Defining Salience in Case-Based Arguments*

Kevin D. Ashley
29 Westvicw Terrace
Easthaiiipfcon, Massachusetts 01027

Abstract

A largely unrealized goal of AI has been to design systems that can tailor descriptions of knowledge base objects to suit alternative points of view and contexts. This problem is examined here with respect to case based argu merits, arguments in which past cases are cited as justifications for a decision. We define the features of a case that are salient in describing it from various viewpoints and in a variety of argument contexts employed in HYPO, an adversarial case based reasoning program. We demonstrate HYPO's ability to slant case descriptions from opposing viewpoints and to tailor the descriptions to fit contextual roles in an argument, including filing the case, distinguishing it, posing it as a counterexample or as the target of a hypothetical.

1 Introduction

A recurring theme of A I has been to design systems that can tailor descriptions of knowledge base objects to suit alternative points of view and contexts. In 1975, for example, Minsk}' proposed frame representation schemes, in part, as a way of efficiently handling changes in viewpoint, either an observers physical perspective on a room or his metaphorical perspective in a nonvisual domain [Minsk v 7f>]. By 1995, researchers on the "Knoe sphere" project speculated that their comprehensive, interactive encyclopedia would be able, at the reader's choice, to describe the same basic reference information from a variety of stylized viewpoints such as "Fmpha sizeTheory", "NoMath," or "Ideology" [Penat 82].

As this perennial theme illustrates, one hallmark of intelligence is knowing what aspects of an object are salient from a particular viewpoint. Salience means strikingness or emphasis. As viewpoints change, different features of an object stand out.

I lie goal has been to design systems in which the context guides selecting which features to emphasize while, at the same time, minimizing the computational complexity of the selection process. Some progress has been made. Research in explanation has focused on slanting descriptive explanations to reflect the questioner's goals as evidenced by the question's perspective [McKeown 85]. Ordering objects in a scene according to "visual" salience has been explored as a guide to planning natural language descriptions of the scene [McDonald 82].

Determining salience is especially important in making arguments. Given the reality that most positions have weaknesses as well as strengths, a debater must know what facts help the argument or hurt it, what facts should be emphasized, mentioned or simply ignored. Attorneys are notoriously adept at describing (some- might say twisting) the same facts from opposing vantage points to support their clients' position. For example, a city commissioner hires the daughter of the judge who presides over the commissioner's lover's divorce. The prosecutor calls it a bribe; defense counsel calls it a coincidence. Beside characterizing the facts differently, the attorneys must decide which facts to emphasize (e.g., did the daughter have relevant employment experience or had she never before been able to hold down a job.)

A number of AI researchers have approached the problem oT salience in arguments by identifying various structures to elucidate support and attack relations and guide inference, (e.g., argument graphs and "molecules" [Birnbaum 82], argument trees [Cohen 83], and conversational moves [ReichmanAdar 84].) In general, the techniques have not been computationally powerful enough to generate realistic domain arguments.

Case-based reasoning techniques allow salience decisions to be made in the context of a more limited kind of argument, case based arguments, arguments in which prior cases are cited as justifications for deciding a problem situation. Salience has special significance in case-based reasoning where cases are used either as models for problem solving, examples for explanation, or argument justifications. For instance, Kolodner and Simpson

'Currently, the author is a Visiting Scientist at the IBM Thomas .1. Watson Research Center, P.O. Pox 218, York town Heights, New York 10598. Formerly, the author was a postdoctoral researcher and graduate student at the Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts 01003. This work was supported in part by: the Advanced Research Projects Agency of the Department of Defense, monitored by the Of fice of Naval Research under contract no. N00014 84 K 0017; the University Research Initiative, award no. N000M 86 K 0764; and an IBM Graduate Student Fellowship.

demonstrated how prior cases can focus a problem solver on those features of the problem situation that are salient for its solution [Kolodner 85].

Interestingly, the reverse is also true: when a past case is to be used in an argument about how to decide a problem, the features of the case that are salient depend on the particular facts of the problem situation, the viewpoint of the arguer and the argument context in which it is used.

A case-based reasoning program, HYPO, describes objects in a case knowledge base (''CKB") in a contextual way so that differing views of a case are supported and reasoned with and decisions about salience are made without complex inference. HYPO is an *adversarial* case-based reasoning program described in detail in [Ashley 87a]. It makes arguments how to decide a problem by citing relevantly similar past cases as justifications. In the past we have described the ways in which HYPO selects the most relevant cases and employs them in arguments [Ashley 87b,Ashley 88].

Here we focus on the narrower issue oT how HYPOfmakes determinations of salience, that is, how it determines what to say about a case that it cites in an argument. Since HYPO draws factual analogies between a cited case and the problem, one aspect of salience is determining what factual similarities (or differences) are relevant to the analogy. Another aspect of salience, how ever, is determining which of the relevant similarities and differences to emphasize in a given context and how to characterize them to convey that emphasis.

2 Views and Contexts

In HYPO, determining salience of features depends on the viewpoint and context of the description of the case.

HYPO supports two points of view in an argument corresponding to the two warring sides in a lawsuit: the plaintiff and the defendant. The plaintiff complains that the defendant's actions are illegal and caused injuries to the plaintiff for which it seeks relief From the court. HYPO deals with lawsuits involving trade secrets claims in which the plaintiff complains that the defendant has gained an unfair competitive advantage by obtaining il legally plaintiff's commercial sec rets.

HYPO describes cases in four basic contexts common to case-based arguments in law, political decision making, evaluating property and ordinary discourse (See dis cussion in [Ashley 87a].) Here we describe the contexts in terms of legal arguments:

J. Cited Case to make a point in a legal argument: HYPO makes an argument that a side should win the lawsuit described in the problem situation by citing a past case (the cited case) as an authority. In other words, HYPO adopts the viewpoint of, let us say, the plaintiff, argues that the plaintiff should win because the problem is similar to the cited case where a plaintiff won and emphasizes the relevant factual similarities.

2. Distinguished Case in response to a point in a legal argument: HYPO contests the argument point by distinguishing the cited case. That is, HYPO adopts the view of the opponent, the defendant, and asserts that the differences between the problem and the cited case are so great that the defendant should win. In this context, HYPO emphasizes the relevant factual differences between the problem and the cited case, which we will refer to as the distinguished case.

3. Counterexample to a cited case: Another way that HYPO contests an argument point is by citing a past case as a counterexample to the cited case. That is, taking the view of the defendant, HYPO cites as a counterexample a case won by a defendant. HYPO emphasizes the features of the counterexample that counteract the force of the plaintiffs argument based on the cited case. As described below, there are four kinds of counterexamples.

1. Target Case that motivates a hypothetical modi fixation: HYPO suggests how to strengthen a side's argument by suggesting hypothetical variations of the problem. Taking the view of either plaintifT or defendant, HYPO suggests how to bolster the argu ment by focusing on new features that would make the problem more similar to a past case (the target case) which could then be cited in support of that side.

I he program computes the most salient features to emphasize in describing a case from the CKB when it is used in any of the above contexts as a cited case, distinguished case, counterexample or target case. The features that are salient in any of these contexts can be described in succinct computational terms. In order to present those definitions, we define relevant similarities and introduce a bit of formalism.

3 Defining Relevant Similarities

In HYPO, the features of a case that matter are defined in terms of factors. A factor is a general collection of facts that tends to favor or hurt the plaintiff's argument in a particular kind of lawsuit. Factors have magnitudes, that is, a particular case may be a more or less extreme example of a factor. Factors are represented in HYPO with dimensions [Ashley 87a]. Each dimension has a set of prerequisites used to test if the factor applies to a case and a range of values over which the magnitude of the factor in a particular case may vary. Figure I shows examples of three factors in the trade secrets domain.

f^{1} = Sccrcts-Disclosed-Outsiders (Helps defendant *(6)* to extent that plaintiff has disclosed secrets to outsiders)

f^{2} = Outsider-Disclosures-Restricted (Helps plaintiff (PIE) if all disclosures to outsiders are confidential)

f_3 = Disclosure-in-Negotiations (Helps defendant *(6)* if plaintiff disclosed secrets to defendant in negotiations)

Figure 1: Three Sample Factors

Crown Industries
Outcome: Defendant won
Factors: { $f^1 f^3$ }
Magnitudes:
$M(f^1_{Crown})$ = disclosures to 7 outsiders

Midland Ross
Outcome: Defendant won
Factors: { f^1 }
Magnitudes:
$M(f^1_{MidlandRoss})$ = disclosures to 100 outsiders

Data General
Outcome: Plaintiff won
Factors: { $f^1 f^2$ }
Magnitudes:
$M(f^1_{DataGeneral})$ = disclosures to 6000 outsiders
$M(f^2_{DataGeneral})$ = all disclosures restricted

Crown Industries – Variant
Outcome: Defendant won
Factors: { $f^1 f^2 f^3$ }
Magnitudes:
$M(f^1_{Crown-var})$ = disclosures to 7 outsiders
$M(f^2_{Crown-var})$ = all disclosures restricted

Figure 2: Four Sample Cases

Cases in the database are represented in HYPO cs sentially as historical collections of factors, each with a particular magnitude, to which some authoritative de cision maker (i.e., the judge) has assigned an outcome, either for the plain tiff or the defendant. Figure 2 shows examples of four cases. Fach case is described abstractly in terms of an outcome, a set of applicable factors, and the magnitudes of certain of those factors.[1]

Relevant similarities among cases are defined computationally in terms of sets of shared factors. The set of relevant similarities, $5_{Ci,C2}$, between two cases, (\ and c_2, is the intersection of the sets of factors that apply to each case. In symbols:

$$S_{c_1,c_2} = F_{c_1} \cap F_{c_2} \quad (1)$$

Intuitively, the relevant similarities are reasons why the two cases should be decided the same way. Since the same combination of factors applied to each, the cases presented similar strengths and weaknesses from plaintiff's viewpoint.

Relevant differences are defined computationally in terms of factors not shared and differences in magnitudes of shared factors. Assume that the outcome of c_2 was in favor of plaintiff. The set of relevant differences between c_1 and c_2, $D_{c1,c2}$» is the union of three sets, the pro-defendant factors that apply only to C_1, the pro-plain tiff factors that apply only to c_2 and the shared factors which favor plaintiff more strongly in c_2 than in c\. In symbols:

$$\begin{aligned} D_{c_1,c_2} \equiv & \\ & [(F_{c_1} \setminus F_{c_2}) \cap F_-] \\ & \cup [(F_{c_2} \setminus F_{c_1}) \cap F_+] \\ & \cup \{f \mid (f \in S_{c_1,c_2}) \wedge (M(f_{c_2}) > M(f_{c_1}))\} \end{aligned} \quad (2)$$

Intuitively, the relevant differences are reasons why the two cases should be decided differently. All of these relevant differences make c_1 a weaker case for plaintiff than c_2. As a result, c_2 is a weaker justification in a legal argument that plaintiff should win in c\.

4 Selecting Most Relevant Cases

In HYPO, a case is analogous or "on point" to a problem if it shares some relevant similarities to the problem, P, (i.e., $S_{p,c_i} \neq \emptyset$.) One case, c_1, is "more on point" than another case, C_k, relative to a problem if the set of relevant similarities between the problem and c_i, is a superset of the set of relevant similarities between the problem and c_k (i.e., $S_{p,c_i} \supset S_{p,c_k}$.) The set of on point cases that are most relevantly similar to the problem, that is, the set of "most on point" cases, *MOP,* is defined as the set of all cases, C_i, such that for each c,, there is no case that is more on point than c,. To put that another way, *MOP* is the set of all cases, c,, such that c, is on point and for all other on point cases c_k, either c_i and C_k have no similarities in common with the problem situation or C_i, is as or more on point than c_k, or neither is more on point than the other.

Intuitively, *MOP* contains the cases that are candidates for the best cases to cite for either side in an argument how to decide the problem. Of all the cases in the CKB, they are closest in that they share with the problem the greatest overlap of factual strengths and weaknesses as represented by factors. For a given problem situation, having identified the factors that apply to it, HYPO computes *MOP* using the usual set operations.

5 Defining Salience within Contexts

With these concepts, one may succintly describe the features of a case that are salient when it is used in any of the four argument contexts. To illustrate how the salience expressions work, Figure 3 presents HYPO's descriptions of a single case, the *Data General* case of Figure 2, used in all four contexts, as a cited case, distinguished case, counterexample and target.

Cited Case: The salient features of a cited case (cc) are its relevant similarities to the problem situation. More specifically, the set of relevant similarities

[1] The following notation will make it possible to speak more succinctly about factors and cases. Let:

$F_+ \equiv$ the set of all factors, f, such that f generally favors plaintiff.

$F_- \equiv \{f \mid f$ generally hurts plaintiff $\}$.

$F_{c_i} \equiv \{f \mid f$ applies to $c_i\}$, (i.e., the set of all factors, f such that f applies to case c_i.)

For a given factor f^k :

$M(f^k_{c_i}) \equiv$ the magnitude of factor f^k in case c_i.

"$M(f^k_{c_i}) > M(f^k_{c_j})$" means that the magnitude of factor f^k is greater for the plaintiff in case c_i than in case c_j.

" \ " $\equiv$ set difference.

1. As Case Cited by Plaintiff:
   Where: Even though: Plaintiff disclosed its product information to outsiders, plaintiff should win a claim for trade secrets misappropriation.
   Cite: Data General
2. As Case Distinguished by Defendant:
   Data General is distinguishable because: In Crown Industries, Plaintiff disclosed its product information in negotiations with defendant. Not so in Data General. In Data General, plaintiff's disclosures to outsiders were restricted. Not so in Crown Indus tries.
3. As Counterexamples Cited by Plaintiff:

   Trumping:
   Data General is more on point [than Midland Ross] and held for plaintiff where it was also the case that: plaintifFs disclosures to outsiders were restricted.

   Boundary:
   Data General held for plaintiff even though in Data General plaintiff disclosed its product information to more outsiders than in Midland Ross.

A. As Target Case for Plaintiff:
Plaintiff's response would he strengthened if: plaintiff's disclosures to outsiders were restricted. Cf. Data General.

Figure 3: Descriptions of *Data General* Case in Four Argument Contexts

between *cc* and *p* is the intersection of the sets of factors that apply to each (i.e., $S_{p,cc} = F_p \cap F_{cc}$. See equation 1.) Figure 3, No. 1 shows how HYPO describes the *Data General* case as cited case (cc), where the problem situation (*p*) is that of the *Crown Industries* case of Figure 2.

HYPO's description of *Data General* as a cited case in No. 1 emphasizes the factor shared with the problem (f^1 involving disclosures) and ignores the unshared factor (f^2 involving restrictions on disclosures.) HYPO computes the salient features using expression 1:

$$S_{Crown,DataGeneral} = \{f^1 f^3\} \cap \{f^1 f^2\} = \{f^1\}$$

Intuitively, this makes sense. To call attention to f^2 at this point would not be tactically sound since f^2 is not a relevant similarity but a relevant difference (i.e., it is a pro-plaintiff factor that *Data General* does not share with the problem.) It is up to the defendant to point out the differences (which HYPO does as we see below.)

Note that HYPO prefaces the description in No. 1 with the word "Eventhough" in recognition of the fact that although f^1 is a relevant similarity, plaintiff's disclosures to outsiders are not normally reasons for deciding in plaintiff's favor. Had there been some relevant similarities that favored the plaintiff, for example, that the plaintiff had taken substantial precautions to secure its secrets, HYPO would have stated them first. That is, HYPO would say, "Where: plaintiff took substantial security precautions, Fventhough: plaintiff disclosed its product information to outsiders, plaintiff should win

Distinguished Case: The salient features in distinguishing a case are the relevant differences between the problem and the distinguished case, $D_{p,dc}$. As Figure 3, No. 2 shows, when HYPO describes the *Data General* case as distinguished case (de), where the problem situation *(p)* again is the Crown case, it focuses on factors that *Data General docs* not share with *Crown* (J^7 involving restrictions on disclosures and f^3 involving disclosures in negotiation.) Making the substitutions in definition (2) of relevant differences, the three terms are:

$$(F_{Crown} \setminus F_{DataGeneral}) \cap F_- = \{f^1 f^3\} \setminus \{f^1 f^2\} \cap F_- = \{f^3\}$$

$$(F_{DataGeneral} \setminus F_{Crown}) \cap F_+ = \{f^1 f^2\} \setminus \{f^1 f^3\} \cap F_+ = \{f^2\}$$

$\{\}$ $(M(f^1_{DataGeneral}) \not> M(f^1_{Crown})$ (6000 disclosures is worse, not better, for plaintiff than 7 disclosures.)

Taking the union of the terms yields the set of salient differences between the *Crown* case and *Data General*, $D_{Crown,DataGeneral}$:

$$\{f^3\} \cup \{f^2\} \cup \{\} = \{f^3 f^2\}$$

HYPO argues that these differences not only account for the outcome in *Data General but* also show that the *Crown* case need not have the same result.

Note that the analysis allows HYPO to select those dif fere rices in factor magnitudes that are salient in a given context. Although the magnitudes of *Data General* and *Crown* along f^1 are different, HYPO sensibly does not call attention to this difference. The difference does not help, indeed it hurts, the defendant's argument. *Crown* is much stronger for plaintiff than *Data General in* terms of the numbers of disclosures to outsiders.

The analysis also allows HYPO to point out those factors that are salient in a context because they do *not* apply to a case, for example, in the context of distin guishing *Data General, f^3* is significant because it does not apply to *Data General.* The importance of this "nonfeature" becomes apparent only in light of the process of distinguishing *Data General* from *Crown.*

Counterexample: HYPO supports four kinds of counterexamples, each with a somewhat different definition of the features of a case that are salient when citing it as a counterexample. Each kind of counterexample involves a case that is "contrary" to the cited case in the sense that it had the opposite outcome, focuses on different salient features, and can be used to disparage the impact of the salient features of the cited case. The four kinds of counterexamples and the associated salient features are discussed below:

(1) Trumping: a contrary case that has more (in the sense of a superset of) relevant similarities to the problem than the cited case. This kind of counterexample

"trumps" the cited case. It has all the cited case's similarities and then some. HYPO has a method for computing the set of trumping counterexamples, $TCX_{s,cc}$, that can be cited by a side, s, in response to a case, cc, cited by s's opponent. $TCX_{s,cc}$ is the set of all most on point cases, c_i, such that the outcome of c_i favors s and c_i is more on point than cc.

The salient features, the ones that HYPO will emphasize in describing a trumping counterexample, are the extra similarities, that it shares with the problem that the cited case does not, specifically, the set difference of the set of relevant similarities between p and ccx and that between p and cc, (i.e., $S_{p,ccx} \setminus S_{p,cc}$ where $S_{p,ccx} \supset S_{p,cc}$.)

(2) Partial: a contrary case that has the same set or some nonempty subset of similarities to the problem as the cited case. A partial counterexample shows, some what more weakly than a trumping counterexample, that the shared similarities do not always lead to the same outcome as the cited case. The set of all cases, $PCX_{s,cc}$, that a side s can cite as partial counterexamples to a cited case cc is the set of all cases, c_i, such that the outcome of c_i favors s and c_i shares some of the same factors with the problem as cc.

The salient features of a partial counterexample are the similarities that both counterexample and cited case have to the problem, namely, the intersection of the set of relevant similarities between p and ccx and that between p and cc, (i.e., $S_{p,ccx} \cap S_{p,cc}$ where $S_{p,ccx} \subseteq S_{p,cc}$.)

(3) Boundary: a contrary case which is an extreme example of some factor, f^i, that both is a relevant similarity between the problem and the cited case and favored the winner of the cited case. The set of all cases, $BCX_{s,cc}$, that a side s can cite as boundary counterexamples to a cited case cc is the set of all on point cases, c_j, such that the outcome of c_j favors s and there is some factor, f^i, such that c_j and cc share f^i with the problem, and the magnitude of f^i in c_j is greater in favor of s's opponent than in cc.

For a given boundary counterexample, ccx, the salient features are the factor, f^i, of which the counterexample is a more extreme example than the cited case and the relative magnitudes of f^i in the counterexample and the cited case. Specifically, the salient features are: (Assume that defendant won the cited case.)

$$\{f^i \mid [f^i \in F_-] \wedge [f^i \in S_{p,cc}] \wedge [M(f^i_{cc}) > M(f^i_{ccx})]\} \quad (3)$$

By emphasizing that the plaintiff won despite that fact that the counterexample was so much worse for the plaintiff along factor f^i than the cited case, a boundary counterexample tends to show that f^i is not a very significant factor. (Like human arguers, HYPO uses counterexamples and distinguishing to deal symbolically with the problem of weighting competing factors. See [Ashley 88].)

(4) Potentially trumping: a contrary case that would be a good counterexample with which to trump the cited case *if* the problem had certain additional factors. The set of all cases, $PTCX_{s,cc}$, that a side s can cite as partially trumping counterexamples to a cited case cc is the set of all cases, c_j, such that the outcome of c_j favors s and there is some hypothetical variant of the problem, *pvar*, such that c_j is more on point with respect to *pvar* than cc. In HYPO, we have identified a number of heuristics for hypothetically modifying a problem situation to generate *pvar*. See [Rissland 86].

The salient features of a potentially trumping counterexample are the extra similarities between it and the hypothetical variant of the problem (*pvar*), namely the set difference of the set of relevant similarities between *pvar* and ccx and that between p and cc, (i.e., $S_{pvar,ccx} \setminus S_{p,cc}$ where $S_{pvar,ccx} \supset S_{p,cc}$.) These salient features are the difference between the problem situation and a hypothetically modified problem in which a side's arsenal of cases to cite is strengthened by the addition of a new trumping counterexample.

As Figure 3 illustrates, HYPO describes a case differently depending on how it is used as a counterexample. No. 3, shows how HYPO describes the *Data General* case as a trumping and boundary counterexample (ccx). Here, the problem situation (p) is a variant of the *Crown* case in which all 7 disclosures to outsiders were subject to confidentiality restrictions (i.e., $F_{Crown-var} = \{f^1 f^2\}$ and $M(f^1_{Crown-var}) = 7$) and the case cited for the defendant (cc) is the *Midland Ross* case of Figure 2.

In the first sentence of No. 3, HYPO cites *Data General* as a counterexample to trump *Midland Ross* and emphasizes the extra similarities that led to the opposite result in *Data General*: namely f^2. This comes from the above definition of a trumping counterexample's salient features where $\{f^1 f^2\} \setminus \{f^1\} = \{f^2\}$.

In the second sentence of No. 3, HYPO cites *Data General* as a boundary counterexample, emphasizing the fact that the plaintiff won in *Data General* eventhough it was a much worse case than *Crown* in terms of disclosures to outsiders (6000 outsiders versus 7.) When used as a boundary counterexample, the salient feature of *Data General* is f^1, because, making the substitutions in formula (3):

$$[f^1 \in F_-] \wedge [f^1 \in S_{Crown-var,Midland}] \wedge [M(f^1_{Midland}) > M(f^1_{DataGeneral})].$$

Target Case: The salient features of a case that is a target (t) for a hypothetical variation of another case to be modified (*cmod*) are those factors in $F_t \setminus F_{cmod}$ that are added, or those factors in $S_{cmod,t}$ whose magnitudes are changed, to make the hypothetical more like or more extreme than the target. As Figure 3, No. 4, shows, when HYPO describes the *Data General* case as a target for a hypothetical modification of the *Crown* case, it emphasizes the extra factor (f_2 dealing with restrictions on disclosures) which, if added to *Crown*, would make *Data General* a counterexample with which to trump *Midland Ross* ([Rissland 86] describes more completely HYPO's posing of hypotheticals.)

To summarize, HYPO generates a variety of descriptions for any given case in its knowledge base depending on the viewpoint and context of the argument it makes. As the descriptions of the *Data General* case in Figure 3 illustrate, HYPO emphasizes various aspects

of the case, sometimes focusing on similarities, sometimes differences, sometimes factor magnitudes, sometimes factors significant because they do not apply. Each of HYPO's thirteen implemented factors has an associated "canned" phrase for describing it. HYPO composes the phrases corresponding to the salient factors into the brief case summaries. In this way, HYPO tailors descriptions of the thirty cases in its CKB to fit the argument contexts, viewpoints and problem situations. The descriptions compare favorably with those found in court opinions citing the same cases. Sec [Ashley 87a]. Although HYPO's text generation scheme is primitive, the salience information could be used as inputs to a so phisticated natural language generator for planning more elaborately tailored case descriptions.

6 Conclusion

In this paper, the problem of designing a system that can tailor descriptions of knowledge base objects to suit alternative viewpoints and contexts has been examined for a limited kind of discourse: case-based arguments. Determining the features of a past case that are salient when the case is cited as an argument justification depends on the particular facts of the problem and of the other cases in the knowledge base, the viewpoint of the arguer and the contextual role of the case in the argu ment.

HYPO, an adversarial case-based reasoning program, tailors its descriptions of a case to emphasize those features, specifically factors, that arc salient depending on viewpoint and context. HYPO describes cases in arguments from two conflicting viewpoints, that of plaintiff or defendant, and in four contextual roles, as cited case, distinguished case, counterexample or target case. HYPO makes further decisions about salience depending on the kind of counterexample a case is used as: trumping, partial, boundary, or potential trump. For each of the contexts, we have defined the salient features and illustrated how the definitions apply to a past case described in each context.

References

[Ashley 87a] Ashley, Kevin I). *Modelling Legal Argument: Reasoning with Cases and Hypothetical*.* PhD thesis, Department of Computer and Information Science, University of Massachusetts, 1987. COINS Technical Report No. 88 01. To be published by The Mil Press.

[Ashley 87b] Ashley, Kevin D, and Rissland, Edwina E. Compare and Contrast, A Test of Expertise. In *Proceedings AAAI 81,* American Association Tor Artificial Intelligence. August, 1987. Seattle.

[Ashley 88] Ashley, Kevin D. and Rissland, Edwina E. Waiting on Weighting: A Symbolic Least Commitment Approach. In *Proceedings AAAI 88,* American Association for Artificial Intelligence. August, 1988. Minneapolis.

[Birnbaurn 82] Birnbaum, L. Argument Molecules: A Functional Representation of Argument Structure. In *Proceedings AAA] 82,* American Association for Artificial Intelligence. Pittsburgh, PA, August, 1982.

[Cohen 83] Cohen, Robin. *Computational Model for the Analysis of Arguments.* Technical Report OSRG-151, Department of Computer Science, University of Toronto, 1983.

[Kolodner 85] Kolodner, Janet L., Simpson, Robert L., and Sycara-Cyranski, Katia. A Process Model of Case-Based Reasoning in Problem Solving. In *Proceedings IJCAI 85,* International Joint Conferences on Artificial Intelligence, Inc. Los Angeles, CA, August, 1985.

[Lenat 82] Lenat, 1). B., Boming, A., McDonald, D., Taylor, C, and Weyer, S. Knoesphere: Building Expert Systems with Encyclopedic Knowledge. In *Proceedings IJCAI 82,* International Joint Conferences on Artificial Intelligence, Inc., pages 167 169. Karlsruhe, August, 1982.

[McDonald 82] McDonald, D. 1). and Conklin, E. J. Salience as a Simplifying Metaphor for Natural Language Generation. In *Proceedings AAAI 82,* American Association for Artificial Intelligence, pages 75 78. Pittsburgh, PA, August, 1982.

[McKeown 85] McKeown, Kathleen R., Wish, Myron, and Matthews, Kevin. Tailoring Explanations for the User. In *Proceedings IJCAI 85,* International Joint Conferences on Artificial Intelligence, Inc. Eos Angeles, CA, August, 1985.

[Minsky 75] Minsky, Marvin. A Framework for Representing Knowledge. In Winston, Patrick (editor), *The. Psychology of Computer Vision.* McGraw-Hill, New York, 1975.

[ReichmanAdar 84] Reichman-Adar, R. Extended Person-Machine Interface. *Artificial Intelligence* 22:157 218, 1981.

[Rissland 86] Rissland, Edwina L. and Ashley, Kevin D. Hypothetical as Heuristic Device. In *Proceedings AAAI 86,* American Association for Artificial Intelligence. August, 1986. Philadelphia, PA.