Distinguishing – A Reasoner's Wedge [1]

Kevin D. Ashley [2]
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003

Telephone: (413) 545-0332
CSNET address: ASHLEY@UMass

Abstract

In this paper we focus on the *Distinguisher's Wedge*, an intellectual tool for responding to an argument that two cases are alike by asserting reasons why they are different and why the differences matter. We characterize the wedge as involving a search for *distinctions*, factual differences between the cases that tie into justifications for treating them differently. We show how the wedge can be modelled computationally in a Case-Based Reasoning ("CBR") system using *precedential* justifications and describe how the model is realized in our *HYPO* program which performs legal reasoning in the domain of trade secret law. Legal argument, with its emphasis on citing and distinguishing precedents and lack of a strong domain model, is an excellent domain for studying the wedge. We show how HYPO uses "dimensions", "case-analysis-record" and "claim lattice" mechanisms to cite and distinguish real cases and suggest how the model may be extended to cover more sophisticated kinds of distinguishing.

1. Introduction

Distinguishing is making explicit what is different about two cases and why that difference justifies *not* treating them the same way. In everyday argument, we are often called upon to distinguish cases. For example, a twelve year old, whose birthday is next month, demands that his over-protective parents tell him why *he* can't go to the movies to see "Little Shop of Horrors" but his fifteen year old sister can. The parents' task is to *distinguish* the brother's case from the sister's. The short answer, as if it were enough, is, "Because your sister is three years older." But, one might ask, or – if one is twelve, demand – why *that* difference makes a difference, a question to which there are many dubious responses, among them: "Your sister is more mature"; "The movie gets out too late – it's past your bedtime."; "The movie is rated PG-13 – it's for teenagers."..., to which there are at least as many indubitable retorts: "But I *am* a teenager! I'm going to be 13 next month."; or "I could see it if it was on a VCR but not at the movies!"; or "Noah's parents

[1]This work was supported (in part) by: the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract no. N00014-84-K-0017, and an IBM Graduate Student Fellowship.

let *him* go see it." Dubious as the responses are, there are plenty of other differences that meet the issue even less: "Your sister is a girl and you're a boy."; "Your sister has more money than you do." (One can imagine the issues for which *these* differences would make a difference.); and some differences that appear to favor the opposite conclusion: "Your sister has more homework than you do."

Distinguishing is so fundamental it is like one of mechanics' simple machines, like a reasoner's *wedge*, a tool that uses the point of the factual difference to pry the two cases apart. One hears this kind of argument in many diverse domains: in arguments about mathematics (If figure 6 is a polyhedron then why isn't figure 7? [Lakatos, 1976]); in scientific research (If two persons were exposed to the same virus why did only one get the disease?); in historical political analysis (America's Nicaragua policy is leading to another Vietnam.); And, of course, in legal arguments [Levi, 1949]. In each domain, the distinctions (or lack of them) elucidated in the responses are the bedrock of effective analysis.

2. The Distinguisher's Wedge

Distinguishing involves searching for *distinctions*, factual differences between the two cases that tie into *justifications* for treating them differently. In its most general form, the wedge works like this:

> (1) A reasoner *must* respond to an assertion in an argument by analogy that since case α is just like case β, α should be treated in the same way as β;
>
> (2) Among the tools in the reasoner's kit is the wedge. In searching for distinctions, the would-be distinguisher has two alternatives that he pursues in parallel:
>
> 1. Search for factual differences with which to take advantage of known justifications. If possible, the reasoner will try to find out or infer new facts about α or β that constitute such differences.
> 2. Search for possible justifications with which to take advantage of known or possible factual differences between α and β.
>
> (3) Assuming that the search for distinctions is successful, the reasoner must still evaluate how useful the distinctions are as responses. By citing the distinction, the reasoner changes the state of the argument and may introduce new facts or justifications that will have to be defended.

The need to distinguish drives both factual inquiry and theory building. It prompts the distinguisher to investigate for new factual differences and to find or create theories that assign the desired significance to the differences. The parents' responses, for example, point out factual differences, in the childrens' ages, maturity and bedtimes, as well as theories why the differences matter, ranging from invoking the rule that movies are only for their rated age group to an implied theory about movies' effect on children. The factual difference and the justification are, of course, intimately related. The sex difference, for example, may not appear relevant except, arguably, in light of a theory that girls mature faster than boys.

A purported distinction can be disputed at a range of levels, including disagreements about the asserted factual differences, their significance, or the validity of the justifications. Among the ways

to respond to a distinction are: (1) to cite another case that minimizes the difference asserted by the distinguisher, in effect granting that the difference does matter (e.g., citing the example of twelve year old Noah with the permissive parents. One can almost hear the parent's next distinction: "Just because Noah's parents allow him to do certain things does not mean ..."); (2) to assert that the justification applies to α just as much as it applies to β, (e.g., "I am a teenager!"), in effect granting that the justification is valid but questioning the meaning of the justification. This is an example of how the wedge puts pressure on a definition; (3) to grant the difference but show that it actually leads to the opposite of the distinguisher's conclusion. For example, the supposed difference in maturity may imply that the sister has greater responsibility – like homework – and drive a factual query whether sister does not have a big report due tomorrow and shouldn't be going to the movies after all. In other words, the asserted distinction backfires; or (4) to cite a counter example to the justification (e.g., Noah saw the PG-13 movie despite being less than thirteen.)

There are many ways to evaluate who wins the battle over whether there is a distinction between the cases. Good ways to lose such arguments, at least in some domains, are to lead off with an example α for which the opponent can cite numerous equally analogous counter examples, or to assert a distinction that can be turned against the distinguisher's position. Good arguers take these evaluation criteria into account in strategically planning their arguments.

3. Computationally Modelling the Wedge – a Case-Based Approach

One major problem for computationally modelling the wedge is controlling the inferencing necessary to find a distinction between cases α and β. The justification part of a distinction may be an explanatory chain of inferences in terms of some causal theory, for example the "maturity" theory of the initial example, or more simply, invocation of the rule that PG-13 movies are only for teenagers, whose conclusion applies to β but not, because of some factual difference, to α. In light of the theory, that difference is crucial and worthy of a distinction. *But how does a system find the difference and the theory?* There may be many differences between α and β, not all of which are relevant to the issue posed. The relevance of even a crucial difference may not become apparent except after a possibly long chain of inferences, and it may "cut" the wrong way, that is, lead to a chain of inferences that hurts the distinguisher's position. There may also be many theories that would lead to the conclusion that that α and β should be treated differently. Which one should the system try? How far should it backchain before deciding that there is no crucial difference between α and β along the lines of that theory?

Another problem is in dealing with the "open texture" of the predicates in the rules representing the theory. It is not possible simply to define the predicates and hope that there's an end to it. The boy's response is quite reasonable. Depending on the context, "Teenager" sometimes means "$\geq$ 13" but sometimes it may also mean "of the level of maturity of at least an average 13 year old." Add this wrench into the machinery of controlling inferencing and things get really messy.

A case-based approach to modelling the wedge takes a short cut across the problem of controlling inferencing. A central element of a CBR approach is the use of *precedential justifications* instead of chains of inferences to simplify the control problem. A CBR system assumes that the facts that a prior case had a certain cluster of features, and that the decision of the prior case was made because of some of those features and inspite of others, are a basis for a precedential justification for coming to the same conclusion in a future case with a similar combination of features.

Consistently with its use of precedential justifications, CBR assumes that when a decision is

made about a case, the decider assigns credit or blame to some of the case's factual features as either contributing in favor of or against the decision. In effect, the decision of a case: (1) *Selects* certain features that are important enough for purposes of credit assignment; (2) *Clusters* the selected features; and (3) "Weights" them, ranking the features in the cluster that favor the decision higher than those against it, at least in that case. The factual similarities and differences among cases that are important in the domain are the ones that previous cases have found to be important; they are predefined for the system and used as the basis for *indexing* the cases in the CKB. In other words, the cases in the CKB are indexed by the same features that are involved in credit assignment and precedential justifications.

In essence, the CBR approach controls the inferencing problem by flattening out the depth of inferencing needed to come up with precedential justifications. Although the depth of inferencing is shallow, the depth of analysis is not by virtue of the breadth of the index, the size and diversity of the CKB, and the mechanism for selecting the best cases to use as precedents.

4. The HYPO Program

HYPO is a CBR program that reasons about a fact situation by critically comparing it to precedent cases. Its domain is trade secret law [Rissland & Ashley, 1986; Rissland & Ashley, 1987; Ashley & Rissland, 1987]. The law is an excellent domain for studying the wedge, since distinguishing and the use of precedential justifications are primary components of legal argument [Levi, 1949].

The main sources of legal knowlege in HYPO are contained in HYPO's CKB and its library of dimensions. Dimensions represent the legal relationship between various clusters of operative facts and the legal conclusions they support or undermine. Dimensions provide not only indices into lines of cases but a scale for comparing cases in terms of important factual differences that affect the strength, or weakness, of a fact situation with respect to that line of reasoning. For instance, one line of trade secret cases focusses on the degree to which the "cat (i.e., secret) has been let out of the bag", even by the complaining plaintiff, himself: that is, how many disclosures of the putative secret were there and of what kind? This way of looking at a trade secret case (captured by the *Disclose-Secrets* dimension) provides one approach to resolving a misappropriation dispute and was used in the *Data General* and *Yokana* cases discussed below. Another approach might emphasize the competitive advantage gained by the defendant at the plaintiff's expense or the switching of a key employee from the plaintiff to the defendant [Rissland & Ashley, 1986]. Each dimension has: *prerequisites*, expressed in terms of *factual predicates*, that tell whether a dimension applies to a case or not; *focal slots* that single out the particular facts making a case stronger or weaker along the dimension and range information that tells how a change in the focal slot affects that strength (e.g., for *Disclose-Secrets*, the focal slot is the number of disclosees. Increasing that number weakens the plaintiff's position.) See generally [Ashley, 1986].

5. HYPO's Model of Distinguishing

HYPO's model of distinguishing uses *dimensions*, *case-analysis-records* and *claim lattices* to find both relevant factual differences and precedential justifications. Dimensions provide HYPO's handle on factual differences and a means for extrapolating from prior cases. By virtue of dimensions' definitions, HYPO knows that a particular difference is significant and which side, π or δ, the

difference favors. Dimensions also provide a precedential justification for why the difference matters – because the difference represents either a strengthening or weakening of features that mattered in prior cases indexed by the dimension. The case-analysis-record and claim lattices enable HYPO to compare cases in terms of multiple dimensions' cumulative effects as well as along individual dimensions. The differences they deal with represent either a strengthening or weakening of the *closeness of the analogy* between a prior case and the cfs.

Here is how the model works. *First*, in analyzing a new cfs, HYPO runs through the library of dimensions and produces a **case-analysis-record** that contains: (1) applicable factual predicates; (2) applicable dimensions; (3) near-miss dimensions; (4) potential claims and (5) relevant cases from the CKB. **Near-miss** dimensions are those for which some, but not all, of the prerequisites are satisfied. The combined list of applicable and near-miss dimensions is called the **D-list**. Figure 1 describes a cfs based, for purposes of illustration, on *Data General v. Digital Computer Controls, Inc.*, a real case in the CKB (but with one difference: the cfs involves 12000 disclosures of plaintiff's secret while *Data General* involved only 6000.) Figure 2 shows the case-analysis-record for the cfs.

Second, HYPO uses the case-analysis-record to construct the **claim lattice**, which is a lattice such that: (1) the root is the cfs together with its D-list; and (2) successor nodes contain pointers to cases that share a subset, usually proper, of the dimensions in the cfs's D-list. Figure 3 shows the claim lattice actually generated by the HYPO program for analyzing the cfs of Figures 1 from the viewpoint of a trade secrets misappropriation claim [Ashley & Rissland, 1987]. (There is a separate claim lattice for each possible claim.)

The ordering scheme enables claim lattices to capture a sense of closeness to the cfs of cases in the CKB. Those sharing more dimensions are nearer to the cfs. Those nodes closest to the root whose subsets of the cfs's D-list do not contain near-miss dimensions can be considered most-on-point-cases "**mopc's**" to the cfs; leaf nodes are the least-on-point. All of the cases displayed are relevant to the cfs because they all share some legally important strengths or weaknesses with the fact situation as represented by the dimensions shared with the cfs.

Third, HYPO uses the cases in the claim lattice to make and respond to arguments by analogy about the cfs citing those cases as precedents. Different major branches of the lattice indicate different ways to argue the case, effectively one way for each group of mopc's. HYPO can argue the case for *side 1*, let us say the plaintiff ("π") in the cfs, by *citing* a pro-π mopc. As it happens, in Figure 3, the only mopc is the pro-π *Data General* case, so HYPO cites it in favor of the plaintiff:

[a] $\hookrightarrow$ For *Side 1*: (point) π wins.
Case Cited: *Data General*
(π in cfs and cited case disclosed secrets
but disclosures subject to restriction.)

Since mopc's share the most legally important strengths and weaknesses with the cfs (i.e. mopc's are the closest analogies to the cfs), *Data General* is the most persuasive case HYPO could cite for the defendant as *side 1*. (For purposes of illustration *Data General* is also the basis of the cfs in Figure 1. It is encouraging that after analyzing the cfs, HYPO has found the nearly indentical case to be most on point!) *Telex v. IBM*, for example, is not a mopc because, although it is very close to the root, the *Competitive-Advantage* dimension which applies to *Telex*, and which would help π if it applied to the cfs, is only a near-miss for the cfs. (Note that *Competitive-Advantage* is *'d in Figure 3.) In support of the citation, HYPO draws the analogy between the mopc and the cfs. The relevantly similar facts are just those summarized by the dimension[s] that apply to both

(i.e., *Disclosed-Secrets* and *Restricted-Disclose*, the latter capturing the idea that if the disclosees agree to maintain confidentiality, then the secret is safe.)

HYPO responds to points like [a] by ***distinguishing*** the cited case using three basic methods: (1) Comparing the strengths of cfs and cited case along the dimensions they share in common; (2) Finding strengths or weaknesses, represented by dimensions, that cfs and cited case do *not* share. (3) Finding other cases that are more on point than the cited case.

Responses [b], [d] and [f] below illustrate these methods. As an example of the **first** method, consider how HYPO distinguishes *Data-General* from the cfs, on behalf of *side 2*, the defendant, by comparing values of the focal slots of the shared dimension:

[b] $\hookleftarrow$ For *Side 2*: (response to [a])
Case Distinguished by δ: *Data General*
(π in case disclosed to 6000 outsiders; '
6000 $\ll$ 12000 disclosures in cfs.)

HYPO knows from the claim lattice and the range information about the *Disclose-Secrets* dimension, that *Data General* presents a stronger case because π disclosed the confidential information to fewer outsiders.

The **second** method involves focussing on facts associated with ***unshared*** dimensions that helped a party (π or δ) in the cited case or hurt the corresponding party in the cfs. Suppose the *Telex* case were cited on behalf of plaintiff as *Side 1*:

[c] $\hookrightarrow$ For *Side 1*: (point) π wins.
Case Cited: *Telex*.

In Figure 3, *Telex* is a *potential mopc*, that is, a case close to the root, some of whose applicable dimensions are only near-misses with respect to the cfs. HYPO distinguishes *Telex* on behalf of defendant as *Side 2* by pointing out the facts that help π in *Telex* and that are associated with dimensions that are either near-misses (*Competitive-Advantage*) or inapplicable (*Bribe-Employee*) to the cfs:

[d] $\hookleftarrow$ For *Side 2*: (response to [c])
Case Distinguished by δ: *Telex*
(π in case gained competitive advantage;
δ in case bribed π's employees.)

The **third** method of distinguishing involves finding a case, favorable to the responding side, whose *overall factual difference* from the cfs is less than that of the cited case. A mopc, for example, distinguishes all pro-opponent cases in any successor nodes of the claim lattice because it is more on point. (A mopc does not distinguish opponent's cases in other branches of the lattice – that would be like comparing apples and oranges.) If one were starting the argument from the other side and cited the *Midland-Ross* and *Yokana* cases on behalf of defendant as *Side 1*:

[e] $\hookrightarrow$ For *Side 1*: (point) δ wins.
Cases Cited: *Midland-Ross*, *Yokana*
(π in cfs and cited cases disclosed secrets to outsiders.),

HYPO responds by pointing out the pro-π strength that *Data General* shares with the cfs but which is missing from *Midland-Ross* and *Yokana*, namely that the disclosures were subject to restrictions to maintain confidentiality (a feature captured by the *Restricted-Disclose* dimension that applies to the cfs and *Data General* but not to *Midland-Ross* or *Yokana*). In other words, HYPO's response for plaintiff as *side 2* is:

[f] $\hookleftarrow$ For *Side 2*: (response to [e])
Cases Distinguished by π: *Midland-Ross*, *Yokana*
(Although π in cases and cfs made disclosures,
disclosures in cases were not on confidential basis;
Data General is more on point.)

HYPO uses the knowledge of how a case may be distinguished to prompt the user to find new factual differences and new justifications that would strengthen or weaken the argument. For example, HYPO uses the distinction between the *Telex* case and the cfs elucidated in response [d] to pose a hypothetical variant of the cfs in which the factual difference between the two is reduced [Rissland & Ashley, 1986; Ashley & Rissland, 1987]. The hypothetical prompts the user to investigate whether π in the cfs gained a competitive advantage or δ in the cfs bribed π's employees, in which case the *Telex* case would become a powerful precedent for the π that could not be so readily distinguished.

6. Evaluating HYPO's Model of the Distinguisher's Wedge

HYPO's model allows it to distinguish cases in a manner similar to what is actually done in court opinions in cases involving similar issues to our cfs. In *Mixing Equipment Co. v. Philadelphia Gear, Inc.*, 436 F.2d 1308, 1315 (3d Cir., 1971), the court distinguished the *Yokana* case as follows:

> [Another case] and *Midland-Ross Corp. v. Yokana*, 293 F. 2d 411 (3 Cir. 1961) cited by appellants are inapposite. They involve situations in which restrictive covenants had not been utilized by the former employer.

In *Data General Corp. v. Digital Computer Controls, Inc.*, 357 A.2d 105, 109 (Del. Ch., 1975), the real case on which the cfs is based, the court took pains to point out, with respect to the drawings that had been disclosed to customers:

> Such drawings bore a proprietary notice or legend, and the machine itself was accompanied by a ...confidentiality agreement limiting the use of such drawings to maintenance, as opposed to manufacture, which was stated to be forbidden without plaintiff's consent in writing.

In *National Rejectors, Inc. v. Trieman*, 409 S.W. 2d 1, 40-42 (Sup. Ct. Mo., 1966) the court said:

> [W]e do find some significant parallels between the facts of this case and those of *Midland-Ross Corporation v. Yokana*, (D.C. N.J.), 185 F Supp. 594. ...What was lacking in *Yokana* as in this case, was any evidence that, prior to defendant's competition,

plaintiff considered the information which Yokana sought to use trade secrets. The court pointed out that plaintiff's blueprints in *Midland-Ross* were furnished plaintiff's suppliers and customers and potential customers. The court found an absence of precautions on the part of plaintiff to keep secret information regarding its machines.

7. Extending HYPO's Model of the Wedge

HYPO's model of distinguishing does not yet support all of the kinds of disputes that attorneys have about justifications. Attorneys argue about the significance of a court's holding in a prior case in more abstract terms, using predicates that obtain wide currency in the analysis of particular claims (e.g., what attorneys would call the *elements* of a claim, generalized statements that purport to define the necessary requirements of a claim.) Consider, for example, alternative ways for defendant's attorney to state the response in [b]:

> *Data General* does not help the plaintiff. With 12000 disclosures, plaintiff doesn't *have* a secret to protect anymore, regardless of whether the disclosees agreed to maintain confidentiality or not.

More interestingly, the attorney might address the court:

> If you hold in favor of a plaintiff who has disclosed its "secret" to 12000 outsiders then, regardless of whether they have agreed to maintain confidentiality or not, you are effectively doing away with the requirement imposed in *Midland Ross* and *Yokana* that trade secrets be secret.

This kind of response (reminiscent of the "I *am* a teenager!" response) is common enough in the law. HYPO avoids representing detailed definitions of predicates used in justifications like "secret" so as not to compound the inference control problem. Compare this to the classic rule-based approach of Waterman and Peterson [1981] who would use ever more refined rules to define legal predicates. As Gardner points out, there is no way in jurisprudence to specify logical definitions from which it is possible to deduce whether the predicate is satisfied [Gardner, 1984]. Gardner proposed to use cases to resolve "hard" issues about the meanings of predicates, but her main approach and her "cases" were very rule-like.

HYPO's model does, however, point to a middle ground for providing predicates with a kind of operational meaning, not in terms of a logical definition, but in terms of the boundaries of fact situations to which the predicate has been held to apply or not and knowledge of how those boundaries may be stretched. HYPO does know, for example, that *Data General* is an extreme case for plaintiffs in some sense. It is a *boundary* case for plaintiff along the *Disclose-Secrets* dimension, the weakest case along the dimension that the plaintiff still won. A boundary case may or may not be as on point as a mopc, but it is still useful as a precedent. A boundary case may be used to convince a court that it is not making new law by showing that, even if the court *were* to rule for the plaintiff in the cfs, the cfs would still not be the *worst* case, at least in some sense, that a plaintiff has won.

Of course, the meaning of a predicate in a justification like "secret" is not one-dimensional but then neither is the meaning of the claim of trade secrets misappropriation. HYPO represents the latter by associating various claims with clusters of dimensions. In a sense, the cases indexed by the

dimensions scope out the boundaries of fact situations that have or have not been deemed to present winning trade secrets claims. In a similar way, HYPO could "tag" certain important predicates such as the elements of a claim, with cases that scope out the boundaries of fact situations that have or have not been deemed to satisfy the predicate. HYPO would then know what kinds of factual changes effect not only the strength of the claim, but also the meaning of the predicates. Whatever a "secret" is, HYPO would know what factual circumstances would make something more or less of one.

8. Conclusion

In this paper we have described the Distinguisher's Wedge, a tool for responding to an argument that two cases are alike by asserting distinctions, that is, factual differences and justifications why the differences matter. We have shown how a computational model of the wedge can be incorporated into a Case-Based Reasoning (CBR) system that makes and responds to arguments using precedential justifications. We have described how the HYPO program realizes this model and have suggested an extension of the model useful in generating more sophisticated disputes about the terms of justifications.

9. Figures

Plaintiff Data General (π), who developed and marketed the Nova 1200 minicomputer, complained that defendant Digitial (δ) developed a competing minicomputer, the D–116, by misappropriating π's trade secrets. Specifically π complained that δ copied π's drawings of the design of the Nova 1200 and used them to design the substantially identical D–116. The drawings appeared in a maintenance manual that π distributed to 12,000 customers who purchased the Nova 1200. The drawings and manual contained a legend that prohibited their copying except by written permission of π.

Figure 1: Current Fact Situation (cfs) based on *Data General Corp. v. Digital Computer Controls, Inc.*

Applicable Factual Predicates:
exists-corporate-claimant, exists-confidential-info, exists-disclosures ...

Applicable Dimensions: *Disclose-Secrets*, *Restricted-Disclose*

Near-Miss Dimensions:
Competitive-Advantage, *Vertical-Knowledge*

Potential Claims: Trade Secrets Misappropriation

Relevant CKB cites: See claim lattice, Figure 3

Figure 2: Case-Analysis-Record for CFS

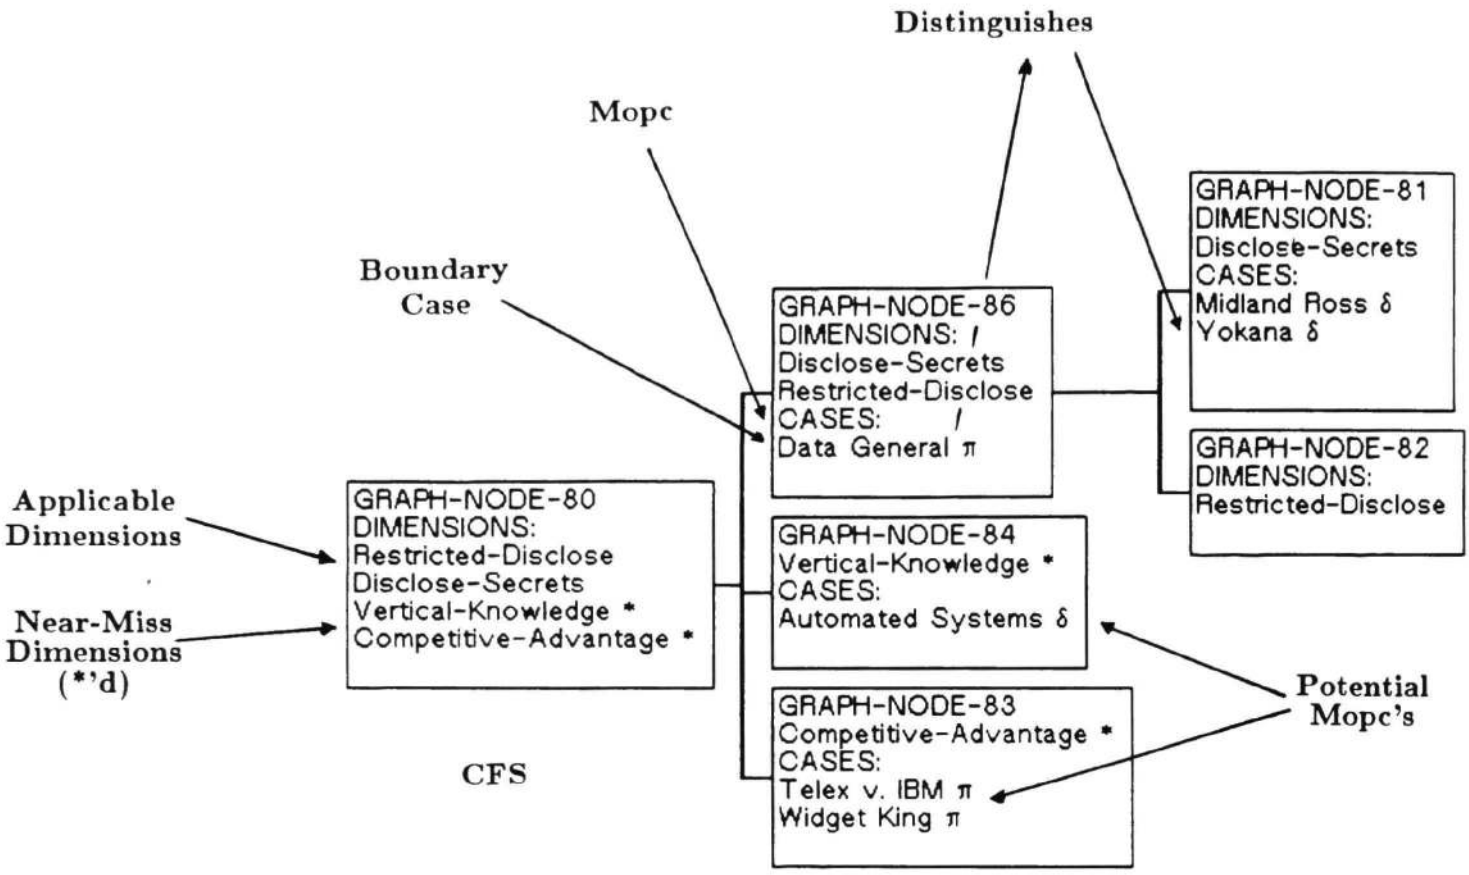

The root node represents the cfs and its D-list. (Dimensions that are near-misses for *cfs* have *'s.) Successor nodes contain pro-plaintiff (π) or pro-defendant (δ) cases, involving trade secrets misappropiation claims, that are on point to cfs. Nodes closest to root that do not have near-miss dimensions contain **mopc**'s; otherwise they may contain **potential mopc**'s. Leaf nodes are least-on-point. Each major branch of lattice that contains mopc's represents one way of arguing the cfs. Mopc's **distinguish** cases in successor nodes. **Boundary** cases are examples of extremes along particular dimensions. Hypothetical **hybrid mopc**'s combine features of different mopc's that hold for π and δ. Potential mopc's suggest fruitful hypothetical variants of cfs.

Figure 3: A Claim Lattice.

References

[1] Kevin D. Ashley. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals - A Thesis Proposal.* Project Memo 10, The COUNSELOR Project, Department of Computer and Information Science, University of Massachusetts, 1986.

[2] Kevin D. Ashley and Edwina L. Rissland. Creating Neighborhoods of Cases, Projections Through a Case Space. Submitted: IJCAI-87, 1987.

[3] A. vdL. Gardner. *An Artificial Intelligence Approach to Legal Reasoning.* PhD thesis, Department of Computer Science, Stanford University, 1984.

[4] I. Lakatos. *Proofs and Refutations.* Cambridge University Press, London, 1976.

[5] Edward H. Levi. *An Introduction to Legal Reasoning.* University of Chicago Press, 1949.

[6] Edwina L. Rissland and Kevin D. Ashley. HYPO: A Case-Based Reasoning System. Submitted: IJCAI-87, 1987.

[7] Edwina L. Rissland and Kevin D. Ashley. Hypotheticals as Heuristic Device. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, American Association for Artificial Intelligence, August 1986. Philadelphia, PA.

[8] D. A. Waterman and M. Peterson. *Models of Legal Decisionmaking.* Technical Report R-2717-1CJ, The Rand Corporation, Santa Monica, CA, 1981.