An Internal Contradiction of Case-Based Reasoning[1]

David B. Skalak[2]
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003
SKALAK@cs.umass.edu

Abstract

In a case-based reasoning system, one simple approach to assessment of similarity of cases to a given problem situation is to create a linear ordering of the cases by similarity according to each relevant domain factor. Using Arrow's Impossibility Theorem, a result from social welfare economics, a paradox is uncovered in the attempt to find a consistent overall ordering of cases by similarity that satisfactorily reflects these individual rankings. The implications of the paradox for case-based reasoning are considered.

1. Introduction

1.1 Framework for this Paper

This paper attempts to demonstrate that one underlying model of case-based reasoning ("CBR") is internally contradictory. This model, called the "Simple Model" here, assumes that cases can be ranked in order of their similarity to a current problem situation. The conclusion to be drawn from this work is that the model of CBR used by a system must take care to avoid the assumptions that give rise to the contradiction, or must use an entirely different approach, such as one based on numeric weighting of cases, which has its own well-trodden pitfalls.

The path to this result is somewhat circuitous, for it involves translating a result, Arrow's Impossibility Theorem[3], from social welfare economics to CBR. But first as background, case-based reasoning is briefly reviewed, and the problematic model is then presented. An example of a possible CBR contradiction is given.

The first leg of this circuitous path is the presentation of Arrow's Theorem in its original form. In broad scope, the Theorem yields a contradiction that arises in the following circumstances. A group of people rank their preferences for some set of candidates or goods or some other resource. The goal is to derive an overall ranking for the group of individuals as a whole that consistently reflects the individual rankings. Certain apparently reasonable conditions are placed on the compromise ordering and how it should reflect the individual rankings. The Theorem's conclusion is that it is impossible within certain limitations to create a preference procedure to derive the "social" ordering from the individual rankings.

Arrow's original, surprising result is phrased in terms of an "individual" "preferring" one "social state" to another. To establish the contradiction for CBR, it

[1]This work was supported in part by the National Science Foundation, contract IRI-8908481, the Office of Naval Research under a University Research Initiative Grant, contract N00014-87-K-0238, and a grant from GTE Laboratories, Inc., Waltham, Mass.

[3]Arrow originally referred to the theorem as the General Possibility Theorem (Arrow 1963). Consistent with common usage, I have called it the "Impossibility" Theorem.

will remain to show that Arrow's Theorem can be duly translated to the realm of CBR and the Simple Model. The term "individual" in the Theorem is translated into "domain factor", "preference" into "ranks as more similar to the current problem case", and "social state" into "case". Translated in this way, the conditions of Arrow's Theorem are shown to apply. Thus the conclusion of the Theorem is yielded: that there exists no function to achieve a consistent overall ranking of cases by each individual factor that appropriately reflects the individual rankings. After this result is established, its implications for CBR are discussed.

1.2 Case-Based Reasoning

Case-Based Reasoning is using past problem-solving episodes to analyze or solve a new problem. If necessary, a previous analysis or solution will be adapted to reflect the differences between the new and old cases.

CBR involves several steps, including: (1) accept a new experience and analyze it in order to retrieve relevant cases from case memory; (2) order the retrieved cases to select a set of "best" cases from which to craft a solution or analysis of the problem case; (3) derive the solution by modifying that of the most similar case or formulate an analysis of the problem case in view of the most "on-point" cases; and (4) optionally, store the newly solved or interpreted case in case memory, adjusting indices into memory as appropriate.

1.3 The Relevance of this Result to CBR

When confronted with a situation where a variety of conflicting factors are at work, one may appeal to CBR. In this circumstance, it is difficult to assign blame or credit for the results of a case to a single factor or even to cluster of them. Cases may be viewed from conflicting vantage points that can point to different results (Ashley 1989b). When viewed from the perspective of one factor, one case may seem to be the most similar; another factor may point to another case as being more salient. Or, in a mixed paradigm system or blackboard architecture that uses a variety of knowledge sources to suggest lists of similar cases, those knowledge sources will probably yield different rankings of cases by similarity. *Cf.* (Rissland and Skalak 1989). According to Arrow's Theorem, duly translated, no collective ranking process that meets its preconditions can take the linear orderings of cases by the individual knowledge sources and combine them into a consistent collective similarity ordering.

In the variety of CBR sometimes called "problem-solving" CBR, one case is usually selected as the most similar, and its solution (plan, diagnosis, *etc.*) is modified to solve the original problem posed. If modifying that case turns into a dead end, the next most similar case may be processed and its solution modified. In the variety of CBR referred to as "interpretive" or "precedent-based" CBR, cases are usually presented in an ordering according to their similarity to the current problem case. Both varieties of CBR require that cases be ranked by similarity in some fashion. That different rankings are available to the system, either by virtue of distinct knowledge sources or contrasting factor vantage points, necessitates collecting these rankings into a single ranking. Applying Arrow's Theorem to a simple CBR model suggests that this attempt to construct an overall ordering may involve internal contradictions. Either no procedure will be able to create such a ranking, or we shall have to give up some of the requirements normally placed on the individual rankings and how the collective ranking is to reflect them.

2. The Simple CBR Model

2.1 Description of the Model

Consider the following simplified model of CBR, here called the "Simple Model" . We are given a set of cases, a set of relevant factors, and a problem case called the "current fact situation". The relevant factors have been pre-designated to determine case similarity and to serve as indices for retrieval of similar cases from the case knowledge base. In this model, the goal of CBR is to rank the cases in a linear ordering in view of their "overall similarity" to the current fact situation.

The Simple Model's apparent solution is to begin by linearly ordering the cases according to similarity with respect to each individual factor. "Ties" are permitted: two or more cases may be equally similar according to a factor. Once these factor-by-factor rankings are assembled, one could try to combine each factor's ranking of cases into a single ordering by overall similarity. I attempt to demonstrate that it is not possible to produce this overall ranking within certain apparently reasonable constraints.

The objection may be made that this Simple Model is unrealistic --- no current CBR system in practice employs it. Several responses to this objection may be made. (1) Elements of the Simple Model are present in some systems, as discussed in Section 5.2. (2) If the Simple Model is not used in its pure form, it may be useful to say why this obvious approach has been avoided. (3) The relevance of the Simple Model may be seen more realistically in the context of a blackboard implementation of CBR, where a number of knowledge sources are available to rank cases. It is likely that individual knowledge sources would suggest contrasting similarity rankings of cases. One would think therefore not in terms of factors ranking cases, but in terms of knowledge sources (that may themselves consider a variety of factors) ranking cases. (4) The factor-by-factor approach is found elsewhere in artificial intelligence. For a recent example see the analysis of the Independent Credit Assignment assumption by Subramanian and Feigenbaum (1986).

2.2 An Example of a Similarity Contradiction

Suppose we have the following three cases and a current fact situation ("CFS") dealing with whether the law permits a driver to overtake an automobile traveling in the same direction. Let us suppose also we have identified three factors as relevant:

1. type of pavement center line (with values none, dotted, single, and double),
2. visibility (poor, fair, good), and
3. distance from position of car to oncoming car in the opposite lane (in feet).

Case	Line	Visibility	Distance	Result
CFS	single	good	800	?
A	single	poor	900	No
B	dotted	good	1000	Yes
C	none	fair	850	Yes

Then, from the point of view of the --

line factor, **case A** is more similar to the **CFS** than **case B**, and **case B** is more similar to the **CFS** than **case C** (write: ApB & BpC, where the relation p may be thought of as "is preferred to");

visibility factor, **case B** is more similar than **case C**, and **case C** is more similar than **case A** (BpC & CpA);

distance factor, **case C** is more similar than **case A**, and **case A** is more similar than **case B** (CpA & ApB).

A majority of the factors --- two out of three in each instance --- indicate that **case A** is more similar than **case B** and that **case B** is more similar than **case C** (ApB & BpC). By transitivity of the preference relation, a majority of the factors therefore would say that **case A** is more similar than **case C** (ApC). However, a majority of the factors (visibility and distance) actually rank **case C** as more similar than **case A** (CpA). This is a contradiction. This theoretical contradiction presents a practical quandary for a driver's selection of the most similar case in view of the opposite results of **case A** and **case C.** In **case A** passing is not permitted; in **case C**, it is.

3. Arrow's Impossibility Theorem

The above case-based example is actually a reflection of a voting paradox known for over a century (Nanson 1882). Paradoxes like this one have given rise to research by Kenneth Arrow to investigate the conditions under which the ranking by individuals of preferences for "social choices" (for example, for political candidates) may fail to yield a satisfactory "social" (collective) ranking that reflects the individual preference orderings.

Two conditions are placed on the notion of a preference to ensure its accordance with its common meaning. A *preference* (a) must be transitive, and (b) must permit the comparison of any two choices; that is, for any two choices x and y, (x is preferred to y) or (y is preferred to x), or we are indifferent as between x and y (Arrow 1963).

Arrow's Impossibility Theorem on the collective ranking of social choices by individuals is stated in terms of five conditions that might reasonably be placed on a social preference ordering. These conditions are quoted below directly from (Mueller 1979) citing (Vickrey 1960) and are more transparent than Arrow's original renderings. The exception may be Condition IV, which is stated in a more plausible way by Arrow's original work and I have used that formulation[4] (Arrow 1963). The Theorem denies the existence of a *social welfare function*, defined as "a process or rule which, for each set of individual orderings $R_1,...,R_n$ for alternative social states (one ordering for each individual), states a corresponding social ordering of alternative social states, R". (Arrow 1963, p.23). A *social state* is a catch-all term that would include such diverse persons and objects as political candidates, governmental forms, welfare entitlement schemes and access to certain goods.

<u>The five Preconditions to Arrow's Impossibility Theorem:</u>

I. Unanimity (The Pareto Postulate). If an individual preference is unopposed by any other contrary preference of any other individual, this preference is preserved in the social ordering.

II. Nondictatorship. No individual enjoys a position such that whenever he expresses a preference between any two alternatives and all other individuals express the opposite preference, his preference is always preserved in the social ordering.

III. Transitivity. The social welfare function gives a consistent ordering of all feasible alternatives. [That is, the social ranking must be transitive.]

IV. Range. (Unrestricted domain). Among all the alternatives there is a set S of three alternatives such that, for any set of individual orderings $T_1,...,$ T_n of the alternatives in S, there is an admissible set of individual orderings $R_1,..., R_n$ of all the alternatives such that, for each individual i, x R_i y if and only if x T_i y for x and y in S.

---

[4]Mueller's statement of Condition IV posits the existence of a "universal" alternative, which assumption he admits is not crucial (Mueller 1979, p.186).

V. Independence of Irrelevant Alternatives. The social choice between any two alternatives must depend on the orderings of the individuals over only these two alternatives, and not on their orderings over any other alternatives.

Arrow's Impossibility Theorem:
No social welfare function satisfies conditions I through V.

The reader is referred to (Arrow 1963) and (Mueller 1979) for proofs of the Theorem.

Arrow's original rendering of the Range Condition may benefit from a little explication. It requires that "every logically possible set of individual orderings of a certain set S of 3 alternatives can be obtained from some admissible[5] set of individual orderings of all alternatives." (Arrow 1963, p.24). The assumption assures that the collective, overall ranking process is not performed in a way that necessarily rules out some possible collective orderings (Mueller 1979). The purpose of the condition is to ensure that the social welfare function is not biased against certain rankings.

4. The Application of Arrow's Theorem to the Simple CBR Model

To establish the result of this paper, it remains to translate Arrow's Theorem from the realm of individual preferences for social states and social welfare functions into the language of CBR: the language of domain factors, similarity, and case ranking procedures.

Let CFS be a given problem case. Instead of saying that an individual prefers one social state to another, say that a *factor prefers* Case 1 to Case 2 if considering that factor alone would lead one to say that Case 1 is "more similar" than Case 2 to the problem CFS. A factor may rank two (or more) cases as equally similar -- in the language of public choice, the factor is *indifferent to* them. This definition of preference may be seen to satisfy the above two criteria placed on a preference. The translation of "social welfare function" is to the procedure that ranks cases by overall similarity from the factors' similarity orderings.

It can be seen that each of the five conditions is reasonably imposed in the translated setting provided by the Simple CBR Model.

For example, the translated Unanimity Condition would state: if a factor's preference between two cases (*i.e.*, that factor's favoring one case as being more similar than another case to a problem situation) is unopposed by any other contrary preference of any other factor, this preference is preserved in the overall ranking of cases. If each factor says that one case is more similar than another to a problem situation, then the collective ordering must preserve this similarity rating.

The Nondictatorship Condition ensures that no factor has so much importance attached to it that it alone can veto the similarity assessment of the remaining factors.

The Transitivity Condition requires that the collective ranking of cases by similarity be transitive. (This condition appears reasonable at first glance, but may be suspect in some applications.)

The Range Condition is perhaps the most opaque, but seems to provide the "twist" that makes the proofs of the Theorem work. In the CBR realm, it requires that there are three cases (A, B and C) that have the property that in the overall ranking any of the relative similarity rankings of those cases is possible for some conceivable

---

5 An *admissible* set of orderings is a set that is in the domain of the social welfare function, the process that performs the collective ranking.

individual rankings of the cases. This condition requires that these three cases may be ordered in any way, as long as the individuals rankings require it. The upshot of the Range Condition is that the rankings are unbiased to the extent that they do not rule out the possibility of any relative ordering of the three cases. Any of case A, B or C may be the most on point; any may be the least on point. Their relative similarity depends solely on what the individual factors require.

The last condition, the Independence of Irrelevant Alternatives, can also be seen to reasonably apply to the CBR realm. The overall relative similarity of two cases depends only on each factor's ranking of those two cases and not the factor rankings of other cases. If case A is more similar overall than case B, that result does not depend on the individual factor rankings of any other case.

The power of the Arrow's result for public choice is partly due to the self-evident nature of its preconditions. The conditions retain their manifest attraction when translated to the CBR realm. Once one is assured that the five conditions are met, the Theorem's original proofs apply (duly translated), and we have the

Impossibility Result for Case-Based Reasoning: Under these assumptions, there is no process that performs the collective ranking of cases by similarity that satisfies Conditions I to V.

5. Limitations and Implications of the Result

5.1 Arrow's Assumption of No Cardinal Utilities

There is a somewhat obscured assumption to Arrow's results that is not contained in his statement of the Theorem, but is referred to elsewhere in (Arrow 1963). In his own work on social welfare functions, Arrow has assumed that one does not rely on "cardinal, interpersonal comparisons of utility" to obtain the social ranking (Arrow 1963, p.9). In other words, Arrow posits that one may not assign numerical values to gauge the strength of the individual preferences, and use a numerical social welfare function to yield a collective preference ranking for all the individuals. This assumption is rooted in his expressed desire to avoid the difficulties created by assigning numerical values to personal preferences and the conceptual ambiguity of comparing degrees of preference across individuals.

The upshot of Arrow's assumption for the Simple CBR Model is that the Theorem assumes that we are not using cardinal assessments of case similarity for each factor, and then applying an evaluation function to combine these numerical assessments of similarity into an overall ranking. One such constraint is a proscription upon using cardinal numbers to gauge the degree of each factor's preferences. Only ordinal preferences are permitted by the Simple Model. That is, it can be said whether one case is more similar to than another to the current fact situation, but no numerical value can be used to quantify the relative similarities of the cases to the current problem. In particular, this model would not permit using a numerical evaluation function that weights each factor's cardinal preference for a case and sums those weighted preferences into a real number that reflects the collective preference for that case.

Following Arrow's interdiction against cardinal utilities is reasonable in many CBR domains. Numerical weighting schemes for credit assignment may be avoided in view of a litany of conceptual and pragmatic objections against their use. See, *e.g.*, (Ashley 1989a), (Ashley and Rissland 1988), (Kolodner 1988). A number of these objections are collected in (Ashley 1990): (1) lack of cognitive validity for many domains and domain experts; (2) uncertainty and arbitrariness in weight assignment and lack of an authoritative source of numerical weights, with the consequent disagreement among experts as to specific weights; (3) the failure of static weights to reflect the context-dependent importance of factors; (4) pitfalls of a premature commitment to a weighting scheme; and (5) reduction of information into a single number entails a representation that loses information that might otherwise be used to advantage.

5.2 Three Weighting Models

Ashley (1989a) identifies three models of factor weighting that underlie current CBR systems: numerical, precedent-based, and heuristic. The consequences for each of these models of the Impossibility Result are discussed below.

One way for a CBR system to avoid the paradox of the theorem is to use numerical functions to assess case similarity. See, *e.g.*, (Kibler and Aha 1987). Numerical weights may be appealed to in domains where an analytic or statistically-based model is accessible. See (Ashley 1989a). In domains where numeric weighting of factors is reasonable, the Nondictatorship Condition of the Theorem may be violated, however. A factor with a sufficiently large weight may determine the outcome of the similarity ranking, regardless of the rankings by lesser weighted factors. However, an evaluation function that is used merely for its ordinal, and not its cardinal, properties --- used just to rank cases by similarity in a total order --- may be subject to the strictures of the Theorem nonetheless.

In some areas of application, models of any sort are simply not available, and one must appeal to previous cases to assess the relative importance of factors (Ashley and Rissland 1988). The lack of some domain models is apparently a classical impetus for the use and development of case-based reasoning techniques in the first place. In such applications of CBR, one must rely merely on the patterns of combinations of factors that have appeared in previous cases to perform credit assignment. It is difficult to identify the factor(s) that determine the results in such instances, and one must avoid embracing a ranking scheme that reduces to a voting arrangement. In such precedent-based schemes where no numeric or heuristic weighting is available, special care must be taken to avoid the siren song of the Theorem. The scheme used must violate at least one of the Theorem's preconditions, explicit or implicit, or be potentially subject to the paradox. In the work of Rissland and Ashley, for example, (a) cases are ordered by clusters of factors and not by individual factors, (b) that ordering is partial and not total, and (c) the Range Condition may not be satisfied.

The appeal to an analytic model provides an intermediate point between numeric weights and the assessment of the importance of factors based on precedent. Where a causal or other analytic model is available, it may be used to rank goals and the features that influence their attainment (Koton 1988), (Bareiss 1987), (Kolodner 1985), (Hammond 1987). See generally (Ashley 1989a). Where factors can be ranked in a preference ordering according to their importance, the Theorem's paradox is circumvented through the violation of the Nondictatorship Condition. The similarity ranking may be dominated by the most important feature. For example, Kolodner's PARADYME system avoids the paradox by using an ordered set of preference heuristics to choose the most useful cases from those retrieved from a previous partial matching (Kolodner 1989). In general, systems that use a model to rank goals (and the factors that influence goal attainment) may avoid the potential for internal contradiction by avoiding the Theorem through the violation of the Nondictatorship Condition.

6. Conclusion

CBR is inherently subject to a dilemma. On one hand, it is problematic to assign numerical values and weights to similarity preferences, and thereby hide the resulting "value" judgements in trading off disparate and fundamentally incommensurate factors. On the other hand, one must avoid embracing a symbolic or implicit weighting scheme that can be construed as a voting procedure of the sort here presented with its attendant paradoxes.

References

Arrow, K. J. (1963). Social Choice and Individual Values. Cowles Foundation for Research in Economics at Yale University. New Haven, Yale University Press.

Ashley, K. D. (1989a). "Assessing Similarities Among Cases." Proceedings of the Second DARPA Case-Based Reasoning Workshop, May 1989, Pensacola Beach, FL.

Ashley, K. D. (1989b). "Toward a Computational Theory of Arguing with Precedents: Accommodating Multiple Interpretations of Cases." In Proceedings of the Second Annual Conference on AI and Law, June 1989, Vancouver, B.C.

Ashley, K. D. (1990). Modelling Legal Argument: Reasoning with Cases and Hypotheticals. Cambridge, MA, MIT Press (in press).

Ashley, K. D. and Rissland, E.L. (1988). "Waiting on Weighting: A Symbolic Least Commitment Approach." Proceedings AAAI-88, American Association for Artificial Intelligence, August 1988, Minneapolis.

Bareiss, E. R., Porter, B.W., and Wier, C. C. (1987). "Protos: An Exemplar-Based Learning Apprentice." In Proceedings of the Fourth International Workshop on Machine Learning, pages 12--23, June 1987, University of California at Irvine.

Hammond, K. J. (1987). "Explaining and Repairing Plans that Fail". In Proceedings IJCAI-87, International Joint Conferences on Artificial Intelligence, Inc., August 1987, Milan.

Kibler, D. and Aha, D.W. (1987). "Learning Representative Exemplars of Concepts: An Initial Case Study." Proceedings of the Fourth International Workshop on Machine Learning, pages 24--30, June 1987, University of California at Irvine.

Kolodner, J. L. and Simpson, R. L., and Sycara-Cyranski, K. (1985). "A Process Model of Case-Based Reasoning in Problem Solving." In Proceedings IJCAI-85, International Joint Conferences on Artificial Intelligence, Inc., August 1985, Los Angeles.

Kolodner, J. L. (1988). "Retrieving Events from a Case Memory: A Parallel Implementation." Proceedings of the DARPA Case-Based Reasoning Workshop, May 1988, Clearwater Beach, FL.

Kolodner, J. L. (1989). "Judging Which is the 'Best' Case for a Case-Based Reasoner." Proceedings of the DARPA Case-Based Reasoning Workshop, May 1989, Pensacola Beach, FL.

Koton, P. (1988). "Reasoning about Evidence in Causal Explanations." In Proceedings of the First DARPA Case-Based Reasoning Workshop, May 1988, Clearwater Beach, FL.

Mueller, D. C. (1979). Public Choice. Cambridge Surveys of Economic Literature. Cambridge, England, Cambridge University Press.

Nanson, E. J. (1882). Transaction and Proceedings of the Royal Society of Victoria. **19**: 197-240.

Rissland, E. L. and Skalak, D. B. (1989). "Combining Case-Based and Rule-Based Reasoning: A Heuristic Approach." Proceedings IJCAI-89, International Joint Conferences on Artificial Intelligence, Inc., August 1989, Detroit.

Subramanian, D. and Feigenbaum, J. (1986). "Factorization in Experiment Generation." In Proceedings AAAI-86, American Association for Artificial Intelligence, August 1986, Philadelphia.

Vickrey, W. (1960). "Utility, Strategy, and Social Decision Rules." Quart. J. Econ. **74**(Nov.): 507-35.