were in terms of the proportion of decrease or increase in contrast of T with reference to a standard. The standard was the contrast of T after adaptation to condition B. Subjects had practice sessions until they could give consistent ratings for the standard. For each adaptation field, six consecutive trials were presented before a new adaptation field was presented. Each adaptation stimulus was presented for two sets of six consecutive trials for each ISI; the order of presentation of adaptation stimuli was random. Stimuli were viewed with one eye only. Since the estimate of error variance is computed from the subject interaction, not from the within-cells replications, the presentation of six consecutive trials does not artificially lower the error term in the statistical analysis. On the other hand, such a procedure permits a thorough testing of adaptation effects, since any decay from a previous adaptation field would enter into only the first one or two presentations, and any buildup of an adaptation effect would be included.

When the geometric mean contrasts of T for each ISI and each adaptation condition (Fig. 2) are analyzed one finds that the apparent contrast for GC is reduced and that the reduction is due neither to the effect of prolonged viewing of a cube alone (C), nor to adaptation to mean luminance (B). This depression also differs in both amplitude and recovery course from that occurring for G. There is a general depression of apparent contrast at the early ISI's, which is no doubt caused by a general masking-by-flashes effect (*6*).

The differences between conditions G, GC, and B and conditions G, GC, and C are highly significant. A three-way analysis of variance showed all measurable main and interaction effects significant at $P < .001$. For this design, the error term is the interaction including subjects; thus, the within-cells variance (which would be computed from each of six sets of 12 estimations of magnitude) does not enter into the calculation. A Duncan's multiple range test showed all main effects significantly different from each other at $P < .001$ except those between B and C.

The depression in apparent contrast for GC cannot be due to artifacts caused by eye movement. Nor can it be due to a difference in mean spatial luminance, since if this were so, there should be an identical depression for C. Thus, this depression in apparent contrast may indicate that certain subsets of neurons are active not in the presence of the physical stimulus, but to some internal representation of the meaning of that stimulus. If the object blocked from view (the grating) is in some sense completed by neural activity, that is, if neurons symbolizing grating fire as if to a grating, it would be likely that these neurons would fatigue or adapt upon prolonged viewing, just as neurons active when a grating is presented fatigue or adapt. A corresponding depression in apparent contrast would result.

The effect may be more general, however. Whenever a grating is present, neurons may respond in portions of a scene where the grating is blocked, whether or not that portion contains an object which is clearly perceived to be in front of that grating. Hence, in order to test that activity in response to nongrating portions of the stimuli used in this experiment actually symbolizes "in front of," it must be shown that, with scenes in which gratings are simply interrupted, such as a picture of a grating with a hole in its middle, there is no adaptation effect.

Although it is known from electrophysiological work (*5*) that there is activity in the visual cortex throughout the process of recognition and learning, there has previously been little in the way of psychophysical measures of neural activity beyond that of the initial registration of stimuli. The results of this experiment indicate that activity beyond this—higher-order-activity—can be measured. The possibilities for investigation of higher-order stages in pattern recognition implied by these findings are very broad.

NAOMI WEISSTEIN
*Department of Psychology, Loyola University, Chicago, Illinois 60626*

**References and Notes**

1. C. Blakemore and F. W. Campbell, *J. Physiol.* **203**, 237 (1969).
2. N. Weisstein, *Psychol. Bull.* **72**, 157 (1969).
3. A. Guzman, *Proceedings of the 2nd Hawaii International Conference on System Sciences* (1969), vol. 2, p. 479; M. Minsky, in *Computers and Thought*, A. E. Feigenbaum and J. Feldman, Eds. (McGraw-Hill, New York, 1963), p. 406; U. Neisser, *Cognitive Psychology* (Appleton-Century-Crofts, New York, 1967).
4. K. H. Pribram, D. N. Spinelli, M. C. Kamback, *Science* **157**, 94 (1967).
5. R. M. Steinman, *J. Opt. Soc. Amer.* **55**, 1158 (1965).
6. D. Kahneman, *Psychol. Bull.* **70**, 404 (1968).
7. Supported by grant 8 RO1 EY 00143-02 from the National Eye Institute.

10 March 1970

## Assessment of Multiattribute Preferences

Abstract. *Operational assumptions are made concerning the preferences of a decision maker. Functional forms of multiattribute utility functions that satisfy these assumptions are stated. These forms provide operational methods for assessing preferences over multiattribute consequences and have a wide range of practical application.*

A decision problem is one with more than one available course of action. A consequence will eventually result from any particular course of action the decision maker chooses to follow, and he must choose a "best" course of action from the alternatives. For example, the manager of a blood bank must choose an inventory ordering policy that best satisfies the objectives of his particular blood bank. For a certain policy, a consequence might be described by "five percent of the blood requested by doctors cannot be supplied from stock."

In such decisions, there is the ever-present problem that consequences can rarely be described completely in terms of one attribute, such as the "percent of unsupplied demand" in our example. In the general case, one might describe these consequences in terms of several attributes. For our example, the consequences of a particular inventory policy might be summarized adequately by the number of units of outdated blood, the percent of unsupplied demand, and the age of transfused blood. When more than one attribute is necessary to describe the consequences, they are called multidimensional consequences.

To complicate the problem further, uncertainty is often associated with the consequences. Again referring to our example, for a particular ordering policy one could probably not state beforehand the exact percent of unsupplied demand, the exact number of units outdating, and so forth. Therefore, with each specific course of action, various consequences would have various probabilities of occurring.

The general decision problem is summarized as follows. The decision maker has a number of alternative courses of action, each of which will eventually result in a multidimensional consequence. However, at the time this choice must be made, there is uncertainty as to what consequence will actually result. The decision maker's problem is to select a course of action in view of this uncertainty and his preferences for the various possible consequences.

Decision analysis is a systematic procedure for analyzing such problems. With this approach, it is assumed that the decision maker's preferred course of action will depend on (i) the probabilities that this course of action will result in the various possible consequences, and (ii) his preferences for those consequences.

To conduct a systematic analysis, the decision maker must quantify his judgment by using probabilities that represent his present state of knowledge and his preferences. Thus, one of the most important problems in decision analysis concerns quantification of the decision maker's preferences. For this purpose, a utility function (*1*) over all possible consequences is required.

At present, the only general method for evaluating consequences over more than one attribute is the additive utility function (*2*), which requires restrictive assumptions on the decision maker's preferences. Consequently, the basis of many of the methodological difficulties in applying decision analysis to complex problems is that appropriate techniques do not exist for systematically assessing multidimensional utility functions valid for decision making under uncertainty. This is the problem of concern in this research.

The basic approach used in this study is to make assumptions about the preferences of the decision maker and then to investigate the restrictions that these assumptions place on his utility function. Assumptions are chosen that are operationally significant and are relevant to many decision problems. Before stating the results, the concepts of utility independence and conditional utility functions must be introduced.

Let $(x_1,x_2,\cdots,x_n) \epsilon X$ represent a consequence and $u(x_1,x_2,\cdots,x_n)$ be a utility function over $X$. Vector attributes $Y$ and $Z$ are defined such that $X = Y \times Z$, where $y = (x_1,x_2,\cdots,x_m)$ and $z = (x_{m+1},x_{m+2},\cdots,x_n)$ represent specific amounts of $Y$ and $Z$, respectively. A consequence may now be written as $(y,z)$. Setting $z$ equal to a specific value $z_0$, the conditional utility function for $Y$, given $z = z_0$, is $u(y,z_0)$.

Given $u(y,z)$, $Y$ is said to be utility independent of $Z$ if the decision maker's relative preferences for different amounts of $Y$, when $Z$ is held fixed at $z_0$, are the same regardless of the amount of $z_0$. When this is the case, since utility functions are unique up to positive linear transformations, the conditional utility function for $Y$, given any value of $Z$, is a positive linear transformation of the conditional utility function for $Y$ given any other value of $Z$. Mathematically stated, if $Y$ is utility independent of $Z$, then for any $z_0$ and all $z$,

$$u(y,z) = c_1(z) + c_2(z)\, u(y,z_0)$$

If $Y$ and $Z$ are utility independent of each other, they are said to be mutually utility independent. Similarly, if each of the $X_i$'s is utility independent of all the others, they are mutually utility independent.

The results of this research (stated below) provide operational methods for simplifying the assessment of a multidimensional utility function appropriate for making decisions under uncertainty, provided the requisite assumptions hold. In the $n$-dimensional case, if the conditional utility function for $X_i$ is denoted by $u_i(x_i,x_{\bar{i}})$, where $x_{\bar{i}}$ represents a fixed amount of all the other attributes, and if $x_i{}^*$ and ${}^*x_i$ are arbitrarily chosen subject to

$$u_i(x_i{}^*,x_{\bar{i}}) > u_i({}^*x_i,x_{\bar{i}})$$

then one can prove theorem 1.

*Theorem 1.* Given $X = X_1 \times X_2 \times \cdots \times X_n$ and the $X_i$ are mutually utility independent, $u(x_1,x_2,\cdots,x_n)$ is completely determined by

(a) $u_i(x_i, x_{\bar{i}})$ for arbitrary $x_{\bar{i}}$, for each $X_i$

and

(b) $u(x_1{}^a, x_2{}^a,\ldots, x_n{}^a)$ for all $x_i{}^a = x_i{}^*$ or ${}^*x_i$

For the case of $X$ partitioned into $Y$ and $Z$, there is the corollary.

*Corollary.* If $Y$ and $Z$ are mutually utility independent, $u(y,z)$ can be evaluated from

$$u(y,z) = u(y,z_0) + u(y_0,z) + ku(y,z_0)\, u(y_0,z)$$

where $y_0$ and $z_0$ are arbitrarily chosen and $k$ is an empirically evaluated constant.

Another important result for the case where $X = Y \times Z$ is theorem 2.

*Theorem 2.* If $Z$ is utility independent of $Y$, subject to consistent scaling, $u(y,z)$ can be evaluated from

$$u(y,z) = u(y_0,z)\, u(y,z_1) + [1 - u(y_0,z)]\, u(y,z_0)$$

where $y_0$, $z_0$, and $z_1$ are arbitrarily chosen.

The usefulness of these results is that they reduce the problem of assessing an $n$-attribute utility function to one of assessing $2^n - 2$ empirically evaluated constants and $n$ single-attribute utility functions. Adequate techniques exist to evaluate these constants and to assess one-attribute utility functions (*3*). Also, the requisite assumptions require that the decision maker specify his relative preferences for consequences with only one attribute varying at a time. I have found that decision makers are able to make this choice, whereas they find it quite difficult to directly specify relative preferences for consequences with more than one attribute varying. Therefore, one can readily determine whether the utility independence assumptions hold for a particular problem. The value of these results is that they provide operational methods for assessing multiattribute preferences.

The results summarized here have been applied to three problems in (*4*). The first concerns preferences over service levels for two types of customers in a telephone system; the second is the preferences for shortage and outdating of blood in a hospital blood bank; and the third involves preferences over the cost and accuracy for different surveys of an inland waterway. Yntema and Klem (*5*) applied these ideas to the assessment of the safety of landing aircraft in different weather situations with three varying attributes. In each of those problems, the appropriateness of the assumptions was verified and the decision maker's preferences were assessed.

Currently, these and similar ideas are being used to quantify the preferences of a metropolitan fire department for response to fires. The different attributes represent response times of various fire department apparatus. The preferences will be used with a simulation model of the fire department's operations in examining alternative operating policies of the department.

RALPH L. KEENEY
*Civil Engineering Department and Operations Research Center, Massachusetts Institute of Technology, Cambridge 02139*

### References and Notes

1. A utility function assigns a real value to every consequence in a manner such that expected utility can be used as a guide to making decisions under uncertainty. An excellent reference on utility theory is P. C. Fishburn, *Manage. Sci.* **14**, 335 (1968).
2. A utility function, $u(x_1,x_2,\cdots,x_n)$, is additive if

$$u(x_1,x_2,\cdots,x_n) = \sum_{i=1}^{n} u_i(x_i)$$

where $u_i$ is a utility function over the $i$th attribute.
3. See, for example, R. Schlaifer, *Analysis of Decisions under Uncertainty* (McGraw-Hill, New York, 1967).
4. R. L. Keeney, *Operations Research Center Technical Report 43* (Massachusetts Institute of Technology, Cambridge, 1969).
5. D. B. Yntema and L. Klem, *IEEE Trans. Hum. Factors Electron.* **6**, 3 (1965).
6. I thank Professors H. Raiffa of Harvard and A. W. Drake of M.I.T. for their suggestions during this research. This work was partly supported by the U.S. Army Research Office (Durham) under contract DA-31-124-ARO-D-209, and by the Office of Naval Research under contract Nonr-3963 (06), NR 276-004.

13 March 1970