchy.

We recognize, of course, that a system such as ours necessarily misses many of the more subtle points of a story. For example, in the story "Birmingham Riots," King argues that violence harms the aggressor more than the victim, and that it is better to be beaten and have a clear conscience than to beat someone else and be guilty. There are many more interesting conclusions that are suggested in King's speech that we lack the mechanisms to deal with. Given its limitations, however, the model exhibits rather sophisticated behavior with a minimum of knowledge-intensive programming.

One trend at ILS (as opposed, say, to the CYC project [Lenat and Guha, 1990]) has been to move away from large, inference-intensive AI reasoning systems,

with their detailed and complex representations. The emphasis in projects like this one is on building AI systems that require a less fine-grained knowledge representation. We believe this is essential if we are to meet our goals of scalability. Our goal is not to build a system that functions very well on a small set of interesting examples. Rather, we are trying to create a system that can exhibit interesting behavior over a broad range of cases, drawn from a wide variety of domains.

References

- [Alvarado *et al.*, 1986] S.J. Alvarado, M.G. Dyer, and M. Flowers. Editorial comprehension in oped through argument units. In *Proc. AAAI-86*, pages 250–256, Philadelphia, PA, August 1986. AAAI.
- [Bareiss and Slator, 1992] Ray Bareiss and Brian M. Slator. From protos to orca: Reflections on a unified approach to knowledge representation, categorization, and learning. Technical Report #20, Institute for the Learning Sciences, Northwestern University Evanston, IL, January 1992.
- [Bareiss *et al.*, 1991] Ray Bareiss, William Ferguson, and Andy Fano. The story archive: A memory for case-based tutoring. In *Proceedings of the 1991 DARPA Case-Based Reasoning Workshop*. DARPA, 1991.
- [Birnbaum, 1982] Lawrence Birnbaum. Argument molecules: A functional representation of argument structure. In *Proceedings, AAAI 1982*, pages 63–65, Philadelphia, PA, August 1982. AAAI.
- [Ferguson *et al.*, in press] William Ferguson, Ray Bareiss, Lawrence Birnbaum, and Richard Osgood. Ask systems: An approach to the realization of story-based teachers. In *The Journal of the Learning Sciences*, in press.
- [Flowers *et al.*, 1982] Margot Flowers, Rod McGuire, and Lawrence Birnbaum. Adversary arguments and the logic of personal attacks. In Wendy G. Lehnert and Martin H. Ringle, editors, *Strategies for Natural Language Processing*. Lawrence Erlbaum Associates, Morristown, NJ, 1982.
- [Kass, 1990] Alex M. Kass. Developing creative hypotheses by adapting explanations. Technical Report #6, The Institute for the Learning Sciences, Northwestern University Evanston, IL, November 1990.
- [Lenat and Guha, 1990] D. Lenat and R. V. Guha. *Building Large Knowledge-based Systems*. Addison-Wesley, Reading, PA, 1990.
- [McGuire *et al.*, 1981] R. McGuire, Lawrence Birnbaum, and M. Flowers. Opportunistic processing in arguments. In *Proceedings of the Seventh IJCAI, Vancouver, B.C.*, pages 58–60, 1981.

- [Oren *et al.*, 1990] T. Oren, G. Salomon, and K. Kretzman. Guides: Characterizing the interface. In B. Laurel, editor, *The Art of Human-Computer Interface Design*. Addison-Wesley, Reading, PA, 1990.
- [Rosch, 1988] Eleanor Rosch. Principles of categorization. In A. Collins and E. Smith, editors, *Readings in cognitive Science: a perspective from Psychology and Artificial Intelligence*, pages 312-322. Morgan Kaufman, San Mateo, CA, 1988.
- [Schank and Osgood, 1990] Roger Schank and Richard Osgood. A content theory of memory indexing. Technical Report # 2, The Institute for the Learning Sciences, Northwestern University Evanston, IL, March 1990.
- [Schank, 1977] Roger C. Schank. Rules and topics in conversation. *Cognitive Science*, 1:421-441, 1977.