Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

Fuzzy Signal Detection Theory: Basic Postulates and Formulas for Analyzing Human and Machine Performance

Raja Parasuraman and **Anthony J. Masalonis,** Catholic University of America, Washington D.C., and **Peter A. Hancock,** University of Minnesota, Minneapolis, Minnesota

Signal detection theory (SDT) assumes a division of objective truths or "states of the world" into the nonoverlapping categories of signal and noise. The definition of a signal in many real settings, however, varies with context and over time. In the terminology of fuzzy logic, a real-world signal has a value that falls in a range between unequivocal presence and unequivocal absence. The definition of a response can also be nonbinary. Accordingly the methods of fuzzy logic can be combined with SDT, yielding fuzzy SDT. We describe the basic postulates of fuzzy SDT and provide formulas for fuzzy analysis of detection performance, based on four steps: (a) selection of mapping functions for signal and response; (b) use of mixed-implication functions to assign degrees of membership in hits, false alarms, misses, and correct rejections; (c) computation of fuzzy hit, false alarm, miss, and correct rejection rates; and (d) computation of fuzzy sensitivity and bias measures. Fuzzy SDT can considerably extend the range and utility of SDT by handling the contextual and temporal variability of most real-world signals. Actual or potential applications of fuzzy SDT include evaluation of the performance of human, machine, and human-machine detectors in real systems.

If a man will begin with certainties, he shall end in doubts; but if he will be content to begin with doubts he shall end in certainties.

– Francis Bacon, *The Advancement of Learning* (1605, bk. 1, v. 8)

INTRODUCTION

$\mathbf{A} \cap \bar{\mathbf{A}} = 0$. This seemingly irrefutable mathematical expression asserts that a statement and its opposite can never coexist. Something either *is* or *is not.* A tumor is either cancerous or benign; a new car is either reliable or faulty; a politician is either honest or dishonest, and so on. Confidence in the truth of this mathematical expression and of these representative statements pervades both academic and everyday thinking.

A little reflection reveals that although the logical expression is often true, it need not always be true. After all, tumors, cars, and politicians come in all shades, not just black and white. To cope with this possibility, Zadeh (1965) developed *fuzzy logic,* sometimes simply called *fuzzy* (Kosko, 1993, 1997). Fuzzy logic allows the possibility that the intersection of **A** and $\bar{\mathbf{A}}$ is nonzero. Is it or isn't it? The truth lies somewhere in between. To the extent that an event *is* somewhere in between, forcing its categorization into nonoverlapping sets of black and white can result in the loss of useful information and less sensitive analysis. Rather, if we follow Bacon's admonition to "begin with doubts" and express those doubts mathematically in fuzzy terms, then the analyses that follow may very well "end in certainties."

If we are doubtful whether a given event is a member of a particular category or not, how can we reach meaningful decisions and take appropriate actions based on our knowledge of the properties of that category? To return to our examples, the tumor needs to be operated

Address correspondence to Raja Parasuraman, Cognitive Science Laboratory, Catholic University of America, Washington, DC 20064; parasuraman@cua.edu. **HUMAN FACTORS**, Vol. 42, No. 4, Winter 2000, pp. 636–659.

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

on, the faulty car returned to the dealer, the politician voted out of office. Is a fuzzy characterization of events a recipe for indecisiveness and inaction? Not necessarily.

As we demonstrate in this paper, fuzzy logic can be combined with a well-known methodology for analyzing decision making: signal detection theory (SDT). The result, which we term *fuzzy SDT*, allows for a broader and potentially more powerful analysis of decision-making performance than do conventional methods. Moreover, as we shall show, use of fuzzy SDT can avoid the possibility of erroneous conclusions that may arise from the application of standard SDT to situations in which signal definition is fuzzy.

SDT was initially developed to quantify the performance of electronic receivers for detecting noisy radio signals (Peterson, Birdsall, & Fox, 1954). It was later extended to describe human detection of threshold-level signals (Tanner & Swets, 1954). Green and Swets (1966) described the modified theory in a landmark book that led to widespread application of SDT (in psychology and related disciplines) to a variety of perceptual and cognitive tasks involving decision making (Swets & Pickett, 1982).

SDT has been shown to provide independent measures of the bias and the accuracy of decision outcomes, thereby granting it many advantages over competing theories and computational methods, such as high-threshold theory (MacMillan & Creelman, 1991). Another advantage is that SDT can be used to analyze human, machine, or joint human-machine performance, thus providing a common metric to describe diverse aspects of detection performance in many application domains (Parasuraman, 1985; Parasuraman & Wisdom, 1985; Sheridan & Ferrell, 1974; Sorkin & Woods, 1985; Swets, 1996).

This well-established theoretical approach is based on a division of objective truths, or states of the world, into one of two nonoverlapping categories: signal or noise. SDT also typically (but not always) assumes binary responses made by the human or machine observer; for example, "yes, a signal is present," or "no, a signal is not present." Traditional SDT requires the mapping of environmental events or sources of evidence into two categorical states of the world. As our opening examples of various events and states indicate, such mapping is often fuzzy rather than exact. The question "What is a signal?" does not usually arise in laboratory studies. The exception is when the experimenter is interested in examining the effects of *uncertainty* regarding some aspect of the signal, such as its duration, starting time, frequency, and so on. We discuss the general problem of uncertain signal detection (Tanner & Birdsall, 1958; Green & Swets, 1966) in relation to fuzzy SDT in a later section of this article.

In the laboratory, the signal is whatever the experimenter defines it to be. In a perceptual experiment on visual discrimination, for example, the signal may be defined as a line oriented at 90°, whereas lines with orientations of 0–80° are defined as nonsignals or noise. A researcher studying recognition memory may define a signal as a face that has been shown to the participant during a prior study period and other previously unseen faces as noise.

In most such laboratory studies of perception, memory, and cognition, the categorization of a physical event as signal or noise is fixed. In contrast, in most real settings the definition of a signal is often context-dependent and varies with several factors. Real-world signals are fuzzy. For example, a sonar operator trying to detect the electronic signature of an approaching submarine on a visual display terminal has to look for a line with a luminance, contrast, and spatial frequency that will depend on the submarine speed, ocean currents, presence of other nearby objects, and so on. The signal will also vary over time as the submarine approaches.

Even when the formal or legal definition of a signal is clearly specified in terms of some measurable physical event, people may treat the event variably in different contexts. For example, the legal definition of a conflict in the flight paths of two aircraft being monitored by air traffic control (ATC) is fixed. According to Federal Aviation Administration regulations, a "signal" in ATC occurs when two aircraft come within 5 nautical miles (nmi) horizontally and 1000 ft vertically of each other.

The acceptable minimum aircraft separation values (vertical and lateral) actually depend on several variables, including the altitude of the aircraft and the environment (cruise, approach

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

for landing, etc.). For the sake of simplicity, however, in this paper we use the 5-nmi/1000-ft minimum standard that prevails in low-altitude (below 29,000 ft) cruise flight.

However, the separation distances that the air traffic controller will consider a signal requiring action will often exceed these minimum values. The controller's definition of signal will vary depending upon the complexity of the traffic, the nature of the ATC sector being controlled, and other factors. In one context, say, with many climbing and descending aircraft, two planes separated by 8 nmi may to some extent represent a signal requiring urgent action because of the potentially high likelihood that the planes will come closer while the controller is distracted by other aircraft. In another situation, say, level cruise flight, two aircraft that approach to within 6 nmi may not particularly perturb the controller.

The definition of a signal in many real-world situations is therefore variable or fuzzy. We introduce the concepts of contextual and temporal variability as key characteristics of signals in real settings and propose that fuzzy methods are particularly well suited to deal with these sources of variability. These considerations suggest that fuzzy logic could be used to extend traditional SDT analysis to situations in which membership of events to signal and noise sets is not strictly dichotomous.

There has been some limited work incorporating both fuzzy logic and SDT. Fuzzy methods have been used to create decision-making algorithms whose results have then been analyzed with traditional SDT or other decision-theoretical analysis; such combinations have been used in the domain of personnel psychology (Alliger, Feinzig, & Janak, 1993; Craiger & Coovert, 1994). In contrast, the present work considers fuzzy logic and SDT simultaneously as opposed to successively, and develops general methods that can be used across a wide variety of applied and basic domains. Furthermore, our methods allow for an analysis of the influence of the fuzziness inherent in the real world at all stages, from signal to response. We develop the basic postulates of fuzzy SDT and derive formulas for the fuzzy analysis of detection performance. We also discuss the implications of fuzzy SDT analysis for understanding human and machine detection performance. Before describing the formulas for fuzzy SDT analysis, we provide brief overviews of fuzzy logic and conventional SDT in the next two sections. Readers familiar with these topics may wish to skip either or both of these sections.

FUZZY LOGIC: A REVIEW

Fuzzy logic represents an alternative method to traditional set theory for assignment of membership of events to sets (Zadeh, 1965). Fuzzy methods have been used in a variety of basic and applied areas (see Kosko, 1993, for numerous examples from many scientific domains). Among the areas that are particularly relevant to the present paper are applications of fuzzy logic in the following domains: (a) cognitive psychology (Brainerd & Reyna, 1993; Campbell & Massaro, 1997; Ellison & Massaro, 1997; Heiser & Groenen, 1997; Massaro, 1988, 1998), (b) human factors/ergonomics and human-computer interaction (e.g., Genaidy et al., 1998; Karwowski & Mital, 1986; Kreifeldt & Rao, 1986; Lehto & Sorock, 1996; Moray, King, Turksen, & Waterton, 1987; Moray, Kruschelnicky, Eisen, Money, & Turksen, 1988), and (c) biomedical engineering and neuroscience (Baumgartner, Winddischberger, & Moser, 1998; Bellazzi, Silviero, Stefanelli, & De Nicolao, 1995; Boston, 1997; Lowe, Harrison, & Jones, 1999; Phelps & Hutson, 1995). Despite the great number of these applications, we must reiterate one point: The application of fuzzy logic in the previously cited works has been either to detect a discrete (nonfuzzy) signal or to analyze a system or result that was derived using nonfuzzy methods. In contrast, in the present paper we present a general model for analysis of fuzzy signals and fuzzy responses.

Fuzzy is best thought of as an extension of traditional set theory, in which elements either do or do not belong to a given set. Because a clear boundary exists between set members and nonmembers, traditional sets are often referred to as *crisp* sets. Suppose we wish to define a set to describe the range of temperatures (t) that a normal human being would consider comfortable. An example of a crisp set is

$$C(t) = [55, 85], \tag{1}$$

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

which would mean that the numbers 55 and 85 (temperatures in degrees Fahrenheit), and all real numbers between these two, are members of the set C, and all numbers below 55 or above 85 are nonmembers. If we plot $C(t)$ versus t, the result would appear as shown in the solid line in Figure 1. However, it would seem more appropriate to distinguish between levels of comfort rather than to assign every temperature to either the *comfortable* or *uncomfortable* sets. We could instead develop a function that permitted a temperature's membership in the set *comfortable* to be somewhere between yes and no, or between 0 and 1. The mapping function could be derived in a number of ways. For example, one could use a scale from 0 (completely uncomfortable) to 1 (completely comfortable) and have a sample of people rate several temperature values on this scale. The average rating for each temperature and the resulting function might be defined as follows:

$$C(t) = \begin{cases} 0 & t < 50 \quad \text{(2a)} \\ (t-50)/20\,, & 50 \le t < 70\,. \quad \text{(2b)} \\ (90-t)/20 & 70 \le t < 90 \quad \text{(2c)} \\ 0 & t \ge 90 \quad \text{(2d)} \end{cases}$$

This function, which is shown in the dotted lines in Figure 1, provides a more realistic assessment of comfort versus temperature and better captures the variability in comfort level. The act of assigning nonbinary membership degrees to a previously binary definition can be called *fuzzification.*

SIGNAL DETECTION THEORY: A REVIEW

SDT assumes two possible states of the world: *signal* (s), in which the event of interest is present, and *noise* (n), in which it is absent (Green & Swets, 1966). At any given time one of these states of the world occurs. The detection system (human or machine, or some combination) makes a *yes* (Y) or a *no* (N) judgment,

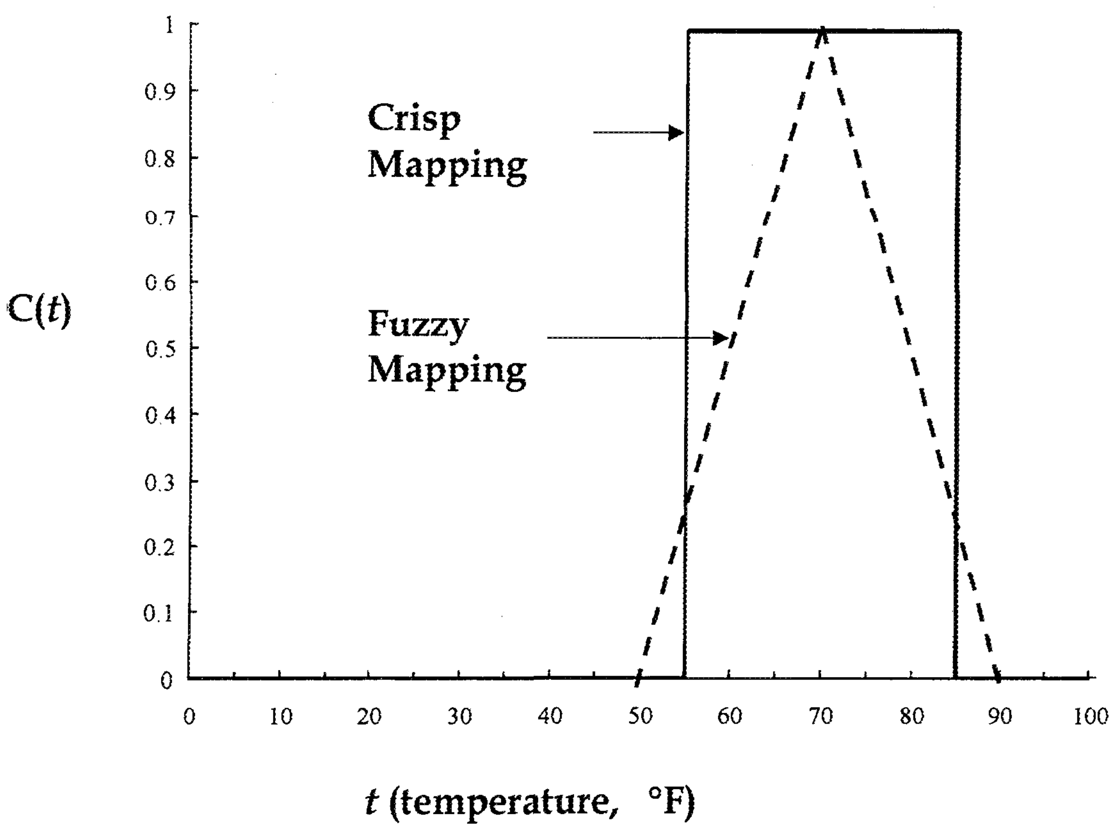

Figure 1. Crisp and fuzzy plots of comfort, *C*, as a function of temperature, *t*.

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

indicating whether or not it is believed that the signal is present or absent. Across many such occurrences or trials, the two possible states of the world and the two possible decisions result in four possible outcomes, each with an associated probability (P):

Hit; signal present, *yes* judgment, $P(Y|s)$ = Hit rate (HR) (3a)
Miss; signal present, *no* judgment, $P(N|s)$ = Miss rate (MR) (3b)
False Alarm (FA); signal absent, *yes* judgment, $P(Y|n)$ = FA rate (FAR) (3c)
Correct Rejection (CR); signal absent, *no* judgment, $P(N|n)$ = CR rate (CRR). (3d)

Only two of the four probabilities are needed for complete characterization of the performance outcomes, because $HR + MR = 1$, and $CRR + FAR = 1$. The convention is to use the probability of a Hit, or Hit Rate (HR) and the probability of a False Alarm, or FA Rate (FAR), to describe the decision outcomes. The hit and false alarm probabilities can then be used to compute various measures of the performance of the detection system. In general, it is necessary to distinguish the *sensitivity* or bias-free accuracy of the detection system from the *criterion* or decision threshold associated with the choice of judgments or responses. In SDT, sensitivity is indexed by the parameter d' and the criterion by the parameter β (Green & Swets, 1966; MacMillan & Creelman, 1991).

BASIC ELEMENTS OF FUZZY SDT

Mapping Functions for Signal and Response

The basis for the notion of fuzzy SDT is that an event or trial can belong to the set *signal* with some degree between 0 and 1. In addition, the response can belong to the set *response* with a degree between 0 and 1. Throughout this paper the parameter s will refer to the degree to which an event is a signal. The parameter r will refer to the degree to which a *yes* (signal present) response was made. In fuzzy SDT, either s or r or both must be continuous variables in the range [0, 1]. One of the two may be binary; that is, $\in \{0, 1\}$. However, s and r cannot both be binary, for this would then reduce fuzzy SDT to crisp SDT. This paper primarily deals with cases in which s and r are both continuous, or s is continuous but r is binary.

In assigning degrees of membership to the signal (s) and response (r) sets, it is necessary to evaluate all possible states of the world (SWs) and each possible response value (RV), and to determine what values will be assigned to each. A mapping function is required to derive a signal value, s, based on some variable or set of variables that describe the SW. Many such mapping functions are possible for the signal.

As noted by Moray et al. (1987), in many applications of fuzzy logic the mapping functions that are chosen are plausible but nevertheless appear somewhat arbitrary (but see Tsoukalas and Uhrig, 1997). This is particularly true if a mapping function has to be developed for some *perceived* characteristic, such as the emotional quality of a face (Ellison & Massaro, 1997) or the heaviness of a lifted object (Genaidy et al., 1998). This problem is somewhat less severe in fuzzy SDT, because the mapping function for the signal can generally be based on objective physical variables corresponding to the state of the world. Given that these variables can be measured, the mapping function can be specified.

Of course the type and complexity of the mapping function will vary with the application. For instance in the ATC example described earlier, a function could map the distance separating two aircraft, call it a, onto s. The mapping function could also be considerably more complex than this, by taking into account not only a, but other variables as well.

The mapping function for the response could be based simply on a confidence rating of signal presence; for example, "80% confident a signal occurred." Such mapping is an extension of traditional SDT methodology (Green & Swets, 1966). Often in standard SDT, in order to derive multiple receiver operating characteristic (ROC) points, judgments of confidence are permitted to be incorporated into the response (e.g., yes-sure, yes-not sure, no-not sure, no-sure). Despite the nonbinary nature of the judgments, the traditional analytic procedure for such rating data has still been categorical rather than fuzzy. Typically each

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

confidence rating is mapped cumulatively into dichotomous yes-no categories (i.e., first an analysis is conducted with *yes-sure* as a *yes* and all other responses as a *no*, next the *yes-sure* and *yes-not sure* are analyzed as *yes* responses with the others being *no*, and so on). Fuzzy SDT differs fundamentally from crisp SDT in that membership assignments to the *yes* and *no* categories remain fuzzy rather than being forced into crisp categories. Finally, the mapping function for r can also be based on reported signal severity, criticality, or intensity.

In SDT, the value of a random variable, the *evidence* variable, represents the strength of the signal as perceived by the operator (Green & Swets, 1966). As the evidence variable is generally continuous, it is amenable to fuzzification. For some applications, in fact, the r values can be assigned based on the evidence variable. This strategy is appropriate if signal detection performance is being measured with regard to how decisions are made based on what is perceived. However, if one is interested in decision-making performance and patterns with respect to some "true" or objective state of the world, then the variable mapped into r must be some measure of the actual objective state of the system about which decisions are being made.

Mapping functions: Discrete states of the world and responses. The following formulas can be used whenever there are a finite number of discrete states of the world (see Tsoukalas & Uhrig, 1997):

$$s = s_S \text{ for } SW = 1 \text{ to } n_S \quad (4a)$$
$$r = r_R \text{ for } RV = 1 \text{ to } n_R \quad (4b)$$

in which SW = each possible value of the state of the world, RV = each possible response value that could be made, n_S = number of possible discrete SW values, and n_R = number of possible discrete RV values. One case in which this mapping approach would be used is when discrete ratings of confidence level of response are taken (yes-sure, yes-unsure, and so on).

Mapping functions: Continuous states of the world and responses. In other situations the original variables representing SW and RV may be continuous. In these cases, SW or RV can take on an infinite number of values, and a function should be constructed mapping the SW or RV values to s or r values, respectively, in the range between 0 and 1:

$$s\,(SW) = \mathrm{f}(SW) \quad (5a)$$
$$r\,(RV) = \mathrm{g}(RV) \quad (5b)$$

in which SW = original value of the variable representing state of the world, on an interval scale, RV = original value of the variable representing response value, on an interval scale, and f(.), g(.) = functions defined according to the needs of the analysis being conducted and whose range must be [0, 1] or included within [0, 1].

Some features of these mapping functions should be noted. First, the function itself need not be continuous, depending on contingencies in the domain being analyzed. For example if some legal cutoff point exists in the SW, a break in the function's continuity can occur at that point (e.g., a jump from r = .5 to r = .9). Also, in fuzzy SDT it is not necessary to choose between Equations 4a–4b and Equations 5a–5b. For example it is permissible to have s be based on a continuous interval value (use Equation 5a) and r on a categorical value (use Equation 4b), or vice versa.

Example mapping function. Consider conflict detection in air traffic control, in which the signal (presence of an aircraft conflict) is coded according to the separation distance, a, between two aircraft potentially headed for a conflict (horizontal separation less than 5 nmi). The function mapping the SW variable, a, to s can be determined according to the following rationale. As the separation distance, a, decreases, the event becomes more signal-like and conversely, becomes less signal-like as a increases. As mentioned previously, the legal FAA definition of a conflict in ATC is 5 nmi horizontal separation. Clearly, though, worse things can happen besides a violation of the legal minimum. For this reason a collision (a = 0) is defined as s = 1 and a marginal violation of the 5-nmi criterion (a = 5) as s = .9. A monotonic decreasing function can be constructed that allows for an increasingly sharp drop-off of s as a increases beyond the 5-nmi cutoff and yields relatively similar (low) values of s at all high values of a.

Table 1 gives some values of a and s, and Figure 2 plots the mapping function. The crisp mapping function that would be used in standard SDT is also shown. For $a > 5$ nmi, the crisp function assigns no value to signal. This could result in some loss of information and potential insensitivity in analyzing detection performance in comparison with fuzzy SDT because it is well known that air traffic controllers typically attend to, and sometimes even act on, pairs of aircraft that will be greater than 5 nmi apart.

If there are many or infinite SW values, or a theoretical basis exists for using a specific mapping function, or both, a continuous mapping function (Equation 5a) should be used. For example, a could be mapped to s using the sigmoid function, $s = 1/[1 + (a/\mathrm{k})^n]$, in which k is a constant and the exponent n can be chosen depending upon the desired sharpness of the mapping function. The function in Figure 2 can be approximated with a sigmoid with k = 10 and n = 5. Other sigmoid functions could also be used, such as the "squashing" functions that are commonly used in connectionist (neural network) models (Rumelhart, Hinton, & Williams, 1986).

Exploration of Methods for Calculating Fuzzy Set Membership

Let us retrace the steps of the analysis so far. First we allowed the definition of both signal, s, and response, r, to be continuous rather than binary. Next, we discussed how mapping functions can be selected for deriving values of s and r. An obvious next step is to determine to which of the four traditional SDT categories (Hit, Miss, FA, and CR) a given s and r pair should be assigned. To which category should the $\{s, r\}$ pair be assigned and how? The solution is to allow a given event pair to belong with some degree to *more than one* of the four categories. In traditional SDT, each event falls completely into one of the four outcome categories, based on two dichotomies (SW – signal vs. noise; RV – yes vs. no). In fuzzy SDT, on the other hand, each event can fall into more than one of the four categories used in traditional SDT.

In evaluating methods for determining an event's degree of membership in each of the four outcome categories, we applied two criteria to develop an appropriate set of functions. First, the functions had to reduce to those of crisp SDT when s and r were made binary, because fuzzy SDT is an extension of the crisp case (or alternatively, crisp SDT is a special case of fuzzy SDT). A second criterion was face validity – if a set of functions gave results that were illogical for any or all inputs, they were rejected.

An important component of our judgment of face validity was the expectation that performance will be considered better the closer that r is to s. For example, if $s = .7$, the optimum degree of response r should also be .7. In the

TABLE 1: Example of a Mapping Function in an Air Traffic Control Application

Separation a (nmi)	Signal (s)	Separation a (nmi)	Signal (s)	Separation a (nmi)	Signal (s)
0.0	1.00	4.1	0.92	9.0	0.65
1.0	0.98	4.6	0.91	10.0	0.50
1.3	0.97	4.9	0.90	11.2	0.33
1.9	0.96	5.0	0.90	12.0	0.25
2.0	0.96	5.8	0.88	12.3	0.23
2.2	0.96	6.0	0.87	13.4	0.15
2.4	0.95	6.9	0.83	14.5	0.12
2.9	0.94	7.0	0.83	15.6	0.09
3.0	0.94	7.1	0.82	16.1	0.08
3.1	0.94	7.3	0.81	16.2	0.08
3.3	0.93	7.4	0.80	17.8	0.04
3.7	0.93	7.7	0.77	19.0	0.02
3.9	0.92	7.9	0.76	19.3	0.01
4.0	0.92	8.8	0.67	20.2	0.00

Note. An Aircraft-to-Aircraft "Conflict" or Signal (s) is Defined in Terms of the Separation Distance (a) Between Aircraft.

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

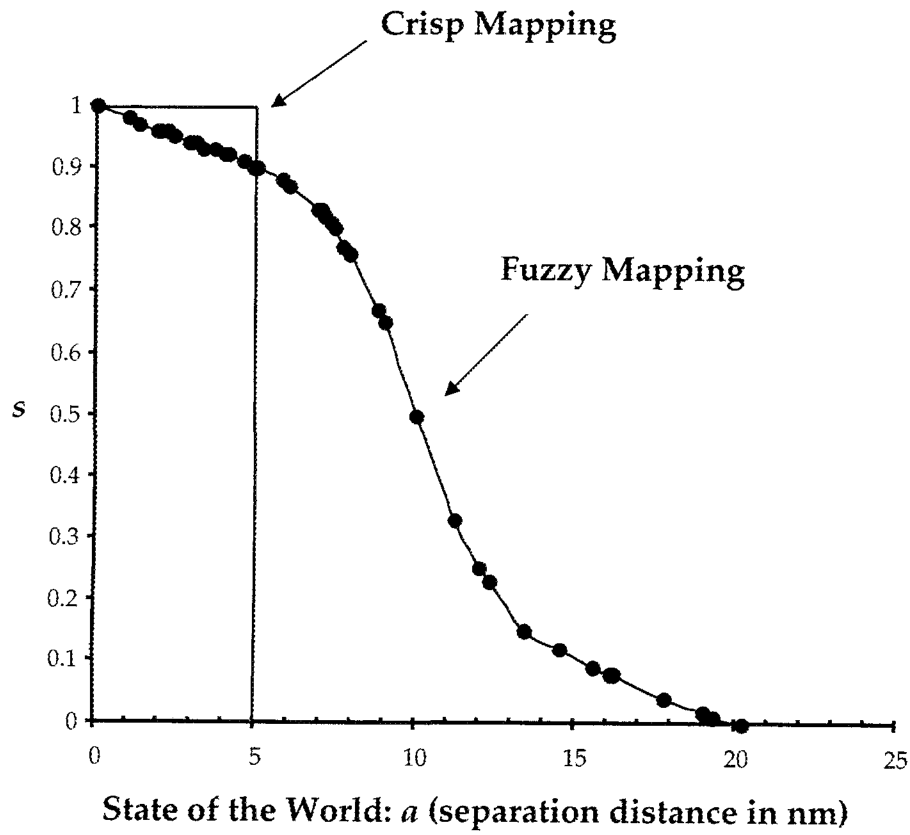

Figure 2. Sigmoid mapping function relating signal value, *s*, to each possible value of aircraft separation distance, *a*, as might occur in an air traffic control case. The rectangular function represents the standard or crisp SDT mapping of *s* to *a*.

same vein, if $r > s$, the ideal set of functions will assign at least some degree of FA membership to the event because the response degree is higher than ideal and the Miss degree for such an event should be 0. The opposite should hold true when $r < s$. The Miss degree membership should be nonzero and that for FA should be 0.

The considerations just discussed are based on the assumption that both *s* and *r* are defined on scales that are mappable to each other. In other words, whatever the application domain, the method of assigning memberships to *s* and *r* should be done such that an *s* of 0 calls for absolutely no response ($r = 0$), and an *s* of .3 calls for whatever degree of response has been defined as $r = .3$, and so on. There is arguably some danger in the assumption that such a mapping is always possible, because the assignment methods for defining *s* and *r* in a fuzzy system can be subjective. However, this does not mean that there will be no defensible basis for membership assignments. This point is well recognized in the literature on fuzzy logic: "Membership functions may represent an individual's (subjective) notion of a vague class...Membership functions may also be determined on the basis of statistical data or through the aid of neural networks...membership functions are primarily subjective in nature; this does not mean that they are assigned arbitrarily, but rather on the basis of application-specific criteria" (Tsoukalas & Uhrig, 1997, p. 15).

It should be remembered that assignments in crisp analyses are frequently arbitrary as well. Take the aforementioned heating example. If a heater operates on a simple crisp thermostat rule in which the heat will be *on* below a certain room temperature and *off* above that temperature, the designer has defined (perhaps arbitrarily, or perhaps because of engineering constraints not related to the end goal of human comfort) what heating strength (e.g. furnace temperature, fuel consumption) corresponds to *on*. This setting cannot always be easily changed by the user of a home or office heating system. Therefore the arbitrary nature of signal/response assignments is not peculiar to fuzzy SDT. As with crisp SDT, some fuzzy SDT *s* and *r* assignments will be more arbitrary than others, and the *s* assignments will be mappable to the *r* assignments more so in some cases than in others. For example in aircraft conflict detection, both signal and response degree might be assigned based on the same function translating distance between aircraft into the range [0, 1]. In other words, the true state of the world might be defined in terms of the closest approach distance between two aircraft, and the response might also be a judgment, by a human or an automated tool, of that distance. This mapping strategy might be considered better than basing *s* on distance and *r* on probability. However, sometimes the nature of the problem will be such that we must use whatever data best reflects the question we are trying to answer about human or machine performance, even at the expense of the best possible mapping. Continuing with the ATC example, the most important performance measure is likely to be the human controller's or machine's degree of confidence that a legal separation violation will occur, rather than the exact distance within which it is believed the planes will approach.

One cautionary note should be considered. If in a fuzzy SDT application there is any suspicion that the selected *s*-to-*r* mapping might be flawed (e.g., whatever has been defined as $r = .9$ is not likely to be the desired response to whatever is defined as $s = .9$), care should be taken in comparing derived outcome parameters of sensitivity, hit rate, and so on, to other systems. The mapping could still be used, however, to compare the performance of different detectors within the same system or to evaluate the performance of the same detector in that system under different conditions.

BASIC FORMULAS OF FUZZY SDT

We turn now to a recommended set of methods for carrying out a fuzzy SDT analysis. Once *s* and *r* are mapped onto [0, 1] using appropriate mapping functions, the next step is to derive event membership in the four outcome categories, Hit, Miss, FA, and CR. In this section we present the methods that met the twin criteria of matching crisp SDT results and face validity. We examined several other methods, all of which were found to be invalid for one reason or another. (A short report describing these rejected methods is available from the authors on request.)

Mixed-Implication Functions for Fuzzy SDT

In standard SDT, the function mapping values of *s* and *r* to the four outcome categories can be specified simply in a truth table format. The truth table, based on logical implication functions (e.g., "If SW = signal and response = yes, then Hit"), is shown in Table 2. Each event is mapped exclusively to only one of the four outcome categories, with a value of 1, whereas the remaining categories have entries of 0. In contrast as mentioned previously, in fuzzy SDT, events will claim nonzero membership in more than one outcome category.

Traditional crisp implication functions have been adapted for fuzzy logic (e.g., Klir & Yuan, 1995; Tsoukalas & Uhrig, 1997). We examined implication functions for fuzzy SDT based on maximum and minimum values, as well as a multiplication function. All had certain undesirable features precluding their use in assigning degrees of category membership for use in fuzzy SDT. However, a mixed set of functions,

TABLE 2: Truth Table for Standard (Crisp) SDT

Signal (*s*)	Response (*r*)	Hit	FA	Miss	CR
0	0	0	0	0	1
0	1	0	1	0	0
1	0	0	0	1	0
1	1	1	0	0	0

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

combining both the maximum and minimum functions, possesses properties suitable for application of fuzzy SDT. The set involves a combination of different implication functions, using maximum functions for the error categories (Miss and FA) and minimum functions for the correct response categories (Hit and CR). The membership values for the four decision outcomes, Hit, Miss, FA, and CR, are defined by the following functions:

Hit;	$H = \min(s, r)$	(6a)
Miss;	$M = \max(s - r, 0)$	(6b)
False alarm;	$FA = \max(r - s, 0)$	(6c)
Correct rejection;	$CR = \min(1 - s, 1 - r)$.	(6d)

Table 3 shows a fuzzy SDT truth table for several selected values of s and r. As an example of the results of this function set, suppose that $s = .8$ and $r = .9$. That is, the state of the world strongly, but not absolutely, points to a signal, and the observer strongly responds that a signal is present. Applying equations 6a–6d, the resulting category memberships are H = .8, M = 0, FA = .1, and CR = .1. Hence the outcome strongly points to a hit, but unlike crisp SDT, there is also some membership in the FA category, representing the fact that the response was stronger than what was called for by the signal. The CR category is also nonzero, reflecting the small membership of the event in the *noise* category and the fact that an unequivocal *yes* response was not made. Note also that if both s and r are forced to be binary (e.g., 0 or 1), then equations 6a–6d will result in a reversion to the crisp truth table shown in Table 2.

To get an intuitive appreciation of Equations 6a–6d, consider the minimum (Hit, CR) and maximum (Miss, FA) functions separately. The minimum functions can be thought of as reflecting the degree of overlap (or intersection) of the fuzzy sets of signal and response (Hits) and of that of nonsignal and nonresponse (CR). For hits, for example, one can think of hit membership as reflecting the degree to which the response set overlaps the signal set. If the overlap is perfect, then the Hit membership takes on the value of the signal (and the response) membership. If the overlap is less than complete, Hit membership is lower and takes on the value of the signal or response, whichever is smaller.

The same reasoning applies to CR membership. For Miss and FA, the function reflects the rationale that the degree of overresponding or underresponding (with respect to the signal) defines the degree of FA and Miss membership, respectively. The rationale for the use of

TABLE 3: Example Truth Table for Fuzzy SDT

Signal (s)	Response (r)	Hit	FA	Miss	CR
0.8	0.9	0.8	0.1	0	0.1
0.1	0.2	0.1	0.1	0	0.8
0.1	0.1	0.1	0	0	0.9
0.1	0.9	0.1	0.8	0	0.1
0.2	0.2	0.2	0	0	0.8
0.2	0.1	0.1	0	0.1	0.8
0.2	0.3	0.2	0.1	0	0.7
0.3	0.2	0.2	0	0.2	0.7
0.3	0.5	0.3	0.2	0	0.5
0.3	0.9	0.3	0.6	0	0.1
0.5	0.2	0.2	0	0.3	0.5
0.5	0.5	0.5	0	0	0.5
0.5	0.9	0.5	0.4	0	0.1
0.75	0.1	0.1	0	0.65	0.25
0.75	0.75	0.75	0	0	0.25
0.75	0.8	0.75	0.05	0	0.2
0.9	0.1	0.1	0	0.8	0.1
0.9	0.9	0.9	0	0	0.1
0.9	0.8	0.8	0	0.1	0.1

the maximum is as follows. If the response is closer to a full *yes* ($r = 1$) than the signal is, then the event membership should fall in the FA category to some nonzero degree but should not have any membership in the Miss category. This follows because $r > s$. On the other hand, if the response is closer to a full *no* than the signal is ($r < s$), then the event membership should fall in the Miss category to some nonzero degree but should not belong with any degree to the FA category.

Table 3 reveals some interesting features of fuzzy SDT membership degrees. First, there is always at least one zero value among the categories. (In contrast, as Table 2 shows, in crisp SDT a given event is always associated with three zero and one unit membership values). In fact, zero values occur only in the Miss category, the FA category, or both, for any given event, and a zero *must* occur in at least one of the two. This outcome follows from the face validity assumption described earlier.

Note that when $r = s$, both Miss and FA have a value of zero, because the response was exactly the correct strength. For example, when $r = s = .9$, the event is mostly a hit (.9), partly a CR (.1), and not at all a Miss or FA (both 0). Also notice that the values of the four categories always sum to 1. This is another sensible and desirable result, because in crisp SDT the four categories are mutually exclusive. In our temperature example it is reasonable for a given value of t to belong, say, to the category *cold* to the extent of .8 and to *cool* with .3, for a total value greater than 1, as these two descriptors can overlap. But it is not reasonable for the memberships of the Hit, Miss, FA, and CR outcomes to sum to greater than 1. Put another way, although fuzzy SDT permits $\mathbf{A} \cap \bar{\mathbf{A}}$ to be nonzero, the sum of the four mutually exclusive outcome categories should be 1, because even in fuzzy SDT the four categories represent the full universe of possible outcomes; together they encompass the "whole truth."

Relating Membership Functions for Decision Outcomes to States of the World (SW)

First, the mapping functions map SW to the signal s and RV to the response r. Next, membership values for decision outcomes are computed by using the mixed-implication functions with the values for s and r. It is then relatively straightforward to examine the relationship between the membership functions for the decision outcomes and the SW. Given an analytical mapping function, a family of curves for the different outcomes can be generated by varying the response variable r.

Figures 3A and 3B plot membership functions for Hits, Misses, FAs, and CRs as a function of the SW variable a. The parameter is the response r. (Discrete r values of 0, .3, .5, .7, and 1 are shown for illustrative purposes only. In actuality a continuous family of curves would be generated, as r takes on any value between 0 and 1.) These outcome categories were generated using the sigmoid mapping function for aircraft separation distances shown in Figure 2 ($s = 1/[1+(a/10)^5]$). Because the a–to–s mapping is nonlinear, so are the functions describing the outcome categories. These four functions take on a shape that reflects the shape of the original function mapping the SW to s, but the value of r puts a limit on the maximum membership in Hit and FA, whereas $(1 - r)$ represents the maximum possible membership for Miss or CR.

The membership functions show that in general, the membership values for Hit and FA trade off against each other, as do those for CR and Miss. Consider Figure 3A. For any given value of r (except $r = 0$), the membership value for Hit decreases with an increase in a (greater aircraft separation), whereas that for FA increases. This makes intuitive sense because as a increases, the event (aircraft separation) becomes less signal-like and more noise-like, so that the outcome of a response, irrespective of its strength, will become less Hit-like and more FA-like. Also as Figure 3A shows, for any given value of a, both Hit and FA memberships increase with increased r. This is again reasonable because for any given signal value, as response strength increases, the outcome should become both more Hit-like and more FA-like. These patterns are similar for the CR and Miss memberships, except that the variations with a and r are inverted (see Figure 3B).

Figure 3 also shows that membership values for Hit and Miss do not sum to 1; that is, $H + M \neq 1$. By the same reasoning, $CR + FA \neq 1$. Recall, however, that the sum of the member-

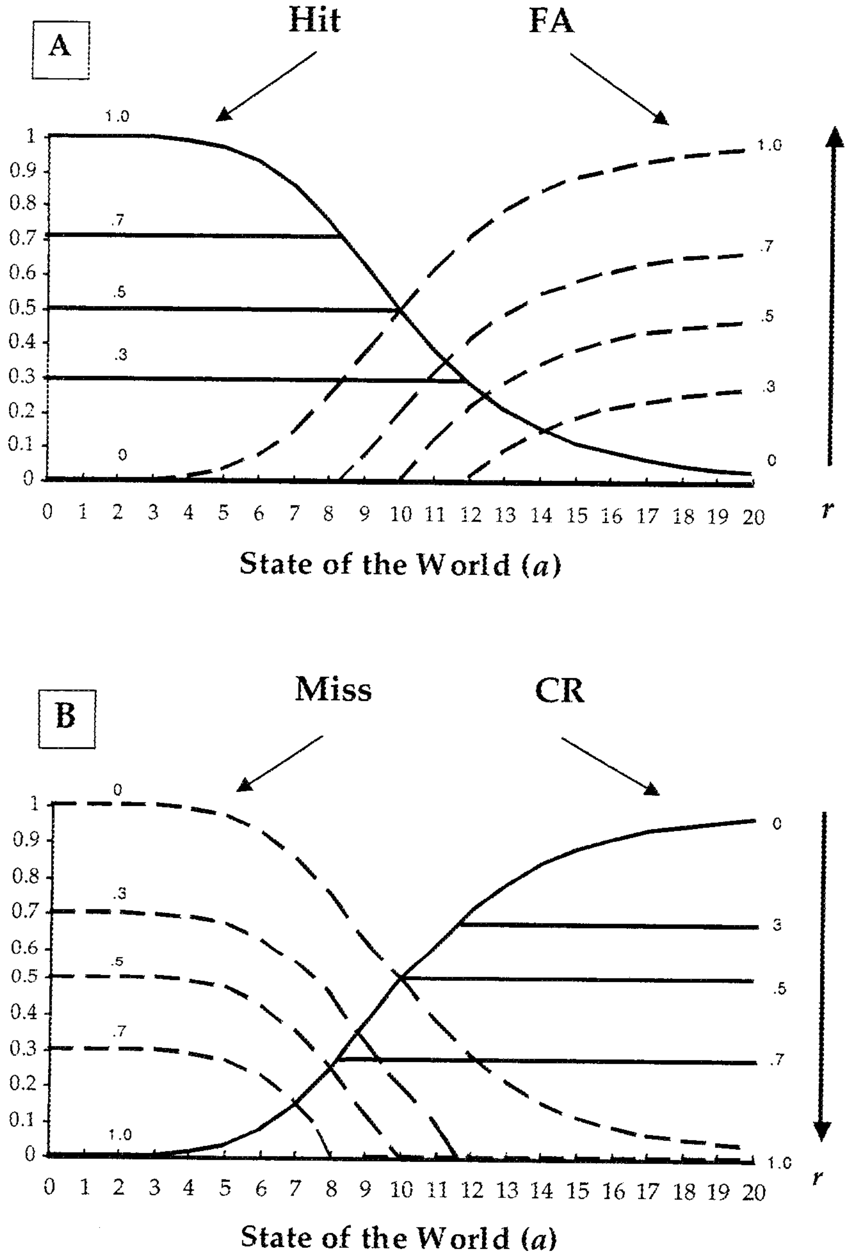

Figure 3. Fuzzy Hit, Miss, FA, and CR membership values with sigmoid mapping of *s* to *a*.

ship values of all four outcomes is always 1; that is, H + M + CR + FA = 1.

Figure 4 gives the outcomes if the SW-to-*s* mapping is linear. These figures are included for illustrative purposes only and most likely would not reflect a suitable SW-to-*s* mapping for air traffic control (but they might be appropriate for another application). Both Figures 3 and 4 show the correspondence between the original mapping function's shape and that of the outcome category functions, as well as showing that the limits set by *r* and (1 – *r*) are independent of the shape of the original mapping function.

Computation of Fuzzy Hit and False Alarm Rates

We now turn to a discussion of how to

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

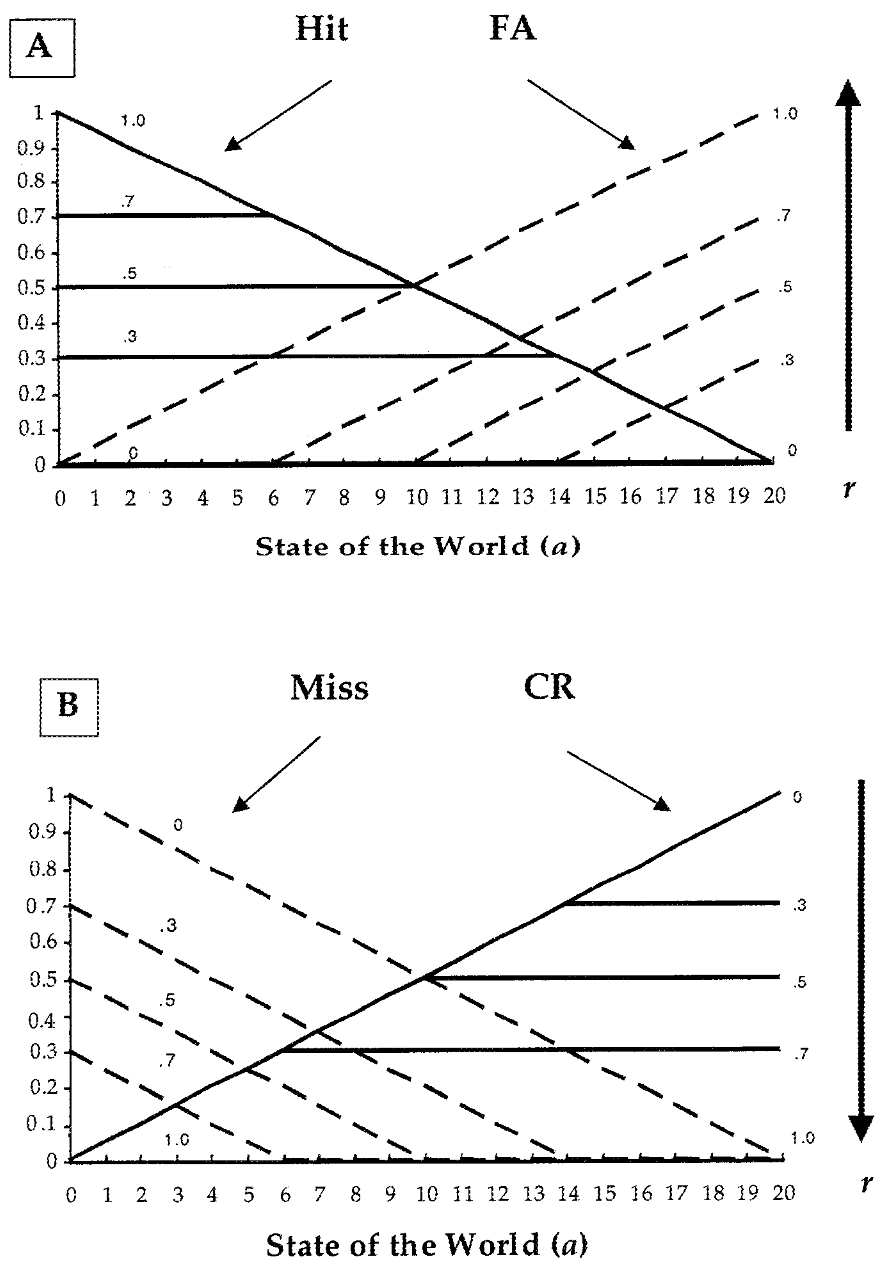

Figure 4. Fuzzy Hit, Miss, FA, and CR membership values with linear mapping of *s* to *a*.

compute some of the standard SDT measures when the signal and response are defined fuzzily. The rates for each of the four decision outcomes (HR, MR, FAR, CRR) can be calculated by summing the membership values in each category over all trials and dividing by the sum across trials of membership values in signal (s) or not-signal ($1-s$). Formally, all four decision outcomes can be computed as follows,

$$HR = \Sigma(H_i)/\ \Sigma(s_i) \text{ for } i = 1 \text{ to N} \quad (7a)$$
$$MR = \Sigma(M_i)/\ \Sigma(s_i) \text{ for } i = 1 \text{ to N} \quad (7b)$$
$$FAR = \Sigma(FA_i)/\ \Sigma(1 - s_i) \text{ for } i = 1 \text{ to N} \quad (7c)$$
$$CRR = \Sigma(CR_i)/\ \Sigma(1 - s_i) \text{ for } i = 1 \text{ to N} \quad (7d)$$

in which i is the trial number, N is the total number of trials and H_i is the degree of Hit for trial i, M_i is the degree of Miss for trial i, and so on.

The $\Sigma(1 - s_i)$ term, although formally com-

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

puted by adding $(1 - s_i)$ for each trial i, can be more efficiently calculated as $N - \Sigma(s_i)$ for $i = 1$ to N. By substituting Equations 6a–6d for the membership values of H_i, M_i, and so on, Equations 7a–7d can be rewritten as:

$$HR = \Sigma(\min(s_i, r_i))/\Sigma(s_i) \text{ for } i = 1 \text{ to N} \quad (8a)$$
$$MR = \Sigma(\max(s_i - r_i, 0))/\Sigma(s_i) \text{ for } i = 1 \text{ to N} \quad (8b)$$
$$FAR = \Sigma(\max(r_i - s_i, 0))/\Sigma(1 - s_i) \text{ for } i = 1 \text{ to N} \quad (8c)$$
$$CRR = \Sigma(\min(1 - s_i, 1 - r_i))/\Sigma(1 - s_i) \text{ for } i = 1 \text{ to N}. \quad (8d)$$

Either Equations 7a–7d in conjunction with Equations 6a–6d, or Equations 8a–8d alone, can be used in any application in which s is manipulated experimentally (or varies naturally) over several trials and separate detection responses are made on each trial. Note that the redundancy between Hit and Miss rates (i.e., HR + MR = 1), and between FA and CR rates (FAR + CRR = 1), holds true for both crisp and fuzzy SDT. (This redundancy should not be confused with the lack of redundancy between membership values on individual trials for Hit and Miss and between values for CR and FA.)

When we use fuzzy SDT Equations 6a–6d and 7a–7d, or Equations 8a–8d to aggregate performance across cases or trials, the method follows directly from that used for crisp SDT. The total Hit and Miss memberships for all trials are summed and divided by the sum of the signal memberships for all trials. For the FA and CR rates, which are dependent on the degree to which each trial did *not* contain a signal, the sum of memberships should be divided by the sum of the extent to which each trial was not a signal (for each trial i, this is equal to $[1 - s_i]$). Applying Equations 7a–7d to the data set presented in Table 1 yields the following. We obtain $\Sigma(s_i) = 9.05$ and $\Sigma(1 - s_i) = 19 - 9.05 = 9.95$. Just as in the crisp case, the Miss rate will be $(1 - HR)$, and the CR rate is $(1 - FAR)$. The sums of the four outcome categories are 7.00, 2.05, 2.35, and 7.60 for Hit, Miss, FA, and CR, respectively. The degrees of membership for each category can now be computed by entering these values into Equations 7a–7d.

$$HR = 7.00/9.05 = .773$$
$$MR = 2.05/9.05 = .227$$
$$FAR = 2.35/9.95 = .236$$
$$CRR = 7.60/9.95 = .764$$

The same calculations could be done directly from the individual s and r values using Equations 8a–8d. With these calculations completed, standard SDT parameters (sensitivity, bias) can be computed just as with crisp SDT. In the next section we continue the analysis begun in this section with our sample data set, and discuss an alternative method for assessing sensitivity in fuzzy SDT.

Sensitivity and Criterion in Fuzzy SDT

In traditional SDT the parametric estimate of operator sensitivity, d', represents the standardized distance between the normal curve approximating the signal distribution and the one approximating the noise distribution. In other words, a curve is approximated for the frequency of each level of the evidence variable given that there is a signal, and another curve is approximated given that there is noise. The d' measure is the horizontal distance, in units of standard deviations, between the curves. Another way of saying this is that the placement of the observed Hit rate on a normal distribution, minus the placement of the observed FA rate on a normal distribution, will be equivalent to the distance between the signal and noise distributions and will provide an estimate of the detector's sensitivity. The d' parameter is most easily calculated by determining the difference between the standardized Hit and FA rates using the following formula:

$$d' = Z(HR) - Z(FAR). \quad (9)$$

Because the fuzziness of the signal has already been captured in the definitions of s, and r, and from them, fuzzy HR, FAR, MR, and CRR, the traditional d' formula can be used in fuzzy SDT analysis. We illustrate the calculation of d' using the sample data from Table 3. For these data, HR is .773 and FAR is .236. Taking the normal deviates of these values and using equation 9 gives $d' = 1.47$.

Another way to assess d' when s is fuzzy is through integration. In traditional SDT, d' is the horizontal distance, in units of standard deviations, between the means of the Gaussian probability density distributions for signal and noise. In crisp SDT the noise curve is essentially the curve for $s = 0$, whereas the signal curve is for $s = 1$. In fuzzy SDT the value of s can fall anywhere between 0 and 1; therefore, another appropriate way to assess sensitivity is to determine how far, in standard deviation units, all the possible curves sit from each other. In other words, theoretically each of the infinite curves between $s = 0$ and $s = 1$ would be approximated. Using the $s = 0$ curve as a reference, the contribution of each of the infinite curves to d' would be weighted by the distance of its s value from 0, and the integral would be calculated:

$$d' = \int_0^1 [\mathrm{Z}(x \text{ value at } s = q) - \mathrm{Z}(x \text{ value at } s = 0)]\, \mathrm{d}q, \tag{10}$$

in which x is the value of the evidence variable or SW. If s had only a finite set of values, then the distances of each of the existing s curves from $s = 0$ would be computed.

$$d' = \Sigma\, [\mathrm{Z}(x \text{ value at } s = q_i) - \mathrm{Z}(x \text{ value at } s = 0)] \text{ for } i = 1 \text{ to N}, \tag{11}$$

in which x is the value of the evidence variable or SW, q_i represents each of the discrete values of s, *and* N is the total number of discrete s values.

The traditional SDT criterion measure, β, can also be calculated using fuzzily derived HR and FAR. We define β in the standard way as

$$\beta = \mathrm{Y}(HR)/\mathrm{Y}(FAR), \tag{12}$$

in which Y(.) represents the ordinate of the normal distribution. For the fuzzy HR and FAR derived in our sample data set, $\beta = 0.98$.

Fuzzy SDT Correlation Analysis

A completely different method for assessing the accuracy of detection performance when s, r, or both, are fuzzy is correlation. The computed correlation coefficient between signal degree and response degree can be an indication of the accuracy of detecting a fuzzy signal. A high correlation would indicate a linear relationship between r and s. Correlational analysis has some advantages. If the mapping from SW to s or from RV to r is erroneous for any reason, then accuracy as computed by fuzzy SDT may be reduced compared with other methods, whereas correlation would still capture the linear nature of the relationship.

The correlational analysis can best be described with an example. Consider the sample s and r data shown in Table 4. The responding system tends to generate r values toward the middle range of the spectrum, even for more extreme values of s. The fuzzy SDT implication functions (Equations 6a–6d) would assign relatively high Miss and FA values for some of the events (e.g., $s = .1$ and $r = .35$ would result in an FA membership of .25). However, because the relationship between s and r is almost perfectly linear, as shown in Figure 5, the computed correlation is very high (greater than .99).

This dissociation between the results of the various fuzzy SDT analysis methods can have different consequences, depending upon the application. It might be that the implication function is a better measure of performance, if indeed it was desirable for s to closely match r. Alternatively it could be that the mapping strategy used to convert RV to r was too prone to assigning middle values and did not make enough use of the extreme (near 1 or 0) r values to be reflective of the question of interest.

Correlation analysis can also pose some danger, similar to that discussed for the implication functions. For example, it too will penalize the detector for overly crisp definitions of signal and response, because a restricted range on one or both variables can deflate a correlation coefficient. Fortunately, methods exist for correcting this problem when it occurs. For example, if the response is always defined in a binary fashion, point-biserial correlations can be used as a more appropriate measure of correlation.

Fuzzy Signals, Crisp Responses

Fuzzy SDT analysis can also be extended to cases in which the signal is fuzzy but the response is discrete or binary; for example, *yes* or *no*. This is particularly useful when considering the analysis of *actions* that are contingent on a particular decision choice. For example, a

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

TABLE 4: Sample *s* and *r* Data to Illustrate Differential Results of Implication Functions Versus Correlation in Measuring Sensitivity in Fuzzy SDT

s	*r*
0	0.3
0.1	0.35
0.2	0.4
0.3	0.45
0.4	0.475
0.5	0.5
0.6	0.525
0.7	0.55
0.8	0.6
0.9	0.65
1.0	0.7

participant in a memory study may decide with a response $r = .8$ that he has been shown a particular face before by the experimenter. If the experimenter then requires the participant to pick the face seen before from a group of distractors, the participant must take a discrete action consistent with his decision choice.

Outside the laboratory, too, human and machine decision choices often need to be translated into discrete actions. For example, an automated fault management system in a nuclear power plant may make a judgment of $r = .9$ (e.g., probability, confidence, or severity) that a fault is present. This fault then needs to

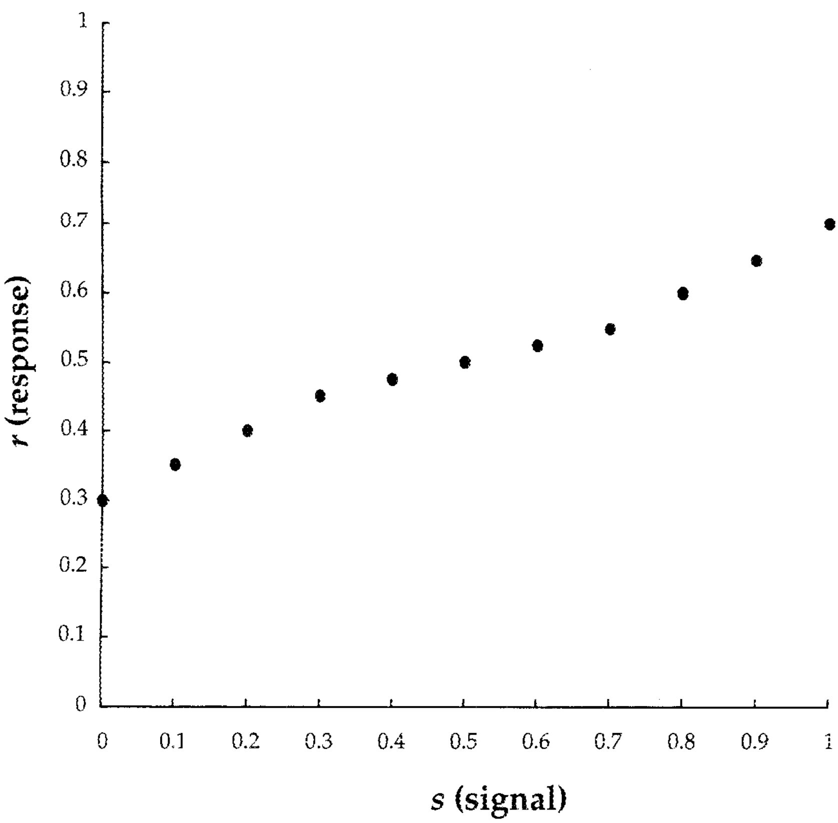

Figure 5. Plot of sample *s* and *r* values illustrating linear relationship.

be translated into a binary action, that is, to turn off the plant or not. In general, an overt response is more difficult to "fuzzify" than the decision about the state of the world upon which the overt action is based.

In a recent comparative analysis of the information-processing performance of humans and automation, Parasuraman, Sheridan, and Wickens (2000) also distinguished decision making from action selection as separate stages of any detection process. Although they allowed action levels to be continuous, they suggested that in most cases actions will be categorical, even when decision making has multiple levels.

Table 5 shows the result of applying the implication functions (Equations 6a–6d) to sample data with $r = 0$ or 1. A binary response to a fuzzy signal always results in the event's having nonzero membership in exactly two of the categories. If the binary response is yes ($r = 1$), then Equations 6a through 6d reduce to:

$$H = s \tag{13a}$$
$$M = 0 \tag{13b}$$
$$FA = 1 - s \tag{13c}$$
$$CR = 0. \tag{13d}$$

Note that the available membership degree of 1 is assigned partly to Hit in the amount of s and the rest, $1 - s$, to FA. If the binary response is no ($r = 0$), then

$$H = 0 \tag{14a}$$
$$M = s \tag{14b}$$
$$FA = 0 \tag{14c}$$
$$CR = 1 - s. \tag{14d}$$

In this case, the available membership degree of 1 is assigned partly to Miss in the amount of s and the rest, $1 - s$, to CR.

Equations 13 and 14 can be used to compute detection statistics for the special case of fuzzy signals with crisp responses. Some additional special cases should also be noted. For example, if r values are allowed to vary along the full range [0, 1] and the observer simply chooses to make binary responses, then the Hit, Miss, FA, and CR rates, as computed using Equations 13 and 14, are appropriate. When aggregated, they will give measures of accuracy and bias that can be validly compared with other observers or systems. However, if the nature of the response is such that its very definition constrains the possible responses to only 0 or 1, then fuzzy SDT analysis could underestimate accuracy. This occurs because the observer is penalized by being forced to use binary responses when the signal is fuzzy, and the penalty increases as s becomes fuzzier (is closer to .5).

To see this, suppose $s = .6$, but r can be only 0 or 1. The permissible response closest to the actual s value is 1, which results in H = .6, M = 0, FA = .4, and CR = 0, or a sensitivity that is lower than could be achieved with a fuzzy response (e.g., $r = .6$). This outcome renders invalid any comparisons of fuzzy SDT parameters between systems or situations in which r is free to vary across the full range [0, 1] and those in which r is only allowed to be 0 or 1.

However, the results from fuzzy SDT analysis of such a situation are, at worst, no less detrimental to calculated accuracy measures than the outcomes of a traditional forced-choice

TABLE 5: Sample Hit, Miss, FA, and CR Outcomes for Fuzzy SDT Analysis When the Response r Is Binary (0 or 1)

s	r	Hit	Miss	FA	CR
0.1	1	0.1	0	0.9	0
0.3	1	0.3	0	0.7	0
0.5	1	0.5	0	0.5	0
0.7	1	0.7	0	0.3	0
0.9	1	0.9	0	0.1	0
0.1	0	0	0.1	0	0.9
0.3	0	0	0.3	0	0.7
0.5	0	0	0.5	0	0.5
0.7	0	0	0.7	0	0.3
0.9	0	0	0.9	0	0.1

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

paradigm in which both signal and response are constrained to binary values. Furthermore, allowing the definition of *s* to be fuzzy will improve estimates of bias by capturing the variability in signal strength, making it possible to assess, for example, the extent to which stronger signals are more likely to generate a response.

The only potential problem remaining with this type of situation is as follows: When the domain involves fuzzy *s* and binary *r*, attempts to compare different operators or different situations within the same system may be confounded by the "luck of the draw." That is, if operator P happens to get a larger number of the "more fuzzy" signals (e.g., *s* values of .4, .5, .6) than operator Q, then it will be more difficult for P to achieve high HR and CRR values overall.

Two solutions exist for this situation. First, one might reconsider whether there might be some way to fuzzify *r*, such as using confidence or probability ratings, even if a binary action is taken. Second, if this is not possible or desirable, the *s* value can be rounded to 0 or 1, and traditional crisp SDT analysis conducted. Note that *rounded* does not have to mean that an *s* of .5 or more rounds to 1, whereas an *s* of less than .5 rounds to 0. The cutoff can be placed anywhere that it makes sense to do so, according to the needs of the particular situation.

It is also possible to conduct an analysis in which trials with middle values of *s* are excluded from analysis, high *s* values are classed as signal-present, and low *s* values as signal-absent. Alternatively, the cutoff can be placed at one point and the traditional SDT analysis conducted, then the cutoff can be moved and another traditional SDT analysis conducted, and so on. Such a moving-cutoff strategy is analogous to the use of confidence intervals in traditional SDT – or for some analyses, not merely analogous but equivalent, because confidence interval ratings can be converted to *r*.

DISCUSSION

Summary of Fuzzy SDT Analysis

The fuzzy SDT analysis developed in this paper involves four main steps: (a) selection and application of mapping functions for states of the world and responses; (b) use of mixed-implication functions to assign degrees of membership in the decision outcomes of hits, misses, false alarms, and correct rejections; (c) computation of fuzzy hit, miss, false alarm, and correct rejection rates; and (d) computation of fuzzy sensitivity and bias measures.

The first step is the selection of mapping functions. The signal mapping function maps variables that describe the states of the world into the fuzzy signal set with membership strength *s*. We described several ways in which such functions could be selected. Mapping functions will vary with the application and can be relatively simple, through the use of a single variable (as in the ATC example; see Figure 2), or could be more complex by being based on many variables.

Mapping functions can also be discrete or continuous and can be derived empirically or based on some theoretical premise. Among continuous analytical functions, the sigmoid may best capture the variability inherent in many applications. Selecting the response mapping function is relatively straightforward and can be based on judgments of confidence, on the reported strength, severity, or criticality of the signal, or both.

The second step is the use of mixed-implication functions for assigning degrees of membership in the conventional SDT outcomes of hits, misses, false alarms, and correct rejections. We used two criteria for choosing appropriate functions: comparability with crisp SDT and face validity. The implication functions we propose, which combine minimum and maximum functions, meet both criteria and seem suitable for a wide range of applications. However, we make no claim that these are the best functions or that other functions may not be more suited to particular applications. Can the implication functions be optimized? Perhaps. One possibility would be to use Monte Carlo methods or simulation with a data set with well-defined states of the world and outcomes, and choose a set of functions that best matches the defined outcomes.

The third step in fuzzy SDT analysis involves computing the mean fuzzy hit and false alarm rates by weighting the membership degrees in these outcome categories by the average

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

degrees of membership in the signal (s) and noise ($1 - s$) sets. The fourth step, computation of sensitivity and bias measures, then follows in exactly the same manner as it would in crisp SDT.

Crisp and Fuzzy SDT

SDT could arguably be viewed as one of the most robust and useful quantitative theories in psychology. The familiar adage, "there is nothing so practical as a good theory," is also particularly well suited to SDT, given the variety and range of the practical applications of SDT in psychology, human factors, and other fields (Swets, 1996). In the present paper we have proposed that combining fuzzy logic and SDT can further enhance the range, power, and utility of SDT. Nevertheless, it is instructive to examine the points of convergence of crisp and fuzzy SDT.

The approach to fuzzy SDT we have outlined in this paper shares some features with the *uncertain signal* problem. This refers to the case in which there is uncertainty regarding some dimension of the signal – for example, its duration, starting time, frequency, location, or other characteristics – so that the signal is known statistically (SKS). The case of no uncertainty is referred to as a signal known exactly (SKE; Tanner & Birdsall, 1958). A number of researchers have compared detection performance for SKS and SKE for a number of different dimensions of signal uncertainty, the general finding being that uncertainty reduces detectability (Egan, Greenberg, & Schulman, 1961; Green & Swets, 1966; Swets, 1984). SKS could be linked to the mapping function concept of fuzzy SDT. For example, one could develop a function that mapped all sources of uncertainty regarding the signal onto the signal strength parameter, s.

In this sense, fuzzy SDT, and the SKS case in SDT can be considered to be related. In another respect, however, the two are distinct, particularly in nonlaboratory conditions in which there is uncontrolled variability in the signal. In the SKS case, the mapping of SW to the signal is nevertheless crisp, and the uncertainty is introduced in its presentation to the observer. For example, an SKS experiment in the laboratory could involve detection of a 1000-Hz tone presented at random times (Egan et al., 1961). Although the time of presentation of the signal is uncertain, the definition of the signal, in terms of the SW-to-signal correspondence, is crisp. In contrast, in the fuzzy SDT cases we consider, the SW-to-signal correspondence is uncertain. An example might be a sonar operator listening to sonar returns consisting of multiple frequencies: A signal indicating the presence of a submarine is not defined crisply as a 1000-Hz tone but as an auditory stimulus whose frequencies vary with context and over time.

Previous Related Work on Fuzzy Logic

Fuzzy logic has been applied to a wide array of problems in psychology (e.g., Massaro, 1998) and human factors (e.g., Moray et al., 1987). For the most part, however, fuzzy logic has not previously been explicitly combined with SDT to develop a general model of fuzzy detection performance, as presented in this paper. Nevertheless, there has been some relevant prior work; for example, a study of word recognition by Meng and Li (1990). Although this study did use fuzzy logic and SDT, it did not present a general model of fuzzy SDT as we do.

There has also been some relevant work on machine classification of noisy signals. Boston (1997) carried out an analysis of the classification of event-related brain potentials (ERPs). He used fuzzy logic to define the degree to which the values of a potential signal on two different ERP features reflected a true ERP signal. He contrasted the judgments regarding signal presence resulting from fuzzy analysis to the performance of a Bayesian detection algorithm. But the end result of his analysis involved rounding out the decision to relatively crisp outcomes: signal present, signal absent, or uncertain.

Although these studies did consider fuzzy logic in the context of particular signal detection problems, they stopped short of developing the general case of fuzzy SDT, which retains the continuous definition of signal, response, or both, resulting from contextual and temporal variability. The methods we presented in this paper allow one to carry the fuzziness through all the stages of analysis, rather than reverting back to categorical characterizations of the outcome. Of course, there are prior exam-

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

ples of the use of nonbinary responses to improve the precision of SDT analysis, such as the use of confidence levels surrounding a yes-no judgment (Green & Swets, 1966), exploring the costs and benefits when a decision criterion is set at different levels (Lehto, Papastavrou, & Giffen, 1998; Parasuraman, Hancock, & Olofinboba, 1997), and Balakrishnan's (1998) method for computing response bias by comparing the types of events (signal vs. noise) corresponding to each self-rated degree of confidence.

Although these and other methods have acknowledged the existence of continua in SW, in the response of a decision maker, or in both, SDT analyses and the resulting interpretations of sensitivity, bias, and payoffs have generally assumed crisp states of the world and crisp responses.

The fuzzy SDT model we have outlined is also conceptually related to fuzzy logic models of perception and cognition that have been developed by cognitive psychologists. Probably the best known of such models is the fuzzy logical model of perception (FLMP) of Massaro (1988; Massaro & Friedman, 1990). The FLMP model assumes that perceptual tasks involve three stages: feature evaluation, integration, and decision making. Fuzzy set memberships are used to derive the relationships between feature values and their integration prior to the decision stage. In this respect, the model is similar to Bayesian models of information integration, although the use of fuzzy membership values (or fuzzy truth values) is unique to the model and distinct from the subjective probabilities that are used in Bayesian analysis. Given that SDT originated as a perceptual theory, links between fuzzy SDT and the FLMP model are to be expected. Moreover, Massaro has advocated the use of "graded" and "expanded" factorial designs in testing models of perception and cognition. This method refers to the use of all possible combinations of graded values of stimulus features, as well as the presentation of stimuli possessing each feature value in isolation (Massaro & Hary, 1986). Massaro and Friedman also recommended testing the FLMP model by using ratings rather than categorical responses.

In some respects, therefore, there are links between the FLMP model and fuzzy SDT. The two are nevertheless quite distinct. FLMP involves the integration of *perceived* information derived from multiple stimulus features that match stimulus prototypes in a fuzzy manner. It uses an integration procedure functionally equivalent to Bayesian analysis. Fuzzy SDT, although it can be based on what is perceived (evidence variable), is generally concerned with the analysis of objectively measured stimuli (states of the world) that map onto the signal set in a fuzzy manner. It uses mixed-implication functions to derive fuzzy hit and FA rates, followed by conventional SDT analyses.

Our intention in proposing fuzzy SDT is not iconoclastic; rather, it is to allow the extension of SDT when it would be useful to do so. As regression is the global case of analysis of variance, so fuzzy SDT is the global case of SDT, improving the precision of signal detection analyses by retaining the information provided by the middle ground, rather than by rounding it into oblivion. Of course, situations still exist in which the simplicity of the variables at hand makes the crisp SDT method more parsimonious without a significant loss of information. We conclude, however, that fuzzy SDT is particularly well suited to evaluation of detection performance when there is significant variability in the state of the world that defines the signal to be detected.

Contextual and Temporal Variability in Signals

There can be both contextual and temporal variability in signal definition. Contextual variability refers to the dependence of signal definition on situation-specific factors (e.g., see the previously discussed example of signals in ATC depending on the type of flights, whether level cruise or climbing/descending). An invariant crisp definition of a signal may result in erroneous conclusions when these factors are present. To the extent that the factors can be encoded into a mapping function as we describe in Step 1 of our fuzzy SDT analysis, this form of analysis can capture the contextual variability better than crisp SDT.

Temporal variability refers to the variation in signal strength over time. In many real settings, the signal is not a discrete event in time but represents an event that unfolds over time.

In industrial process control, for example, the value of a critical system variable (e.g., core temperature in a power plant) will fluctuate over time but may gradually drift toward a dangerously high value if an emergency condition is developing (Moray, 1986). In medicine, to take another example, detection of emergency conditions during surgery requires analysis of temporal patterns and trends in the values of vital patient signs (Lowe et al., 1999). In fuzzy SDT terminology, the value of s may start low and increase over time, or start high and decrease over time.

This type of signal variability over time has not been explicitly modeled in psychological studies. This is a curious omission because psychologists have developed sequential sampling models of the perceived strength of a discrete signal (e.g., Laming, 1968; Link & Heath, 1975). In most such models, the perceived signal strength grows with time following the presentation of the discrete stimulus until some threshold for response is reached. In models of the discrimination between two stimuli, S1 and S2, the evidence makes a "random walk" between S1 and S2 until the threshold for one alternative is crossed.

Sequential sampling and random-walk models have been used to examine the trade-off between detection accuracy and speed of response. The reason that these and related models have not been extended to analysis of actual rather than perceived temporal variability in signal strength is that most laboratory studies of detection and discrimination involve discrete rather than continuous signals.

Models of temporal variability of the states of the world associated with a signal could be developed or empirically defined. The influence of temporal variability could then be examined by extending the mapping function in a particular application by time, T. If the signal mapping function is $s = f(SW)$, then an extended mapping function would be of the form $f(SW, T)$. We propose that the ability to incorporate contextual and temporal variability in signal definition into the analysis of detection performance is one of the major advantages of fuzzy SDT. Conventional SDT can be thought of as providing a snapshot of detection performance at a particular time. Although several such snapshots could be obtained in successive crisp SDT analyses at discrete points in time, fuzzy SDT provides a natural way for analyzing continuous signals by incorporating temporal variability into the mapping function.

Applications

Many domains in which SDT has been applied are amenable to fuzzy SDT analysis. Fuzzy SDT should prove useful in basic studies of perception, memory, and cognition. In many laboratory studies, the strength of a signal is varied as an independent variable. For example, in perceptual experiments, participants may be presented with stimuli of different size, luminance, or orientation. Depending on the question being addressed, it may be appropriate to map each level of an independent variable involving signal strength to an s value and conduct fuzzy SDT analysis on detection rates across conditions, weighted by the s value of the given condition. This would provide a clearer picture of how sensitivity and bias are affected by the other independent variables being studied in the experiment.

Fuzzy SDT is also particularly well suited to applied studies and may better capture the shades of gray inherent in almost any real-world domain. In some cases dichotomization of the signal, the response, or both is necessary. As discussed previously, crisp responses (yes or no) are often necessary when translating decision choices into action. However, as long as the contextual and temporal variability of the state of the world can be captured (e.g., signal fuzzification only), the methods outlined herein can enrich the performance analysis of any decision-making system. Furthermore, as with crisp SDT, fuzzy SDT can be applied to the analysis of the detection performance of humans, machines, or joint human-machine "teams" (Parasuraman, 1987; Sheridan & Ferrell, 1974; Sorkin & Woods, 1985; Swets, 1996).

Some potential real-world applications of fuzzy SDT are briefly discussed here. Consider for example, the evaluation of collision warnings for automobiles and conflict warning systems for aircraft (Hancock & Parasuraman, 1992; Parasuraman et al., 1997). Farber and Paley (1993) discussed the idea of "acceptable" false alarms (FA), those being alerts that activate when the threshold for a collision is *almost* violated, thus reminding the user what the alert looks/sounds like and assuring him or her the system is still working. Traditional SDT would class this kind of alert as an FA, whereas fuzzy SDT would class it primarily as a Hit, with a small degree of FA, which would more closely reflect the fact that the event is not wholly undesirable. The fuzzy SDT analysis is similar to the concept of graded or *likelihood alarms* proposed by Sorkin, Kantowitz, and Kantowitz (1988) but provides a more general basis for determining the alarm degree that is displayed to the human operator.

Similar examples can be drawn from quality control and process control. In these applications a warning system may indicate when the state of a plant or a product has deviated from the norm by more than some cutoff value (Cale, Paglione, Ryan, Timoteo, & Oaks, 1998). However, in evaluating the performance of the detector, one would probably wish to know the extent to which false alarms (in crisp SDT parlance) were generated by events which were near to that cutoff point and would wish the same information for so-called misses. Fuzzy SDT would capture the extent to which false alarms, for example, had some degree of "Hit" by virtue of the event being close to the cutoff for a crisp signal, whereas crisp SDT would classify all alerts generated by events beyond the cutoff equally as false alarms.

An example from industrial psychology serves to illustrate the potential diversity of fuzzy SDT applications. Crisp SDT analysis has been used to evaluate personnel selection tests by ranking individuals who score above or below a cutoff score on a selection test of ability or knowledge, and determining whether they rate as high or low performers according to a subsequent assessment test or performance appraisal. However, a crisp SDT analysis of this type of data ignores gradations within high and low performance as well as ignoring how far above or below the cutoff score the employee scored. The use of fuzzy methods by Alliger et al. (1993) for evaluating employment interview data represented a step in this direction. Additional fuzzy mappings of employment or performance-related variables in order to conduct a full fuzzy SDT analysis would further enrich the techniques those authors have suggested.

CONCLUSIONS

In this paper the basic postulates of fuzzy SDT were presented and simple but powerful formulas were identified for conducting a fuzzy SDT analysis of detection performance. Examples were also provided of applications in basic and applied work. Of course, fuzzy SDT will ultimately stand or fall, to some degree, based on the results it generates in future applications in psychology, human factors, engineering, and other domains. We believe its membership in the set {stand} will be much greater than its membership in the set {fall}.

ACKNOWLEDGMENTS

This work was supported by Grant NAG-2-1096 from NASA Ames Research Center, Moffett Field, California. Kevin Corker was the technical monitor. We thank Neville Moray, Tom Sheridan, and Robert Sorkin for helpful comments on a previous version of this paper.

APPENDIX

Symbol Definitions

SW	State of the World; objective truth specifying a priori some physical state
s	Signal event in standard signal detection theory (SDT); also, degree of membership of state of the world (SW) in the fuzzy set signal in fuzzy SDT
n	Noise event in standard SDT
Y	*Yes*, or response of a detection system that a signal has occurred
N	*No*, or response of a detection system that a signal has not occurred
r	Degree of membership of a response in the *Yes* response set in fuzzy SDT
RV	Response Value; equivalent to *r*
P(*Y*|*s*)	Probability of a Yes response given that a signal occurred; also known as the hit rate (HR)
P(*Y*|*n*)	Probability of a Yes response given that

	a signal did not occur; also known as the false alarm (FA) rate (FAR)
$P(N\|s)$	Probability of a No response given that a signal occurred; also known as the miss rate (MR)
$P(N\|n)$	Probability of a No response given that a signal did not occur; also known as the correct rejection (CR) rate (CRR)
d'	Sensitivity of a detection system in standard SDT
β	Criterion or decision threshold of detection system
n_S	Number of possible discrete SW values
n_R	Number of possible discrete RV values
$f(SW)$	Mapping function for SW
$g(RV)$	Mapping function for RV
H	Degree of membership of a decision outcome in the fuzzy set *Hit*
M	Degree of membership of a decision outcome in the fuzzy set *Miss*
FA	Degree of membership of a decision outcome in the fuzzy set *FA*
CR	Degree of membership of a decision outcome in the fuzzy set *CR*

Fuzzy Implication Functions

Hit:	$\min(s, r)$
Miss:	$\max(s - r, 0)$
False alarm:	$\max(r - s, 0)$
Correct rejection:	$\min(1 - s, 1 - r)$

REFERENCES

Alliger, G. M., Feinzig, S. L., & Janak, E. A. (1993). Fuzzy sets and personnel selection: Discussion and an application. *Journal of Occupational and Organizational Psychology, 66*, 163–169.

Balakrishnan, J. D. (1998). Measures and interpretations of vigilance performance: Evidence against the detection criterion. *Human Factors, 40*, 601–623.

Baumgartner, R., Winddischberger, C., & Moser, E. (1998). Quantification in functional magnetic resonance imaging: Fuzzy clustering vs. correlational analysis. *Magnetic Resonance Imaging, 16*, 115–125.

Bellazzi, R., Silviero, C., Stefanelli, M., & De Nicolao, G. (1995). Adaptive controllers for intelligent monitoring. *Artificial Intelligence in Medicine, 7*, 515–540.

Boston, J. R. (1997). Effects of membership function parameters on the performance of a fuzzy signal detector. *IEEE Transactions on Fuzzy Sets and Systems, 5*, 249–255.

Brainerd, C. J., & Reyna, V. F. (1993). Memory independence and memory interference in cognitive development. *Psychological Review, 100*, 42–67.

Cale, M. L., Paglione, M., Ryan, H., Timoteo, D., & Oaks, R. (1998). *User Request Evaluation Tool (URET) conflict prediction accuracy report* (Technical Report FAA/CT-TN98/8). Atlantic City, NJ: FAA Technical Center.

Campbell, C. S., & Massaro, D. W. (1997). Perception of visible speech: Influence of spatial quantization. *Perception, 26*, 627–644.

Craiger, J. P., & Coovert, M. D. (1994). A fuzzy system for mapping worker attributes to tasks. *Behavior Research Methods, Instruments, & Computers, 26*, 107–111.

Egan, J. P., Greenberg, G. Z., & Schulman, A. I. (1961). Interval of time uncertainty in auditory detection. *Journal of the Acoustical Society of America, 33*, 771–778.

Ellison, J. W., & Massaro, D. W. (1997). Featural evaluation, integration, and judgment of facial affect. *Journal of Experimental Psychology: Human Perception & Performance, 23*, 213–226.

Farber, E., & Paley, M. (1993, April). *Using freeway traffic data to estimate the effectiveness of rear-end collision countermeasures.* Paper presented at the 3rd Annual IVHS America Meeting, Washington, DC.

Genaidy, A. M., Karwowski, W., Christensen, D. M., Vogiatzis, C., Deraiseh, N., & Prins, A. (1998). What is heavy? *Ergonomics, 41*, 420–432.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Hancock, P. A., & Parasuraman, R. (1992). Human factors and safety in the design of intelligent vehicle-highway systems. *Journal of Safety Research, 23*, 181–198.

Heiser, W. J., & Groenen, P. J. F. (1997). Cluster differences scaling with a within-clusters loss component and a fuzzy successive approximation strategy to avoid local minima. *Psychometrika, 62*, 63–83.

Karwowski, W., & Mital, A. (Eds.). (1986). *Applications of fuzzy set theory in human factors.* New York: Elsevier.

Klir, G. J., & Yuan, R. (1995). *Fuzzy sets and fuzzy logic: Theory and applications.* New York: Prentice-Hall.

Kosko, B. (1993). *Fuzzy thinking: The new science of fuzzy logic.* New York: Hyperion.

Kosko, B. (1997). *Fuzzy engineering.* New York: Prentice-Hall.

Kreifeldt, J. G., & Rao, K. V. N. (1986). Fuzzy sets: An application to warnings and instructions. In *Proceedings of the Human Factors Society 30th Annual Meeting* (pp. 1192–1196). Santa Monica, CA: Human Factors and Ergonomics Society.

Laming, D. R. J. (1968). *Information theory of choice-reaction times.* New York: Academic.

Lehto, M. R., Papastavrou, J. D., & Giffen, W. J. (1998). An empirical study of adaptive warnings: Human- versus computer-adjusted warning thresholds. *International Journal of Cognitive Ergonomics, 2*, 19–33.

Lehto, M. R., & Sorock, G. S. (1996). Machine learning of motor vehicle accident categories from narrative data. *Methods of Information in Medicine, 35*, 309–316.

Link, S. W., & Heath, R. A. (1975). Sequential theory of psychological discrimination. *Psychometrika, 40*, 77–105.

Lowe, A., Harrison, M. J., & Jones, R. W. (1999). Diagnostic monitoring in anesthesia using fuzzy trend templates for matching temporal patterns. *Artificial Intelligence in Medicine, 16*, 183–199.

MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide.* New York: Cambridge University Press.

Massaro, D. W. (1988). Ambiguity in perception and experimentation. *Journal of Experimental Psychology: General, 117*, 417–421.

Massaro, D. W. (1998). *Perceiving talking faces: From speech perception to a behavioral principle.* Cambridge: MIT Press.

Massaro, D. W., & Friedman, D. (1990). Models of integration given multiple sources of information. *Psychological Review, 97*, 224–252.

Massaro, D. W., & Hary, J. M. (1986). Addressing issues in letter recognition. *Psychological Research, 48*, 123–132.

Meng, Q., & Li, Y. (1990). An experiment to introduce set-valued statistics into the evaluation method of the signal detectability theory. [In Chinese]. *Acta Psychologica Sinica, 22*, 16–22.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception: Vol. 2. Cognitive processes and performance* (pp. 42.1–42.29). New York: Wiley.

Moray, N., King, K. R., Turksen, R., & Waterton, K. (1987). A closed-loop model of workload based on a comparison of fuzzy and crisp measurement techniques. *Human Factors, 29*, 339–348.

Moray, N., Kruschelnicky, E., Eisen, P., Money, L., & Turksen, I. R. (1988). A fuzzy set analysis of skill, rule, and knowledge-based workload. In *Proceedings of the Human Factors Society 32nd*

Parasuraman, R., Masalonis, A.J., & Hancock, P.A. (2000). Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance. Human Factors, 42 (4), 636-659.

Annual Meeting (pp. 164–167). Santa Monica, CA: Human Factors and Ergonomics Society.

Parasuraman, R. (1985). Detection and identification of abnormalities in chest x-rays: Effects of reader skill, disease prevalence, and reporting standards. In R. E. Eberts & C. G. Eberts (Eds.), *Trends in ergonomics/human factors II* (pp. 59–66). Amsterdam: North-Holland.

Parasuraman, R. (1987). Human-computer monitoring. *Human Factors, 29*, 695–706.

Parasuraman, R., Hancock, P. A., & Olofinboba, O. (1997). Alarm effectiveness in driver-centered collision-warning systems. *Ergonomics, 39*, 390–399.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics, 30*, 286-297

Parasuraman, R., & Wisdom, G. (1985). The use of signal detection theory in research on human-computer interaction. In *Proceedings of the Human Factors Society 29th Annual Meeting* (pp. 33–37). Santa Monica, CA: Human Factors and Ergonomics Society.

Peterson, W. W., Birdsall, T. G., & Fox, W. C. (1954). The theory of signal detectability. *Transactions of the IRE Professional Group on Information Theory, PGIT-4*, 171–212.

Phelps, C. E., & Hutson, A. (1995). Estimating diagnostic test accuracy using a "fuzzy gold standard." *Medical Decision Making, 15*, 44–57.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Vol. 1.* Cambridge: MIT Press.

Sheridan, T. B., & Ferrell, W. (1974). *Man-machine systems.* Cambridge: MIT Press.

Sorkin, R. D., Kantowitz, B. H., & Kantowitz, S. C. (1988). Likelihood alarm displays. *Human Factors, 30*, 445–459.

Sorkin, R. D., & Woods, D. (1985). Systems with human monitors: A signal detection analysis. *Human-Computer Interaction 1*, 49–75.

Swets, J. A. (1984). Mathematical models of attention. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 183–242). San Diego, CA: Academic.

Swets, J. A. (1996). *Signal detection theory and ROC analysis in psychology and diagnostics.* Mahwah, NJ: Erlbaum.

Swets, J. A., & Pickett, R. M. (1982). *Evaluation of diagnostic systems: Methods from signal detection theory.* New York: Academic.

Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d′ and η as psychophysical measures. *Journal of the Acoustical Society of America, 30*, 922–928.

Tanner, W. P., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review, 61*, 401–409.

Tsoukalas, L. H., & Uhrig, R. E. (1997). *Fuzzy and neural approaches in engineering.* New York: Wiley.

Zadeh, L. A. (1965). Fuzzy sets. *Information and Control, 8*, 338–353.

Raja Parasuraman is a professor of psychology and director of the Cognitive Science Laboratory at the Catholic University of America, Washington, D.C. He received a Ph.D. in psychology from the University of Aston in 1976.

Anthony J. Masalonis is a senior systems development engineer at Mitre Corporation, McLean, VA. He received his Ph.D. in applied experimental psychology from the Catholic University of America in 2000.

P. A. Hancock is a professor of psychology at the University of Central Florida, Orlando, FL. He received his Ph.D. in motor performance from the University of Illinois at Urbana-Champaign in 1983.

Date received: November 2, 1999
Date accepted: March 18, 2000