Dynamic Construction of Decisions*

Adele Howe
Paul Cohen

Experimental Knowledge Systems Laboratory
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003 †

February 1987
COINS Technical Report 87-62
EKSL Report 87-05

Abstract

We present an algorithm for iteratively constructing decisions. It acquires knowledge about attributes only as needed to select among alternatives, and emphasizes qualitative comparisons of the alternatives on subsets of their attributes. Thus, it facilitates choices with relatively few data. It is a heuristic decision method that complements the more precise decision theoretic approaches. We have implemented the algorithm as the kernel of a prototype decision support system.

1 Introduction

Much decision-making literature focuses on combining evidence in support of decisions [Howard 66, Raiffa 70, Holst 71]. This paper presents an alternative approach, called ***constructive decision making***, that emphasizes the process of ***acquiring*** just the evidence required for a decision. The constructive decision making approach underlies a decision support system [Howe 86] that incrementally identifies the factors that influence a decision and moves from states of indecision to states in which a choice is clear. It characterizes the current state of the decision with respect to strength of support for each alternative. If no alternative is clearly superior, it seeks information about factors of the decision that could discriminate the alternatives. These include the relative support for alternatives on new attributes and the relative importance of the attributes. The process of acquiring information about attributes continues

*This research is funded by DARPA/RADC Contract F30602-85-C-0014, by National Science Foundation grant IST 8409623, and by University Research Initiative grant 86A1104. An earlier version of this paper was presented at the Phoenix Conference on Computers and Communications, February 1987.

†We are indebted to Carole Beal, Sabine Bergler, David Lewis, and Cynthia Loiselle for their comments on drafts of this paper.

until a clear choice emerges. Constructive decision making incrementally *evolves* decisions by simultaneously identifying the salient attributes of alternatives and assessing the support for alternatives on attributes.

Since constructive decision making is an algorithm for defining and resolving decisions, it is ideally suited to AI programs that make choices based on factors whose relevance cannot be anticipated until run-time. For example, programs with dynamic control structures define and resolve numerous control decisions based on factors whose relevance changes over time [Barr 81, Cohen 87]. Another example is decision support. Currently, a decision analyst defines decisions in terms of alternatives, outcomes, utilities, and probabilities; then a computer program resolves the decision based on this information. For a program to assume responsibility for defining decisions, in addition to resolving them, it would need to construct decisions by asking which alternatives are salient and which of their attributes contribute to a decision. This is what the constructive decision making approach is designed to do.

Constructive decision making complements decision analysis in three ways. First, it stresses the process of constructing the decision by acquiring information as needed. In contrast, decision analysis emphasizes analysis of alternatives given that their outcomes, utilities, and probabilities are already specified. Second, the constructive decision making approach is qualitative and supports decision-making with much less precise data. This can be an important distinction because payoffs, costs and probability assessments are often difficult to obtain [Raiffa 70][1]. However, the price for this advantage is that constructive decision making is heuristic, not normative. Third, constructive decision making facilitates graceful degradation of decision making because it permits some decisions based on relatively few data, but may render more confident decisions if additional data are available. Thus, it provides a qualitative equivalent to the decision-analytic concept of the utility of evidence.

2 Constructive Decision Making

The underlying principle of constructive decision making is that a decision is *constructed* by opportunistically adding attributes of the alternatives to a growing collection of support. It is helpful to view this process in terms of search: the decision under construction is in one of a set of *states*, defined by general characteristics of the collection of evidence. *Actions* are associated with each state to transform it into other states. The goal of constructive decision making is to transform a state that represents a decision in which a choice cannot be made into another in which a choice can be made with confidence. This view is motivated by psychological research that suggests that people search the decision space by conservatively considering attributes and alternatives of the decision – called "choice by feedback processing" [Svenson 79]. Given this, the first step in developing a model of constructive decision making is to identify the possible states and the actions that transform them.

[1] Moreover, according to psychological experiments the cognitive processing required to use subjective expected utility in complex tasks is beyond human capabilities [Biggs 83].

2.1 Decision States

We begin the discussion with simple 2-alternative, 2-attribute problems typified by the *apples and oranges* problem. When you compare apples and oranges in the grocery store, you may find one preferred on the basis of flavor and the other on the basis of quality. If you could combine these attributes on a single composite attribute, then the choice is often clear. But if flavor and quality are not easily combined, then the choice between them is problematic.

Decision alternatives are compared on their salient attributes.[2] A *decision state* in our model is a concise statement of the current combination of salient attributes, including how well the alternatives are distinguished on the available attributes and how important those attributes are. We identified five *dimensions* that describe decision states. In the descriptions that follow, alternatives are referred to as p and q, attributes as A_i and A_j, and values of attributes for specific alternatives as $A_i[p]$. The symbol $\tilde{>}$ indicates preference between two values.

$Sd[A_i]$. A ***significant difference*** on attribute A_i indicates that the values of the two alternatives on this attribute, independent of the rest of the attributes, are distinct. In other words, assuming that the values of all the other attributes for the two alternatives are equal, can a decision be based on this attribute alone? Assessing a value for $Sd[A_i]$ does not require exact values of the alternatives on an attribute, only the subjective judgment that the values are significantly different[3].

$$Sd[A_i] = \begin{cases} 1 & \text{if } A_i[p] \text{ and } A_i[q] \text{ are distinct} \\ 0 & \text{otherwise} \end{cases}$$

Otherwise indicates no significant difference or that we lack evidence to tell whether there is a significant difference.[4]

Sd[A_j] Like Sd[A_i], but for another attribute, A_j.

C[A_i, A_j]. A ***conflict*** exists when A_i and A_j support different alternatives. For example, if A_i supports p and A_j supports q, then there is a conflict.

$$C[A_i, A_j] = \begin{cases} 1 & \text{if } A_i[p] \tilde{>} A_i[q] \text{ and } A_j[p] \tilde{<} A_j[q] \text{ or} \\ & \text{if } A_i[p] \tilde{<} A_i[q] \text{ and } A_j[p] \tilde{>} A_j[q] \\ 0 & \text{otherwise} \end{cases}$$

I[A_i, A_j]. One attribute is often ***more important*** than another though we do not have to say ***why*** we believe this. It may be because the attribute itself is more important, disregarding

[2] Throughout this paper "attribute" is used loosely to refer to features of alternatives that are salient to the task of selecting the best alternative. This definition is vague enough to accommodate *outcomes*, *goals* or *characteristics*. Because each affects the decision differently, the model has been extended to account for these separate types of attributes and how they interact [Howe 86].

[3] A recent formulation of constructive decision-making allows us to represent probability distributions of values for alternatives.

[4] A recent version of the decision typology distinguishes these interpretations of "otherwise."

its value, or the values are so radically different on a particular attribute that the attribute provides much better evidence.

$$I[A_i, A_j] = \begin{cases} 0 & \text{if importance}(A_i) = \text{importance}(A_j) \\ ? & \text{if relative importance is unknown} \\ 1 & \text{if importance}(A_i) > \text{importance}(A_j) \\ & \text{or importance}(A_i) < \text{importance}(A_j) \end{cases}$$

$\tilde{>}[A_i, A_j]$. Assuming that $I[A_i, A_j] = 1$, we need to know which attribute is preferred.

$$\tilde{>}[A_i, A_j] = \begin{cases} 0 & \text{if importance}(A_i) < \text{importance}(A_j) \\ 1 & \text{if importance}(A_i) > \text{importance}(A_j) \\ * & \text{if } I[A_i, A_j] = 0 \end{cases}$$

These dimensions can be illustrated in the context of selecting between apples and oranges: p is apples, q is oranges, A_i is *quality* and A_j is *flavor*. If the quality of apples is "good" and the quality of oranges is "poor," then $\text{Sd}[quality] = 1$ because good and poor are distinct values. Similarly, if one prefers the flavor of oranges to that of apples then $\text{Sd}[flavor] = 1$. Since apples have better quality but oranges taste better, $\text{C}[quality, flavor] = 1$. Finally, if quality is preferred to taste $\text{I}[quality, flavor] = 1$ and $\tilde{>}[quality, flavor] = 1$.

By considering all possible combinations of the values of the five dimensions and pruning out isomorphic states, we obtain 24 basic states that constitute a search space. These are arranged in a table (Fig. 1) that we call the *decision typology*. The apples and oranges problem described earlier is represented by state 23 in this table. Because in this state a conflict exists between the alternatives on their attributes, a decision cannot be made with confidence. The goal of constructive decision making is to search the state space in Figure 1 by transforming states that represent decisions in which no choice can be made confidently into more facilitative states. In the next section, we describe the actions that produce the state transitions within the state space.

2.2 Decision Actions

In the 2-alternative, 2-attribute problem, the decision space is searched by transforming states via five actions: decision, transformation by attribute, transformation by importance, substitution, and combination. **Decision** (abbreviated *De*) means selecting an alternative based on available evidence. Transformations of one decision state into another are appropriate when a decision cannot be made given the available evidence. **Transformation by attribute** (*Ta*) seeks to transform the current decision state by gathering information about one of the attributes. For example, a decision in state 4 (Fig. 1) can be transformed to state 2 or state 5 b *Ta*, depending on whether thhe new information about the attribute implies a conflict between the attributes. **Transformation by importance** (*Ti*) gathers importance information; for example, state 4 can be transformed to state 10, or 17, or 18 by *Ti*. Note that search with these operators is nondeterministic: the state that results from *Ta* or *Ti* depends on whether

State #		0	1	2	3	4	5	6	7
$Sd[A_i]$		0	1	1	0	1	1	0	0
$Sd[A_j]$		0	0	1	0	0	1	0	1
$C[A_i, A_j]$		0	0	0	1	1	1	0	0
$I[A_i, A_j]$		?	?	?	?	?	?	0	0
$\tilde{>}[A_i, A_j]$		*	*	*	*	*	*	*	*
Actions	All Info	Co De	Su De	De		Su,Co De	Co	Co De	Su De
	Part Info	Ta 0,1,4 Ti 6, 12,20	Ta 1,5,8 Ti 7, 13,14	Ti 5,8 21	Ta 3,4,5 Ti 9, 16,22	Ta 2,4 Ti 10, 17,18	Ti 11, 19,23	Ta 6, 7,10	Ta 7, 8,11

State #		8	9	10	11	12	13	14	15
$Sd[A_i]$		1	0	1	1	0	1	0	1
$Sd[A_j]$		1	0	0	1	0	0	1	1
$C[A_i, A_j]$		0	1	1	1	0	0	0	0
$I[A_i, A_j]$		0	0	0	0	1	1	1	1
$\tilde{>}[A_i, A_j]$		*	*	*	*	0	0	0	0
Actions	All Info	Co De	Su Co	Su Co	Co	Co Su	Su De	Su Co	Co De
	Part Info		Ta 9,10,7	Ta 10,11,8		Ta 12,13, 14,17,18	Ta 13,15, 17,19	Ta 14,15, 18,19	

State #		16	17	18	19	20	21	22	23
$Sd[A_i]$		0	1	0	1	0	1	0	1
$Sd[A_j]$		0	0	1	1	0	1	0	1
$C[A_i, A_j]$		1	1	1	1	0	0	1	1
$I[A_i, A_j]$		1	1	1	1	1	1	1	1
$\tilde{>}[A_i, A_j]$		0	0	0	0	1	1	1	1
Actions	All Info	Co Su	Su Co	Su Co	Co	Co Su	Co De	Co Su	Co
	Part Info	Ta 16,17, 13,14,18	Ta 17,19, 15	Ta 18,19, 15		Ta 20,13, 14,17,18		Ta 22,13, 14,17,18	

Figure 1: Typology of Decisions

State numbers appear at the top of each column. For each state, the columns contain the values for the five dimensions followed by the actions that seem most reasonable. The actions are partitioned according to whether all the information has been acquired for the dimensions. If only partial information has been acquired, then transformations are better. The numbers following some of the actions refer to the state transitions that may occur if that action is taken.

significant differences truly exist between alternatives on an attribute, whether they conflict, whether one attribute is more important than another, and which is preferred. Each state in the decision typlogy can be transformed to another by asking the user one or more of these questions, but the destination state depends on the answer to the question.

If, after gathering all information about the attributes of alternatives, a choice cannot be made with confidence, then it is necessary to add a new attribute to the decision. Two actions , **substitution** (Su) and **combination** (Co), add attributes. When an existing attribute doesn't provide evidence (i.e., $Sd[A_i] = 0$), a new attribute may be substituted for it. Alternatively, a new attribute could be combined with existing ones by clustering the attributes according to the alternatives that they support. State 7 represents a decision for which substitution would be an appropriate action. If one preferred, say, apples on flavor but could not distinguish them from oranges based on quality, and neither attribute was more important, then one might substitute a third attribute such as price for quality. Assuming that price discriminates the alternatives, but supports the same alternative as flavor, this action would transform state 7 into state 8. If price supported the other alternative, then state 7 would be transformed into state 11. As before, this search is nondeterministic because one does not know in advance whether an attribute will support one alternative or another.

Combination is appropriate in many states, but especially when the decision is stuck (i.e., in state 11). If the alternatives can be discriminated by a third attribute, say price, then this attribute will cluster with another and break the impasse represented by state 11. For example, if flavor supports apples and quality supports oranges, and neither attribute is preferred, and if price supports apples, then price and flavor form a cluster that is necessarily preferred to quality.[5] That is, one can transform state 11 into state 19 or state 23 by combinination.

Clustering is the key to extending the basic two-alternative, two-attribute situations to two-alternative, N-attribute states, and finally to N-alternative, N-attribute problems. It permits complex decision situations to be constructed iteratively within the framework of our decision typology. We are currently developing an extended typology for multiple alternatives that is based on clustering alternatives on attributes.

3 Issues

We have implemented a decision support system based on the constructive decision making model [Howe 86]. The typology is used to select potential actions that are executed by querying the user for information about attributes and alternatives. Each piece of information causes the system to update its representation of the evolving decision and map it to a new state in the decision typology. If this state does not support a decision with confidence, then a new action is selected and the cycle repeats.

Our experience with the decision support system for constructive decision making has raised several issues. Two in particular are the focus of current work:

[5] If there was no preference between flavor and quality, and if price discriminates the alternatives, then there is necessarily a preference between the price-flavor cluster and quality.

- How should the system be controlled? In particular, if several actions are possible, which one is best? We are using the decision support system to explore control issues. Currently, the actions associated with a state are ranked and executed in that order. More sophisticated strategies are mandated by the extension of the typology for multiple alternatives.

- Should the typology use other information? Currently, it does not include information about the magnitude of differences between alternatives on attributes, nor does it consider probability distributions of these differences. It is relatively simple to extend the typology in these directions. However, the premise of constructive decision making is that information is acquired only as needed to support decisions, so we must carefully specify the states of decisions in which probability distributions and the like are required.

The constructive decision making model is not intended to produce optimal solutions to complex decision problems given complete information, but rather to explore methods for structuring decision problems, performing symbolic comparisons, and reasoning about uncertain decisions. It is the basis of a 2-alternative, multiple-attribute decision support system, and is currently being extended to support decisions among multiple alternatives.

References

Avron Barr and Edward A. Feigenbaum, editors. *The Handbook of Artificial Intelligence*, chapter 5. Volume 1, William Kaufmann, Inc., 1981.

Stanley F. Biggs and Theodore J. Mock. "An Investigation of Auditor Decision Processes in the Evaluation of Internal Controls and Audit Scope Decisions." *Journal of Accounting Research*, Spring 1983.

Paul R. Cohen, David Day, Jefferson Delisio, Michael Greenberg, Rick Kjeldsen, Daniel Suthers, and Paul Berman. "Management of Uncertainty in Medicine." *International Journal of Automated Reasoning*, Spring 1987.

R. A. Howard. "Decision Analysis: Applied Decision Theory." In D.B.Hertz and J.Melese, editors, *Proceedings of the fourth International Conference on Operational Research*, pages 55-71, Wiley, New York, 1966.

Adele Howe, *A Constructive Approach to Decision Making.* Master's thesis, Dept. of Computer and Information Science, University of Massachusetts, 1986.

Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty.* Addison-Wesley, Reading, MA, 1970.

Ola Svenson. "Process Descriptions of Decision Making". *Organizational Behavior and Human Performance*, 23:86–112, 1979.

C.A.S. Stael von Holstein. "A Tutorial on Decision Analysis." In *Third Research Conference on Subjective Probability, Utility, and Decision Making*, London, September 1971.