Theory and Evaluative Reviews

Poisson signal-detection theory: Link between threshold models and the Gaussian assumption

CHRISTIAN KAERNBACH

Laboratoire d'Audiologie Expérimentale, Bordeaux, France

The Gaussian model of signal detection cannot fit asymmetric data as long as the variances of the distributions are kept equal. It is therefore common practice to assume unequal variances in order to fit these data. But this assumption leads to the well-known crossover problem. The present paper provides new arguments for the abandonment of the Gaussian model with unequal variances. In its stead, this paper reevaluates multiple-parallel-threshold models. In particular, the Poisson model turns out to be very useful: it can handle data with any degree of asymmetry, giving a reasonable interpretation of the two parameters of the receiver-operating characteristic. The three-state-threshold model (Krantz, 1969) is given a new interpretation in light of the Poisson model. The slope of Poisson double-probability plots turns out to be much closer to unity than is predicted by the Gaussian approximation.

Classical signal-detection theory (SDT;¹ Falmagne, 1985; Green & Swets, 1974) was left with many open questions concerning the origin of the commonly observed asymmetry of the receiver-operating characteristic (ROC). Although experimental data are commonly described using the Gaussian model with unequal variances, this model should be rejected as it leads to conflict with the general principles of SDT (cf. Green & Swets, 1974). This paper provides new arguments against the unequal-variance assumption and reevaluates the Poisson model of SDT (cf. Egan, 1975), which gives a theoretically consistent way to interpret data of any degree of asymmetry.

The General Principles section below briefly reviews the general principles of SDT that should be respected by all SDT models. This is done by means of a novel concept, the "set of achievable policies." This concept allows the rejection of nonmonotonic ROC curves from purely geometrical arguments. The Classical SDT Models section reviews Gaussian and threshold models. The Parallel-Threshold Models section discusses SDT models based on parallel-threshold processes. In particular, the Poisson model can give a satisfactory explanation of the observed asymmetry. The asymmetry is related to the mean

of the noise distribution, indicating the specificity of the sensory process. The parallel-threshold models connect the Gaussian model with the threshold models, giving both a new interpretation.

GENERAL PRINCIPLES

SDT uses yes-no tasks² to study decision processes in perception. The stimuli are randomly presented with or without signal. The two possible stimuli are labeled "noise" (index n) and "signal + noise" (hereafter "signal" for short, index s). The subject is asked whether there was a signal or not. The answer can be "yes" or "no." With two stimuli and two responses, there are four possible outcomes for each trial: hit, miss, false alarm, and correct rejection. In a series of such trials, one can estimate the subject's conditional probabilities p_n to answer "yes" on "noise" (false-alarm rate) and p_s to answer "yes" on "signal" (hit rate). Thus, the subject's behavior is described by a point (p_n, p_s) in the unit square (see, e.g., point A in Figure 1). To provoke a variety of behaviors, the outcomes of each trial can be rewarded or punished by small amounts of money. For instance, a strict penalty for a miss would cause a less critical policy, expressed in increasing values for both p_s and p_n (point B in Figure 1).

With one and the same signal/noise alternative, the subject has various policies at his/her disposal. The "set of achievable policies" (SOAP) forms an area in the unit square (see Figure 1). This area has some interesting features:

This paper is based partly on a doctoral thesis done at the University of Bonn (Kaernbach, 1988) under the supervision of D. B. Linke. I would like to thank R. Mausfeld for help with literature research and D. M. Green and two anonymous reviewers for comments on earlier versions of the manuscript. Correspondence should be addressed to Christian Kaernbach, Institut für Neuroinformatik, Ruhr-Universität Bochum, Universitätsstrasse 150, Postfach 102148, 4630 Bochum 1, Germany.

- The edge points (0, 0) (say always "no") and (1, 1) (say always "yes") as well as any linear mixture (p, p) of them on the diagonal belong to the SOAP in any case. In general, if point A and point B belong to the SOAP, all points on the line AB belong to it as well. The subject needs only to mix his/her policies. This feature is called convexity (see Figure 1).
- If $A = (p_n, p_s)$ belongs to the SOAP, then $A' = (1-p_n, 1-p_s)$ belongs to it as well. A' represents the same behavior as A, but with "yes" and "no" interchanged. This is called central symmetry.
- With increasing signal strength, the SOAP grows from the diagonal toward the full square, thereby reaching the edge points (0, 1) and (1, 0) of perfect correct or incorrect performance.

The major advantage of the novel SOAP concept is that it derives important features as the convexity from purely geometrical arguments. Convex means curved outward and is used to describe areas where a line segment will lie entirely in the area if its end points lie in the area. That is to say, the area has neither holes nor dents. Because the subject can mix his/her policies (i.e., adopt any policy lying on the line between two other possible policies), the SOAP is necessarily convex. Theories producing a non-convex SOAP will not really describe the SOAP, but something visually similar. Even if they fit experimental data quite well, they can at the most summarize the data empirically but not explain the SOAP. The real SOAP is convex. The central symmetry of the SOAP area means that turning the figure 180° will leave the SOAP the same. This illustrates clearly that policies near (1, 0) (lower right corner) are as hard to achieve as policies near (0, 1) (upper left corner). It is as hard to lie as to tell the truth. Only guessing (i.e., the diagonal) is simple. With increasing signal strength, the area of the SOAP increases from zero

Figure 1. Results of yes-no tasks in the unit square (hit rate, p_s , vs. false-alarm rate, p_n). B represents a less critical response strategy than A. The set of achievable policies (SOAP) is convex and centrally symmetric.

to unity. The area of the SOAP thus provides a suitable measure of the detectability of the signal.

Psychophysics commonly views as most important the upper border of the SOAP, the so-called "receiver-operating characteristic" (ROC). Along the ROC curve, the subject is motivated positively and performing at the limit of his/her ability; the behavior varies only in terms of his/her strategy. The area under the ROC (including the lower right corner) can be shown to be equal to the percentage correct in a two-interval forced-choice task (Green & Swets, 1974). It is therefore often used as a measure of the detectability of the signal. This could mislead one to the assumption that policies near (1, 0) (lower right corner) could easily be adopted for any signal strength. Furthermore, the simple-to-understand, purely geometrical argument "convexity of the SOAP" gets more complicated if expressed in the ROC terminology. It then reads "the slope of the ROC (dp_s/dp_n) monotonically decreasing as a function of p_n " and is commonly derived under the additional assumption of a decision axis that is monotonically related to the a posteriori likelihood ratio (cf. Green & Swets, 1974). The obvious advantages of the novel SOAP concept lead the author to recommend it for teaching purposes. However, to adhere to common usage, ROC terminology will be used henceforth in this paper.

Most SDT models assume a decision axis (for an exception, see Luce, 1959). The stimulus is assumed to evoke an event on this axis, and the subject answers with "yes" if this event was higher than a certain adjustable criterion. The models give probability density distributions for signal and noise. The ROC follows, then, from integration. The ratio of the probability densities for signal and noise is equal to the slope of the ROC, and should therefore increase monotonically. This restricts the set of possible distributions. The ROC can be derived from the distributions, but the distributions cannot be derived directly from the ROC. It is for this reason that up to now much uncertainty has existed about the nature of the underlying distributions.

CLASSICAL SDT MODELS

The Gaussian Model

The best known and most used SDT model is the Gaussian model (Tanner & Swets, 1954). The neural activity for signal and noise is assumed to be distributed Gaussian. Laming (1973) studied the derivation of the Gaussian model from the central-limit theorem and showed that the variances of the distributions for signal and noise should be equal. The equal-variance condition yields symmetric ROC curves with respect to the counterdiagonal. The upper left panel of Figure 2 shows distributions for noise (dotted line) and two signals, and the resulting ROC curves for the Gaussian model with equal variances. Instead of being drawn in the unit square, ROC curves can be given on "double-probability" (DP) paper. Here both coordinates (p_n, p_s) are transformed following the inverse of the cumulative of the Gaussian distribution to (z_n, z_s). The

advantage is that Gaussian ROC curves will appear as straight lines with slope unity. The equation for these straight lines is $z_s = \beta \cdot z_n + i_v$ with slope $\beta=1$. The vertical intercept i_v corresponds to the distance $\Delta\mu$ of the means of the two normalized distributions. The DP plot is illustrated on the upper right panel of Figure 2.

Experimental ROC data (hearing—Tanner, Swets, & Green, 1956; Watson, Rilling, & Bourbon, 1964; vision—Nachmias & Steinman, 1963; Swets, Tanner, & Birdsall, 1961; taste—Linker, Moore, & Galanter, 1964) are usually asymmetric with respect to the counterdiagonal, thus conflicting with a Gaussian model with equal variances. A DP plot shows a slope significantly less than 1, with the tendency of decreasing slope for increasing signal strength. This has led to a modification of the Gaussian model: the assumption of different variances for the distributions for signal and noise. The lower panels of Figure 2 illustrate the ROC curves and DP plots resulting from this model. The DP plots of this model will also give straight lines, but with the slope β being equal to the ratio σ_n/σ_s of the standard deviations for noise and signal. The resulting horizontal and vertical intercepts are $\Delta\mu/\sigma_n$ and $\Delta\mu/\sigma_s$, respectively.

The assumption of unequal variances, however, conflicts with the general principles of SDT (see the General Principles section). For unequal variances, the tails of the distributions cross: the density of the distribution with higher variance exceeds the density of the lower variance distribution for both high and low axis values, whereas for medium axis values it is smaller. Thus, the ratio of the probability densities behaves nonmonotonically on the decision axis, resulting in a nonmonotonic slope of the ROC and a SOAP that is no longer convex. The ROC

even crosses the diagonal, implying a hit rate that is less than a false-alarm rate. Any attempt to save the model by applying nonmonotonic transformations to the decision axis veils the fact that the transformed distributions are no longer Gaussian; in that case, the model should, instead, be described in terms of the transformed distributions. The departures from the regular form of a SOAP or ROC occur usually close to point (1,1), and the experimental errors there are large enough to cover these irregularities. Nevertheless, it should be recalled that such a model will at most summarize the experimental data; it will not explain the nature of the real ROC.

Experimental data tend to show decreasing slope β with increasing vertical intercept i_v . The deviation from unity slope $1-\beta$ is approximately proportional to the intercept i_v . Typical pairs (β, i_v) could be (0.9, 0.4) for a low-level signal or (0.8, 0.8) for a higher level signal. The data tend to hold:

$$\Sigma = (1-\beta)/i_v \approx 1/4. \quad (1a)$$

This relation reveals that the two parameters of the ROC are not completely independent. The ideal parameterization should assign an "asymmetry" parameter to the entire signal-detection task and a "level" parameter to each individual signal level.

Equation 1a provides a graphic description of the data. Better known is its interpretation in terms of the unequal-variance model: $1/\Sigma$ can be shown to be equal to the ratio of the mean distance and the difference of the standard deviations (see, e.g., Green & Swets, 1974):

$$1/\Sigma = i_v/(1-\beta) = \Delta\mu/\Delta\sigma \approx 4 \pm 1. \quad (1b)$$

But this interpretation implies an untenable SDT model, and its formulation does not allow the description of symmetric data. The latter would give a denominator of zero. Furthermore, this interpretation does not explain the origin of this relation.

The Threshold Models

The existence of observer thresholds (for a review, see Krantz, 1969) has always been a point of discussion in SDT. An observer threshold separates internal sensory states. Krantz shows that a threshold model has to assume at least three different states to be consistent with the experimental data. The assumption of only two states (non-detection state $\bar{D}$ and detection state D) will raise the question whether state D results from signals only (high-threshold model; Blackwell, 1953, 1963) or perhaps from noise, too (low-threshold model; Luce, 1963a, 1963b). Both variants lead to predictions that conflict with the experimental data.

Krantz assumes a third state: the superdetection state D^* . State D results from signals and, with smaller probability, from noise. State D^* results from signals only. This model thus has a low threshold (between $\bar{D}$ and D) and a high threshold (between D and D^*). This three-state low- and high-threshold model (3-LHT; see Figure 3) is described by three parameters: the conditional probabilities

Figure 2. ROC curves and DP plots of the Gaussian model with equal and unequal variances. The insets show the distributions for noise (dotted line) and the two signals (solid and dashed lines, corresponding to the two ROC curves). Note the nonconvex ROC curves for the unequal-variance case. They even cross the diagonal.

Figure 3. The three-state low- and high-threshold (3-LHT) model. The distributions for noise (f_n) and three signals (f_{s1} , f_{s2} , f_{s3}) show the probabilities for states $\bar{D}$, D , and D^* . The upper panel shows the resulting ROC curves, and the lower panel, the corresponding DP plots. The distributions were chosen following a Poisson model (see Equation 4).

$f_s(D)$, $f_n(D)$, and $f_s(D^*)$ for signal and noise to result in state D or state D^* . By definition of D^* , $f_n(D^*)$ is equal to zero. The probabilities for $\bar{D}$ follow from the fact that the sum must be unity. The 3-LHT model can fit asymmetric data very well (see Krantz, 1969). It does so at the expense of the inflation of parameters. Krantz himself hoped that “perhaps theoretical interpretation of $\bar{D}$, D , D^* will one day yield a rational relation [between the parameters to reduce their number]” (p. 319). The Poisson model (see below) may serve as a justification of this model as a rather good approximation for certain detection tasks, yielding the desired relation and thus restricting the number of parameters.

The lower panel of Figure 3 shows DP plots of the 3-LHT model. Whereas the Gaussian model produces straight lines in the DP plot, the threshold models produce more or less scalloped curves. In view of this difference, it seems surprising that experimental data could up to now not tell which model was true. But it should be noticed that because of the flat slope of the z -transform, except in the middle region, the binomial error bars are sufficiently small only in the center region of the DP plot. The large error bars outside the center region could only be reduced by more data. However, experiments in SDT have enormous problems in getting a sufficient amount of data in a stable manner, inasmuch as signal-detection situations depend strongly on learning, attention, and fatigue.

PARALLEL-THRESHOLD MODELS

The brain is a massively parallel system. This leads to the suggestion of several threshold mechanisms operating in parallel (Quick, 1974). Egan (1975) reviews the possible detection models based on independent threshold processes. He studies, among others, the binomial and the Poisson distribution (for the Poisson model, see also Nachmias, 1972). The binomial model of signal detection assumes m parallel low-threshold processes. The neural activity is then a number k of the range $[0, m]$ and reflects the number of thresholds exceeded. Given the parameters q_n and q_s for the single process, the neural activity k follows a binomial distribution:

$$f_n(k) = \binom{m}{k} q_n^k (1 - q_n)^{m-k}, \quad f_s(k) = \binom{m}{k} q_s^k (1 - q_s)^{m-k}. \quad (2)$$

The resulting ROC curves can show any degree of asymmetry. A left shift (leading to a DP plot with a slope of less than unity) results from low values of q_n . In this situation, neural activities exceeding the threshold, given noise alone, are rare events. On the other hand, a right shift results from high probabilities of exceeding the threshold even for noise alone. This corresponds to a saturation: the nonexcesses become rare events. The asymmetry observed in most experiments is of the former category. The binomial model helps us to understand this trend: The thresholds under consideration are specific for the signal, that is, they are rarely exceeded by noise alone.

The number m of the parallel-threshold processes determines the characteristic of the binomial model. With $m = 1$, it is equivalent to the low-threshold model. With an increasing number of parallel processes, one may not wish to be concerned by its exact quantity. There are two different ways to approach the limit of $m \rightarrow \infty$: On the one hand, one can keep q_n constant. With increasing m , the distributions of k approach Gaussian distributions. q_s will have to come closer and closer to q_n to produce a similar ROC, and thus the variance of f_s approaches the variance of f_n . This is reminiscent of Laming's (1973) central-limit-theorem argument for a Gaussian model with equal variances (see the Gaussian Model section above). On the other hand, one can keep the product $\mu_n = mq_n$ constant. This leads to Poisson distributions for signal and noise. The value of the binomial model lies in the connection that it establishes between the low-threshold model, the Gaussian model with equal variances, and the Poisson model. And it should be kept in mind that the binomial model is more realistic than the Gaussian or the Poisson model, since the number of possible threshold processes is certainly finite, although admittedly high.

The Poisson distribution is appropriate to describe events that are unlikely to occur in each single case but where the number of independent processes is high. The Poisson distribution is determined by only one parameter.

Figure 4. Poisson model: distribution of threshold excesses k for noise (f_n) and two signals (f_{s1} , f_{s2}). The right-hand panel shows the resulting ROC curves: (a) smaller means ($\mu_n = 0.15$, $\mu_{s1} = 0.8$, $\mu_{s2} = 1.6$); (b) higher means ($\mu_n = 1.0$, $\mu_{s1} = 2.7$, $\mu_{s2} = 5.0$).

ter: the mean μ . The mean and variance are the same parameter, that is, $\sigma^2 = \mu$. Its use for neural firings at the threshold is obvious and has often been recommended (McGill & Goldberg, 1968a, 1968b; Zwilowski & Jordan, 1986). The Poisson model of signal detection assumes an infinite number of parallel low-threshold processes for the same signal. The neural activity k again represents the number of thresholds exceeded. k is Poisson distributed. Thus, one has two parameters to fit ROC data: the mean μ_n of the distribution for noise and the mean μ_s for the signal. The distributions will then be as follows:

$$f_n(k) = e^{-\mu_n} \mu_n^k / k!, \quad f_s(k) = e^{-\mu_s} \mu_s^k / k!. \quad (3)$$

Figure 4 shows the resulting ROC curves. Two signal distributions (indices $s1$ and $s2$) were assumed for each panel. For low means of both distributions (Figure 4a), they resemble 3-LHT ROC curves. For higher means (Figure 4b), they begin to resemble the Gaussian ROC curves. Poisson ROC curves are convex for any chosen pairs of parameters. Any degree of asymmetry can be achieved by the appropriate choice of the two parameters μ_n and μ_s . The rarely found saturation (DP slope of more than unity), however, is not represented. The Poisson model shows decreasing DP slope with increasing intercept as is required by the experimental data (see the discussion of Equation 1).

Equation 1b was the translation of the empirically found trend of data (Equation 1a) to the unequal-variance model. It did not explain anything. For the Poisson model, the constancy of Σ is reflected by simply keeping μ_n constant for all signal levels of one signal-detection task. This is not an arbitrary restriction but follows directly from the con-

struction of the model. The tendency of many experimental data to keep $\Sigma = (1-\beta)/i_v$ in the region of 0.25 ± 0.1 ($\Delta\mu/\Delta\sigma = 4 \pm 1$) is reflected by the Poisson model for one specific noise mean: $\mu_n = 0.10 \pm 0.05$, a rather low value leading to nearly triangular ROC curves. A family of ROC curves with such a noise mean (and increasing signal means; see, e.g., Figure 4a) would thus represent many experimental data. Other values of Σ could be reflected by other values of the noise mean. The Poisson model thus fits asymmetric data with the minimum number of parameters possible. The two parameters are naturally divided into the asymmetry parameter μ_n (constant for all different signal levels) and the signal-level parameter μ_s . The Poisson model offers an illuminating interpretation of its parameters. For a given noise mean, the position of the signal mean is related to the area of the SOAP and describes the detectability of the signal. The position of the noise mean, which is related to the asymmetry (the lower the μ_n , the less the symmetry), describes the specificity of the involved threshold processes. Asymmetric ROC curves indicate that the subject has thresholds at his/her disposal which are quite specific to the signal in question. These thresholds will rarely be exceeded by noise stimuli: μ_n remains small. Symmetric ROC curves are due to high values of μ_n , indicating that the subject is forced to use a lot of less specific detection processes.

For low noise means (e.g., the frequent case $1/\Sigma = 4 \pm 1$, i.e., $\mu_n = 0.10 \pm 0.05$), the Poisson model can be approximated by the 3-LHT model (see the Threshold Models section). For $\mu_n = 0.10$, the probability of state $k=2$ and all higher states together is less than .01 and can be approximated to zero. This gives the 3-LHT model a new interpretation (see Figure 5). The state $k=0$ (no threshold excess) of the Poisson model corresponds to the nondetection state $\bar{D}$. The state $k=1$ corresponds to

Figure 5. The 3-LHT model interpreted in terms of the Poisson model. Two parameters suffice to determine the model, namely $a_n = \exp(-\mu_n)$ and $a_s = \exp(-\mu_s)$.

the detection state D. All other states ($k=2$ and higher) are subsumed under the superdetection state D^* , as only signals contribute noticeably to these states. Let $a_s = \exp(-\mu_s)$ and $a_n = \exp(-\mu_n)$ be the probabilities $f(\bar{D})$ for the non-detection state for signals and noise. Then, from Equation 3 for the Poisson model with low noise means ($\mu_n \leq 0.25$), follow the probabilities for the other states:

$$\begin{aligned} f_n(\bar{D}) &\equiv a_n = e^{-\mu_n}, & f_n(D) &\approx 1 - f_n(\bar{D}), & f_n(D^*) &\approx 0, \\ f_s(\bar{D}) &\equiv a_s = e^{-\mu_s}, & f_s(D) &= -a_s \cdot \ln(a_s), \\ f_s(D^*) &= 1 - f_s(\bar{D}) - f_s(D). \end{aligned} \quad (4)$$

Thus, the 3-LHT model is described by two parameters only: a_s and a_n . This relation was used for the demonstration of this model in Figure 3. This figure may thus serve as an illustration of the ROC curves and DP plots of the Poisson model for low-noise means. It demonstrates the DP slopes that can be produced by that model. The Poisson model serves as a justification of the 3-LHT model as an approximation and yields the desired relation between the three original parameters of that model.

A DP plot of Poisson ROC curves (Figure 6, solid lines) shows a linear behavior of the edge points with a slope of less than unity. This is reminiscent of the Gaussian model with unequal variances (Figure 2). But it should be remembered that very different distributions can cause identical ROC curves. Furthermore, only the edge points are nearly exactly linear; the connecting parts appear to be curved towards the diagonal. The deviation from the straight line is, however, most remarkable for the upper right points, where the Poisson ROC curve does not intersect the diagonal. In any case, the approximate linearity of the Poisson ROC DP plots deserves a comparison to

Figure 6. Comparison of the DP plots for the Poisson ROC curves of Figure 4b (solid lines, $\mu_n = 1.0$, $\mu_{s1} = 2.7$, $\mu_{s2} = 5.0$) and the Gaussian model (dashed lines) with the same means and variances ($\sigma_n = \sqrt{\mu_n}$, $\sigma_{s1} = \sqrt{\mu_{s1}}$, $\sigma_{s2} = \sqrt{\mu_{s2}}$). The slope of the Gaussian model is significantly less than that of the Poisson model. Dotted line: straight line fit to the edge points of the Poisson ROC curve.

the Gaussian model with unequal variances. The dashed lines in Figure 6 derive from Gaussian distributions with the same mean and variance as the Poisson distributions that produced the solid lines. They have a significantly lower slope than the corresponding Poisson plots. That is to say, Poisson distributions reach the unity slope and thus symmetric ROC curves for much smaller values of the means than the Gaussian model with the same means and variances. It is surprising that distributions as similar as Poisson and Gaussian distributions lead to such different ROC curves.³

It is quite remarkable that the observed differences between Poisson ROC shapes and Gaussian ROC shapes correspond very closely to those reported by Greig (1990) for chi-squared distributions. Greig presents an approximation formula that predicts the slopes and the intercepts of energy-detection ROC curves from the means and the variances of the underlying chi-squared distributions. He describes DP straight-line fits by means of two parameters: the slope β and the accuracy measure d_A . The latter is $\sqrt{2}$ times the minimum distance between the origin (0,0) and the straight line: $d_A = \sqrt{2} \cdot i_v / \sqrt{1 + \beta^2}$. The measure d_A is fairly well predicted by the Gaussian approximation: $d_A \approx \Delta\mu / \sqrt{[(\sigma_n^2 + \sigma_s^2)/2]}$. The slope β may be approximated by $(\sigma_n/\sigma_s)^{0.25}$. As a matter of fact, I verified that the same approximation works quite well for the Poisson model, too. Greig concluded that the differences from the Gaussian model should be due to the skewness of the chi-squared distribution. The Gaussian approximation of these distributions drops off too fast. Poisson distributions are more similar to Gaussian distributions than are chi-squared distributions. Skewness could nevertheless be the explanation of the observed phenomenon.

Skewed distributions can sometimes be transformed to symmetric distributions and vice versa by means of a non-linear transformation. $z^{-\ln(z)-1}dz$ will become $e^{-x^2}dx$ by means of the transformation $z = e^x$. Could there be found a single transformation that contributes to the observed differences between Gaussian and Poisson ROC shapes? Imagine three Poisson distributions giving rise to three ROC curves (one for each pair). It should then be possible to find three corresponding Gaussian distributions of unequal variance that would give rise to quite similar ROC curves. Unfortunately, this does not work out. The six parameters of the three Poisson DP plots will nearly always contradict the four degrees of freedom of two Gaussian distributions (the first distribution can be normalized). Important progress in SDT would be made if one could name and quantify those features of the distributions—mean and variance are obviously not appropriate⁴—that are essential to the ROC. That could give us a better understanding of the asymmetry of the ROC.

The exact transformation of the parameters of the DP fit to the means of the Poisson distribution is missing. Table 1 gives the means belonging to a large set of *fit* parameters. Slope β and vertical intercept i_v correspond to straight-line fits to the edge points of the DP plots resulting from a Poisson model with the means as indicated.

Table 1

Poisson Means μ_n and μ_s for Noise and Signal, Respectively, Given the Vertical Intercept i_v and the Slope β of DP Straight-Line Fits

β	μ_n	μ_s	β	μ_n	μ_s
$i_v = .1$			$i_v = .8$		
.99	1.953	2.095	.96	10.14	12.93
.98	.175	.213	.95	5.857	8.019
.97	.018	.027	.94	3.546	5.269
$i_v = .2$			.93	2.307	3.730
.99	10.80	11.47	.92	1.595	2.806
.98	1.894	2.179	.91	1.061	2.075
.97	.515	.660	.90	.726	1.586
.96	.169	.248	.88	.368	1.012
.95	.059	.102	.86	.180	.657
.94	.018	.039	.84	.091	.452
.93	.004	.012	.82	.047	.330
$i_v = .3$			.80	.019	.211
.98	5.364	6.092	.78	.009	.166
.97	1.839	2.269	$i_v = 1.0$		
.96	.732	1.003	.95	9.946	13.46
.95	.327	.506	.94	6.296	9.162
.94	.167	.294	.93	4.156	6.541
.93	.071	.150	.92	2.867	4.900
.92	.031	.081	.91	2.010	3.757
.91	.015	.048	.90	1.470	3.003
.90	.005	.022	.88	.774	1.952
$i_v = .4$			.86	.446	1.395
.98	10.57	11.93	.84	.246	1.002
.97	3.956	4.806	.82	.140	.755
.96	1.771	2.347	.80	.071	.551
.95	.906	1.324	.78	.039	.431
.94	.494	.805	.76	.018	.324
.93	.266	.494	.74	.010	.264
.92	.150	.321	$i_v = 1.2$		
.91	.073	.190	.94	9.558	13.78
.90	.043	.132	.92	4.552	7.617
.88	.012	.058	.90	2.399	4.748
$i_v = .5$			.88	1.347	3.210
.97	6.804	8.200	.86	.783	2.294
.96	3.212	4.192	.84	.492	1.768
.95	1.721	2.451	.82	.294	1.356
.94	.985	1.547	.80	.169	1.055
.93	.571	1.004	.78	.095	.822
.92	.339	.676	.76	.060	.696
.91	.210	.479	.74	.030	.540
.90	.134	.351	.72	.018	.477
.88	.049	.181	.70	.009	.381
.86	.018	.101	$i_v = 1.4$		
.84	.005	.048	.92	6.682	11.00
$i_v = .6$			.90	3.580	6.914
.97	10.38	12.44	.88	2.095	4.791
.96	5.085	6.565	.86	1.276	3.510
.95	2.802	3.925	.84	.768	2.624
.94	1.646	2.524	.82	.500	2.123
.93	1.012	1.714	.80	.313	1.720
.92	.667	1.247	.78	.181	1.332
.91	.428	.902	.76	.117	1.135
.90	.286	.681	.74	.071	.960
.88	.119	.382	.72	.041	.814
.86	.051	.232	.70	.020	.670
.84	.018	.131	.68	.014	.608
.82	.006	.076	.66	.006	.500

Table 1 (Continued)

β	μ_n	μ_s
$i_v = 1.6$		
.92	8.708	14.34
.90	5.051	9.560
.88	3.029	6.712
.86	1.878	4.956
.84	1.189	3.800
.82	.734	2.944
.80	.500	2.498
.78	.313	2.014
.76	.197	1.675
.74	.133	1.472
.72	.074	1.222
.70	.051	1.116
.66	.018	.882
.62	.004	.623
$i_v = 1.8$		
.90	6.780	12.63
.88	4.017	8.780
.86	2.570	6.600
.84	1.622	5.039
.82	1.075	4.051
.80	.725	3.355
.78	.500	2.873
.76	.313	2.385
.74	.204	2.034
.72	.142	1.823
.70	.081	1.556
.66	.031	1.259
.62	.011	1.029
$i_v = 2.0$		
.90	8.250	15.42
.88	5.330	11.39
.86	3.380	8.490
.84	2.231	6.640
.82	1.490	5.338
.80	.985	4.373
.78	.700	3.750
.76	.479	3.221
.74	.313	2.776
.72	.204	2.407
.70	.144	2.189
.66	.061	1.787
.62	.018	1.440
.58	.006	1.205
$i_v = 2.4$		
.90	11.70	21.94
.88	7.766	16.52
.86	5.364	13.07
.84	3.524	10.08
.82	2.430	8.253
.80	1.680	6.864
.78	1.207	5.912
.76	.832	5.075
.74	.582	4.449
.72	.414	3.981
.70	.299	3.626
.66	.136	2.977
.62	.060	2.550
.58	.018	2.017
.54	.006	1.852

For instance, in Figure 6, the fit to the edge points of the Poisson ROC curves is indicated by dotted lines. The parameters of the fit for the upper curve are $i_v = 2.24$ and $\beta = 0.78$. Table 1 contains $i_v = 2.0$ and $i_v = 2.4$. Linear interpolation yields the original parameters $\mu_n = 1$ and $\mu_s = 5$ with less than 1% error. For small means, the resulting DP plots deviate markedly from the straight line. The *fit* parameters obtained might then depend strongly on the exact position of the data points.

Let us finally consider the steps you should take to fit your data to the Poisson model. First of all, you should plot them on DP paper and fit a straight line to them. If the data are compatible with a straight line of slope unity, assume the Gaussian model with equal variances. If the slope is significantly less than unity, look up the corresponding means in Table 1. If μ_n is smaller than 0.3 (highly asymmetric data), you should replot your data in the unit square and fit a two-parameter 3-LHT model as described in the Parallel-Threshold Models section (a_n should come out to be 0.75 or higher). If μ_n is greater than 1, you can rely on the means. In between, it would be best to replot the data in the unit square and directly fit a Poisson model. This might be troublesome because the Poisson DP plots have several dents that will fit the data better or less well depending on their exact position. Thus, a direct fit will have to consist of random searches for well-fitting pairs of parameters alternated with gradual descents. Up to now, insufficient experience exists to estimate the error undertaken by simply relying on the mean values of Table 1. However, these mean values will at least give an approximation of the detection process under consideration.

CONCLUSIONS

A brief review of classical SDT showed an open problem in the treatment of asymmetric ROC data. The well-known Gaussian model with equal variances fails to represent asymmetric data. The application of unequal variances will fit these data, but this leads to ROC curves with nonmonotonic slopes that cross the diagonal, which conflicts with the general principles of SDT. Threshold models can fit asymmetric data equally well, but they need at least three parameters to represent the data. The relation between these parameters has not been determined up to now.

There do exist one-parameter models that produce asymmetric ROC curves and do not have the crossover problem (cf. the exponential model, Green & Swets, 1974). They may be appropriate to fit certain data with a certain degree of asymmetry. Two parameters are needed to fit data of a varying degree of asymmetry. Energy-detection models based on chi-squared distributions (Green & Swets, 1974) have two parameters and produce convex ROC curves. They show a tendency toward decreasing slope for increasing intercepts as required by experimental data

(Greig, 1990). Their major problem is that they predict incremental thresholds different from *Weber's law*. With regard to the asymmetries observed in experiments, this paper reevaluates the Poisson model (cf. Egan, 1975) as an alternative.

The Poisson model can handle data of any degree of asymmetry. It does so with two parameters, one of which is kept constant over the entire signal-detection task, while the other changes with the signal level. Its freedom of curve fitting is thus minimized. The Poisson model produces ROC curves with a monotonically decreasing slope as required. Moreover, it gives helpful insight into the meaning of the two parameters. Asymmetry is interpreted as indicating that the subject has thresholds at his/her disposal that are rather specific to the detection task and will thus rarely be exceeded by noise stimuli. For low noise means (highly asymmetric data), the three-state low- and high-threshold (3-LHT) model of Krantz (1969) is a valid approximation. The Poisson model reduces the number of parameters of the 3-LHT model.

Since it has a nondetection state, the Poisson model is a threshold model. The importance of this may, however, be negligible, and the Poisson model can converge toward the Gaussian model. A critical test of the Poisson model is not easy to find, as it is consistent with all known trends of data. The Poisson model is a neural quantum model, concerned not with possible measurable effects of certainly existing neural quantization, but with the possibility of utilizing it for a decision theory. The Poisson distribution is the appropriate distribution for such rare events as neural firings near the threshold. Its use in SDT is therefore obvious, and the interpretation of ROC data in Poisson terms may provide interesting information about the number of involved processes and their specificity. Thus, the Poisson model may help us to understand the implementation of signal-detection mechanisms in the brain.

REFERENCES

- BLACKWELL, H. R. (1953). Psychological thresholds: Experimental studies of methods of measurement. *Bulletin of the Engineering Research Institute of the University of Michigan*, No. 36.
- BLACKWELL, H. R. (1963). Neural theories of simple visual discriminations. *Journal of the Optical Society of America*, **53**, 129-160.
- EGAN, J. P. (1975). *Signal detection theory and ROC analysis*. New York: Academic Press.
- FALMAGNE, J.-C. (1985). *Elements of psychophysical theory*. Oxford: Oxford University Press.
- GREEN, D. M., & SWETS, J. A. (1974). *Signal detection theory and psychophysics*. New York: Krieger.
- GREIG, G. L. (1990). On the shape of energy-detection ROC curves. *Perception & Psychophysics*, **48**, 77-81.
- KAERNBACH, C. (1988). *Entscheidungsverhalten an der Wahrnehmungsschwelle: Theoretische Überlegungen und ein Experiment zur klassischen Ja-Nein-Aufgabe*. Medical doctor thesis, Bonn University, Bonn, West Germany.
- KRANTZ, D. H. (1969). Threshold theories of signal detection. *Psychological Review*, **76**, 308-324.
- LAMING, D. R. J. (1973). *Mathematical psychology*. New York: Academic Press.
- LINKER, E., MOORE, M. E., & GALANTER, E. (1964). Taste thresholds, detection models, and disparate results. *Journal of Experimental Psychology*, **67**, 59-66.
- LUCE, R. D. (1959). *Individual choice behavior*. New York: Wiley.
- LUCE, R. D. (1963a). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 103-189). New York: Wiley.
- LUCE, R. D. (1963b). A threshold theory for simple detection experiments. *Psychological Review*, **70**, 61-79.
- MCGILL, W. J., & GOLDBERG, J. P. (1968a). Pure-tone intensity discrimination and energy detection. *Journal of the Acoustical Society of America*, **44**, 576-581.
- MCGILL, W. J., & GOLDBERG, J. P. (1968b). A study of the near-miss involving Weber's law and pure-tone intensity discrimination. *Perception & Psychophysics*, **4**, 105-109.
- NACHMIAS, J. (1972). Signal detection theory and its application to problems in vision. In D. Jameson & L. M. Hurvich (Eds.), *Handbook of sensory physiology* (Vol. VII/4, pp. 56-77). New York: Springer-Verlag.
- NACHMIAS, J., & STEINMAN, R. M. (1963). Study on absolute visual detection by the rating-scale method. *Journal of the Optical Society of America*, **53**, 1206-1213.
- QUICK, R. F., JR. (1974). A vector-magnitude model for contrast detection. *Kybernetik*, **16**, 65-67.
- SWETS, J. A., TANNER, W. P., & BIRDSALL, T. G. (1961). Decision processes in perception. *Psychological Review*, **68**, 301-340.
- TANNER, W. P., & SWETS, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, **61**, 401-409.
- TANNER, W. P., SWETS, J. A., & GREEN, D. M. (1956). *Some general properties of the hearing mechanism* (Tech. Rep. No. 30). Ann Arbor: University of Michigan, Electronic Defense Group.
- WATSON, C. S., RILLING, M. E., & BOURBON, W. T. (1964). Receiver operating characteristics determined by a mechanical analog to the rating scale. *Journal of the Acoustical Society of America*, **36**, 283-288.
- ZWISLOCKI, J. J., & JORDAN, H. N. (1986). On the relation of intensity jnd's to loudness and neural noise. *Journal of the Acoustical Society of America*, **79**, 772-780.

NOTES

1. The term SDT often implies the Gaussian assumption. Green and Swets (1974) use it in a more general sense. Threshold theories are not outside this class. This general interpretation of SDT will be used here, whereas the Gaussian assumption will be referred to as the Gaussian model.
2. Two other tasks are of great importance to SDT, namely forced-choice (FC) tasks and confidence-rating (CR) tasks. The results of FC tasks can be predicted from yes-no ROC curves (Green & Swets, 1974). For a detailed discussion of CR tasks and their relation to yes-no tasks, see Krantz (1969). Krantz demonstrates that the curvilinear shape of CR data does not rule out the rectilinear shape predicted for yes-no tasks by threshold theories.
3. Nachmias (1972) compared Poisson models and Gaussian models with unequal variance. His Figure 7 shows only the edge points of Poisson DP plots, which could mislead us to assume a perfectly linear characteristic. He does not plot the lines corresponding to the Gaussian approximation into the same figure, but he assumes that for moderately large values of μ the slope and the intercepts of the Poisson model can be predicted by the Gaussian model. This does not hold true for values as high as $\mu_n = 81$ ($\sigma_n = 9$) and $\mu_s = 100$ ($\sigma_s = 10$). The Gaussian model would predict a slope of 0.9. The Poisson model shows a slope of 0.966. This difference is the more surprising as the Poisson distribution for $\mu = 81$ is fairly well approximated by the corresponding Gaussian distribution.
4. The Gaussian model with unequal variances relates the decreasing slope of DP fits to increasing variances. This cannot hold for distributions other than Gaussian. For instance, a model based on gamma functions will instead give increasing slope with increasing signal, notwithstanding the increasing variance of the signal distribution.