

The Morton-Massaro Law of Information Integration: Implications for Models of Perception

Javier R. Movellan
University of California, San Diego

James L. McClelland
Carnegie Mellon University and
Center for the Neural Basis of Cognition

Information integration may be studied by analyzing the effect of 2 or more sources (e.g., auditory and visual) on participants' responses. Experiments show that ratios of response probabilities often factorize into components selectively influenced by only 1 source (e.g., 1 component affected by the acoustic source and another 1 affected by the visual source). This is called the Morton-Massaro law (MML). This article identifies conditions in which the law is optimal and notes that it reflects an implicit assumption about the statistics of the environment. Adherence to the MML can be used to assess whether the assumption is being made, and analyses of natural stimuli can be used to determine whether the assumption is reasonable. Feed-forward and interactive models subject to a channel separability constraint are consistent with the law.

Many experiments and demonstrations document that the perception of a wide range of stimuli is affected by multiple sources of information. Figure 1 shows a classic example, Ebbinghaus's circles, in which the perceived size of an object is affected by its actual size and by the context in which it appears. Extensive empirical research has been performed in a variety of fields, including cognitive psychology, neuroscience, behavioral ethology, and machine perception, to understand how information is integrated in perception (Clark & Yuille, 1990b; Inui & McClelland, 1996; Partan & Marler, 1999; Stein & Meredith, 1993).

We focus on a specific characteristic of the results that has often been observed in this research. The characteristic pattern arises in experiments examining the effects of variation of two or more different sources of information on perceptual identification responses. For example, one might investigate the effects of visual and auditory sources of information on the identification of a spoken syllable. The pattern is described in detail later, but its essence can be stated in simple terms: The probabilities of identification responses (e.g., saying that a given item is /ba/) given to combinations of inputs from the different sources can be accounted for with a set of stimulus–response factors, in which each factor

represents the support provided by one of the sources to one of the possible identification responses. The pattern indicates that a given input from one source makes the same contribution in support of each identification response, regardless of the input from the other sources.

Two important and related theories have been developed that predict this empirical pattern: one by Morton (1969) and one by Massaro et al. (e.g., Oden & Massaro, 1978). Morton (1969) reviewed a wide range of experiments and proposed a model, the logogen model, that described reasonably well the effects of stimulus and context information on word identification responses. Massaro et al. (Massaro, 1987a, 1989b; Oden & Massaro, 1978) developed a powerful experimental paradigm for the study of information integration in which several levels of each independent variable or information source are factorially combined (see Figure 2 for an example) and the dependent variables are the probabilities of a given set of response alternatives. Massaro et al. (Oden & Massaro, 1978) proposed a model of information integration, the fuzzy logical model of perception (FLMP), to explain the results obtained using this experimental paradigm.

An important feature of the logogen model and the FLMP is that they use autonomous modules that selectively process each of the different information sources and are not influenced by the other sources or by later processing stages (Figure 3). These modules map low-level inputs into quantities that represent the strength of support for high-level perceptual hypotheses (e.g., whether the stimulus is a “G” or a “Q”). As we see later in more detail, the fact that these models are modular leads them to predict that the strength of support for each alternative will factorize as described previously.

Both Morton and Massaro pointed out many cases in which the effects of different information sources can be factored in this way. Morton (1969) showed that his model provided a good fit to a number of experiments examining the effects of context and stimulus information on word identification. Massaro et al. showed

Javier R. Movellan, Department of Cognitive Science and Institute for Neural Computation, University of California, San Diego; James L. McClelland, Department of Psychology, Carnegie Mellon University and Center for the Neural Basis of Cognition.

The research reported here was supported by National Institutes of Health Grants MH00385 and MH47566.

We thank John Anderson, Arthur Jacobs, Dominic Massaro, Ken Paap, and Richard Zemel for discussions and comments on earlier drafts of this article. Javier R. Movellan thanks the Spanish psychologist Angel Riviere in memoriam.

Correspondence concerning this article should be addressed to Javier R. Movellan, Department of Cognitive Science, University of California, San Diego, La Jolla, California 92093-0515. Electronic mail may be sent to movellan@cogsci.ucsd.edu.




Figure 1. The Ebbinghaus circles are a classic example of the effects of context on perception. The perception of the size of the inner circle is affected by the context.

that, as predicted by the FLMP, factorability provides good approximations for response probabilities obtained in a remarkable range of experiments in domains such as word and letter perception, object identification, depth perception, audiovisual speech recognition, memory retrieval, and recognition of emotions. See Massaro's (1989b) Chapter 6 for a review.

The range of cases in which effects of different sources factor is extensive, and we refer to this fact as the Morton-Massaro law. However, it is important to note that the extent of applicability of this law is controversial. Many investigators claimed to find violations (Bülthoff & Yuille, 1996; Samuel, 1981, 1996). Some violations can be disputed (see, e.g., Massaro, 1996), leaving the exact extent of the applicability of the law in doubt. It is not our intention to take a stand on these particular cases, and we do not use the term "law" to endorse the view (Massaro, 1989b, ch. 4) that factorability is a universal principle of human information processing that spans all domains in which information integration is necessary. Instead, we use the term to suggest that the empirical pattern of factorability is worth taking note of, in view of the wide range of cases in which it does hold, and also that the implications of factorability for theories of perception must be seriously considered. Thus, we treat the Morton-Massaro law as a very useful




Figure 2. Stimuli illustrating Massaro's factorial experimental paradigm. On each trial participants see 1 of the 49 stimuli for 400 ms and are asked to categorize it as G or Q. The two sources of information are the angle of the straight line and the openness of the oval gap. From *Experimental Psychology: An Information Processing Approach* (p. 329), by D. W. Massaro, 1989, San Diego, CA: Harcourt Brace Jovanovich. Copyright 1989 by Harcourt Brace Jovanovich. Reprinted with permission.




Figure 3. Schematic representation of the fuzzy logical model of perception. The information sources are represented with the uppercase letter X. Each source of information is selectively processed by a different evaluation module. The output of each module represents the degree of support of an information source for the different response alternatives. These outputs are integrated at a later stage. The response stage operates on the output of the integration stage to produce an external response such as discrete decision or a continuous rating. Note that propagation of information may be cascaded, so that information available to the response stage may change as a function of processing time.

relationship whose confirmation or refutation in particular experiments may inform us about the underlying structure of the perceptual system. Indeed, we present a case of a violation of the law later, and we use it to illustrate how such a violation may reflect important characteristics of the underlying perceptual system that produces it. This way of using factorability of response probability ratios is similar to the way Sternberg (1969) used additivity of reaction times to understand the organization of underlying cognitive processes.

The Morton-Massaro law raises the important issue of its potential incompatibility with interactive models of perception (i.e., models in which all the processing units may be coupled through feedback and lateral connections). According to the FLMP, perceptual processes are feed-forward and modular. They can be decomposed into three stages (see Figure 3). First, there is an evaluation stage in which information sources are evaluated in a modular manner. Next, there is an integration stage in which the outputs of the evaluation modules are integrated. Finally, there is a decision stage in which a response is chosen based on the output of the integration stage. The outputs of the different stages may overlap in time and may be organized hierarchically, but in all such cases information processing is modularized and proceeds in a strictly feed-forward manner. (Note that when the FLMP is implemented as a process model, propagation is thought of as continuous, so that information made available to the integration and to the response stages can change gradually as a function of processing time.)

In neuroscience the feed-forward characterization reflects traditional views about the early stages of visual processing. For example, classic studies on the behavior of single neurons at different levels of the visual pathway showed that, in anesthetized cats and monkeys, neurons in the primary visual cortex have small receptive fields and selectively respond to simple visual stimuli (e.g., gratings of a particular spatial frequency and orientation), whereas neurons in more anterior areas have larger receptive fields and respond to more complex stimuli (Hubel & Weisel, 1962, 1968). The notion that neurons in primary cortex have small local receptive fields is consistent with feed-forward information-

processing models. If feedback and lateral connections had a functional influence on simple neurons in primary visual cortex, their receptive fields should be larger, reflecting the influence of neighboring neurons and of anterior neurons. Feed-forward models, like the logogen model or the FLMP, provide a nice information-processing metaphor for this picture of the neurophysiology.

An alternative approach views perception as arising from a bidirectional propagation of information. In this approach, a perceptual processing system may be subdivided into levels, but there are bidirectional influences between and within levels, so that the outcome of processing at every level is potentially subject to influence from higher and lower levels as well as from other units at the same level. McClelland and Rumelhart's (1981) interactive activation model of word perception is an early example of a model of this type. In this model, a letter unit may influence a word-level unit through feed-forward connections and, conversely, the word-level unit can return the influence back to the letter level through feedback connections (Figure 4). The model was implemented as a connectionist network in which features, letters, and words were represented as local units that interacted according to simple processing principles. Top-down and bottom-up combination of information was achieved through bidirectional connections between the feature and letter units and between the letter and word units, as Figure 4 illustrates. These bidirectional connections

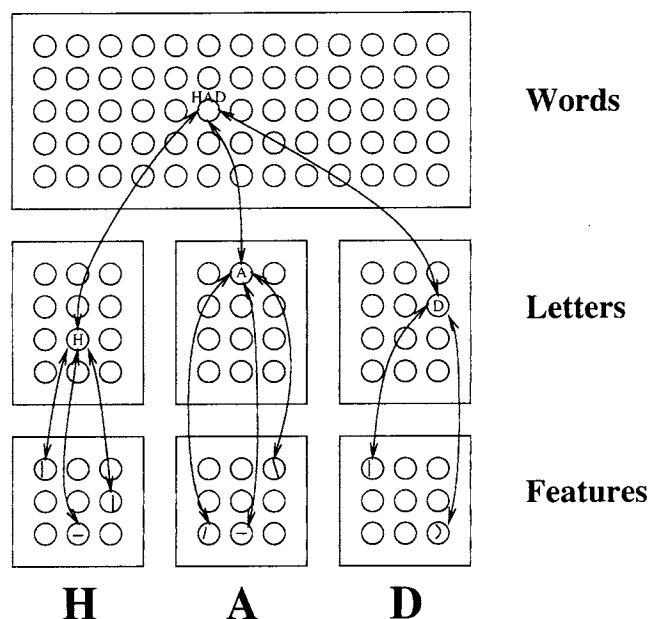


Figure 4. The interactive activation architecture (McClelland & Rumelhart, 1981). The version shown is applicable to three-letter words such as those used in the experiment described later in the text. There is a unit for every three-letter word in English, a unit for each letter in each position, and a set of feature units for each position. Not all units are shown. Bidirectional, excitatory connections for the word HAD are illustrated. Between-level connections are excitatory and symmetric. Within each pool of units, there are symmetric, mutually inhibitory connections between incompatible units (not shown). In some versions of the interactive activation architecture, between-level inhibitory connections have been used, but these are not included in this version.

allowed the network to exhibit a wide variety of word perception phenomena, leading McClelland and Rumelhart (1981) to suggest that such connectionist models might serve as a general paradigm for perception. McClelland and Rumelhart and other authors applied these ideas to a range of topics in perception, attention, memory, word reading, language production, and conceptual knowledge representation (e.g., Dell, 1986; Hinton & Shallice, 1991; McClelland, 1981; McClelland & Elman, 1986; McClelland & Rumelhart, 1986; Phaf, van der Heijden, & Hudson, 1990; Plaut, McClelland, Seidenberg, & Patterson, 1996; Rumelhart, Smolensky, McClelland, & Hinton, 1986).

Interactive concepts also have a long history in neuroscience. Ramon y Cajal (1892) emphasized the existence of feedback connections in his early studies of the brain, and Lorente de No (1922), a pioneer of theoretical ideas about neural computation, proposed "reciprocidad" as a design principle of the brain. Anatomical studies confirmed this principle by showing that, when there is a projection from one brain region to another, there is almost always a reciprocal return projection (Felleman & Essen, 1991). A number of neuroscientists (Peon, 1961; Sperry, 1969; Szentagothai & Arbib, 1974) cautioned against modular and unidirectional views of information processing. Mumford (1992) pointed out that the majority of cortical cells have interarea projections as opposed to exclusively intra-area projections and that this fact runs against modular computing metaphors. More recently, physiological studies on alert-behaving animals showed that contextual attributes far beyond the classical receptive field modulate the response of single neurons in primary visual cortex (Allman, Miezin, & McGuiness, 1985; Lamme, 1995; Lee, Mumford, Romero, & Lamme, 1998). These contextual effects are interpreted as being due to feedback connections from downstream areas and collateral interactions among neurons within the same cortical region. There are now direct demonstrations that contextual influences seen in somatosensory thalamic nuclei arise from feedback connections from somatosensory cortex (Ghazanfar & Nicolelis, 1997). Modular feed-forward models offer little help to understand this picture of the neurophysiology.

The interactive character of connectionist models with feedback and lateral connections appears to be in direct contradiction with the modular character of models such as the FLMP. In interactive models, all units may potentially affect and be affected by the state of all the other units and all the external information sources. It seems counterintuitive that in such models one could possibly have the contribution of a stimulus on the response probabilities being independent of all the other stimuli. Indeed, Massaro (1989c) found that the original version of the interactive activation model did not adhere to the Morton-Massaro law. This raised some interesting questions: Are interactive models intrinsically inconsistent with the Morton-Massaro law? More generally, what, if anything, does the Morton-Massaro law tell us about the organization of perceptual processes? Does it make sense from the point of view of the rationality or optimality of the perceptual process? Is it a consequence of the statistical structure of the environment? Should we expect this law to be violated under some conditions? Our goal is to address these and other related questions. Some of these questions have been previously considered by other authors and are included here to provide a more complete picture of the Morton-Massaro law. In other cases, the answers arise from new

theoretical and empirical results that are presented for the first time in this article.

We consider the phenomenon of factorability in detail, relating it to classical models, including Morton's (1969) logogen model, Oden and Massaro's (1978) FLMP, random utility models, and a statistical technique called logistic regression. Then we identify the conditions under which this phenomenon may represent an optimal strategy of perceptual inference. The analysis suggests that when the Morton-Massaro law holds the environment should have a particular statistical structure. It also suggests that when that statistical structure does not hold the Morton-Massaro law may not hold either. An experiment and a statistical analysis of audiovisual speech signals support these suggestions. Next, we examine whether interactive models may exhibit factorable effects, introduce the concept of channel separability, and consider the relations of this idea to other notions of channel separability and independence.

In the next section, we describe typical experiments in which factorability of information sources has been observed. We formalize the notion of factorability and describe its relationship to the classical models mentioned previously.

Notation and Experimental Paradigm

In the experiments in question, participants are presented with inputs that are combinations of several treatment variables and are asked to report on some property of these inputs by choosing among a set of mutually exclusive response alternatives. In some cases, participants are asked about a perceptual property of one of the treatment variables (e.g., identify the first letter of a three-letter string). That variable is usually known as the stimulus, and the other variables are known as the context. In other cases, participants are asked about the entire input combination. For example, they may be asked about their combined audiovisual experience when a visual presentation of a person saying a word is synchronized with an accompanying acoustic signal. In such cases, the classification of input components as stimulus or context is arbitrary and indicates only that they are different sources of information.

Let X_1, \dots, X_m represent the different treatment factors (sources of information) and R the response obtained for combinations of these treatments. Let the $\mathcal{R} = \{1, \dots, r\}$ represent the set of response alternatives, and \mathcal{X}_i the set of treatment levels of factor X_i . For any set \mathcal{S} , the term $|\mathcal{S}|$ will represent the number of elements in that set, e.g., $|\mathcal{R}| = r$. Participants are repeatedly presented with all possible combinations of the different treatment factors, and the response distribution of each individual participant given each combination of treatments is recorded. If there are only two experimental factors, the response distributions can be organized as a matrix in which the rows represent the treatments of factor X_1 and the columns the treatments of factor X_2 . Each cell in this matrix would contain the probability distribution of response alternatives given a particular combination of treatments.

For example, consider the following hypothetical letter recognition experiment (which we have actually performed and is reported later). On each trial, a three-letter stimulus is presented, followed by a mask, and the participant must decide which of two letters was presented in a particular position (e.g., the second). In this case, there are three sources of information: The stimulus X_2

is the letter presented at the second position, and the other two letters X_1 and X_3 act as the context. The stimulus letters might be E and U, in which case $\mathcal{X}_2 = \{E, U\}$. The contexts might be M_N, R_N, M_D, and R_D, in which case $\mathcal{X}_1 = \{M, R\}$ and $\mathcal{X}_3 = \{D, N\}$. In this case, there would be eight input combinations: MEN, MUN, REN, RUN, MED, MUD, RED, RUD. In our notation, the responses are encoded as numbers, so $\mathcal{R} = \{1, 2\}$, with 1 representing the chosen alternative E and 2 representing the chosen alternative U. For each combination of stimulus and context, we would obtain the probability of Response 1 and Response 2. Thus, there would be a total of $2 \times 2 \times 2 \times (2 - 1) = 8$ independent probability estimates per participant.

In the example just given, there was a one-to-one correspondence between stimuli and response alternatives, but this does not need to be the case. Consider, for example, Massaro and Cohen's (1983) Experiment 2 in which 7 individuals had to identify synthetic consonant sounds presented in the context of other phonemes. There were two response alternatives, $|\mathcal{R}| = 2$, seven stimulus conditions, $|\mathcal{X}_1| = 7$, and four context conditions, $|\mathcal{X}_2| = 4$. The response alternatives were /l/ and /r/, the stimuli were synthetic sounds generated by varying the onset frequency of the third formant followed by the vowel /i/. Each of the seven stimuli was placed after each of four different context consonants: /v/, /s/, /p/, and /t/. Each stimulus by context combination was tested 40 times per individual, thus providing $7 \times 4 \times (2 - 1) = 28$ independent probability estimates per individual.

Similar experiments can also be done with stimulus and context operating in different modalities. For example, Repp, Manuel, Lieberman, & Studdert-Kennedy (1983) presented individuals with all possible combinations of auditory signals of the articulations /ba/, /va/, /ða/, and /da/ and synchronized visual depictions of articulations of the same alternatives, for a total of 16 different audiovisual inputs. Individuals were presented with these inputs and asked to report whether they heard /ba/, /va/, /ða/, or /da/. The experiment provided $4 \times 4 \times 3 = 48$ independent probability estimates per individual (see Massaro, 1989b, for a description of the methods used in Repp et al.'s, 1983, study).

It should be noted that all these experiments are designed to obtain mappings between external input conditions and response probabilities. The experiments do not inform us about the extent to which these mappings reflect perceptual or postperceptual processes. For example, it is theoretically possible that in all these experiments perception itself is not affected by context but that the postperceptual¹ response choices made by participants are. The experiments and the models presented in this article are silent about this issue. What is at hand in these experiments is the nature of the proper characterization of the mappings between information sources and responses, not whether these mappings are perceptual or postperceptual.

Factorability

In this section we describe more precisely the concept of factorability and its relationship to models of perception. Factorability refers to a particular way in which multiple sources of information

¹ By the term *postperceptual* we mean processes that map internal perceptions to external responses.

affect response probabilities. Guided by the logogen model, Morton (1969) proposed that the influence of information sources on perception of words can be accounted for with a factorized version of Luce's (1959) strength model. In Luce's model, a nonnegative strength parameter, denoted with the Greek symbol η , is assigned to each response alternative (indexed here by k) in every combination of information sources. The probability of the response is given by

$$P(R = k|X_1 = x_1, \dots, X_m = x_m) = \frac{\eta(x_1, \dots, x_m, k)}{\sum_{l=1}^r \eta(x_1, \dots, x_m, l)}, \quad (1)$$

where x_1, \dots, x_m represents specific values of the information sources and k and l specific response alternatives. Morton's claim was that these strength parameters can be factorized into nonnegative terms here called η_{X_i} , which are selectively influenced by one factor so that

$$\eta(x_1, \dots, x_m, k) = \eta_{X_1}(x_1, k) \cdots \eta_{X_m}(x_m, k), \quad (2)$$

and

$$\begin{aligned} P(R = k|X_1 = x_1, \dots, X_m = x_m) \\ = \frac{\eta_{X_1}(x_1, k) \cdots \eta_{X_m}(x_m, k)}{\sum_{l=1}^r \eta_{X_1}(x_1, l) \cdots \eta_{X_m}(x_m, l)}. \end{aligned} \quad (3)$$

Here $\eta_{X_i}(x_i, k)$ represents the support provided to the response k when the source X_i takes the value x_i . It is important to note that Equation 3 describes ideal probability values one would obtain with an infinite number of trials, not empirical probability estimates based on a finite number of observations.

Hereafter, we assume that no strength parameter is exactly zero (i.e., we assume that given an arbitrarily large number of trials each response alternative will appear at least once for each input combination). Under this assumption, Equation 3 is equivalent to

$$\begin{aligned} \frac{P(R = k|X_1 = x_1, \dots, X_m = x_m)}{P(R = l|X_1 = x_1, \dots, X_m = x_m)} \\ = \left(\frac{\eta_{X_1}(x_1, k)}{\eta_{X_1}(x_1, l)} \right) \cdots \left(\frac{\eta_{X_m}(x_m, k)}{\eta_{X_m}(x_m, l)} \right). \end{aligned} \quad (4)$$

This says that response probability ratios can be factorized into components selectively influenced by one, and only one, information source. The term $\eta_{X_i}(x_i, k)/\eta_{X_i}(x_i, l)$ is the relative support for responses k and l provided by the source X_i when it takes the specific value x_i . Response biases may be included, but they do not add generality because they can be considered as part of the η_{X_i} terms. Hereafter, we refer to Equation 3 as a factorized strength model.²

A formulation similar to Morton's logogen model arises within the FLMP (Massaro, 1987c; Oden & Massaro, 1978). The FLMP is a factorized strength model in which the strength parameters are constrained to be in the interval (0, 1) and to add up to 1, that is, $\sum_k \eta_{X_i}(x_i, k) = 1$. These provisos occasion no loss of generality. To see why, let $\tilde{\eta}_{X_i}$ represent parameters of a factorized strength model and let

$$\eta_{X_i}(x_i, k) = \frac{\tilde{\eta}_{X_i}(x_i, k)}{\sum_{l=1}^r \tilde{\eta}_{X_i}(x_i, l)}. \quad (5)$$

It is easy to see that the model with parameters $\eta_{X_1}, \dots, \eta_{X_m}$ results in the same response probabilities as the model with parameters $\tilde{\eta}_{X_1}, \dots, \tilde{\eta}_{X_m}$. In addition, the η_{X_i} parameters follow the constraints prescribed by the FLMP, that is, they add up to 1 and are in the (0, 1) range. This shows that response probability matrices conform to a factorized strength model if and only if they conform to the FLMP.

It should be noted that not all factorized models foster the view of human information processing advocated by the FLMP. In particular, Massaro (1987c) interpreted the FLMP parameters in terms of fuzzy logic (Zadeh, 1988) and viewed information integration as a deterministic feed-forward process, involving separate stages of independent evaluation and integration. Evaluation consists of assigning values to the η_{X_i} terms, and integration consists of computing the product of these terms for each response alternative. The only probabilistic part of the process is the decision, which consists of randomly choosing a response based on Equation 3. In addition, there are extensions of the FLMP that make predictions about the changes in choice probabilities as a function of the interstimulus interval between the input and a backward mask (Massaro, 1989b, p. 256). Such extensions are not the focus of this article.

Random Utility Models

Strength models, like the ones we have seen so far, assume strength values are assigned deterministically to each response alternative. Response variability is achieved by choosing randomly with probability proportional to the strength of each response. On the other hand, random utility models assume that choices are made by picking deterministically response alternatives that maximize some internal utility value. In these models, response variability is explained by adding a random component to the utility of each alternative. On each trial, random utility values Y_1, \dots, Y_r are assigned to each response alternative. These random utilities are the sum of an utility function μ_k , which is determined by the external input, and a random noise component N_k , which is independent of the stimulus and context

$$Y_k = \mu_k(X_1, \dots, X_m) + \beta N_k, \quad (6)$$

where $\beta > 0$ is a parameter scaling the effect of noise. On each trial, the alternative with maximum random utility value is chosen. To simplify the presentation in this section, we consider experiments that manipulate only two treatment factors: X_1 and X_2 . The generalization to more than two factors is transparent.

In general, it is difficult to find whether random utility models factorize. However, special cases exist for which analytical results are possible. In particular, McFadden (1978) showed that if the joint cumulative distribution function of the noise is as follows

$$F_N(u_1, \dots, u_r) = e^{-G(e^{-u_1}, \dots, e^{-u_r})}, \quad \text{for } (u_1, \dots, u_r) \in \mathbb{R}^r, \quad (7)$$

² Factorized strength models appear in domains other than perception. For example, the sampling equations of the SAM model of information retrieval (Gillund & Shiffrin, 1984; Shiffrin & Raaijmakers, 1992) conform to our definition of factorized strength models.

where G is a (generator) function satisfying some regularity conditions,³ then

$$P(R = k | X_1 = x_1, X_2 = x_2) = \frac{\partial \log G(e^{\mu_1(x_1, x_2)/\beta}, \dots, e^{\mu_r(x_1, x_2)/\beta})}{\partial \mu_k(x_1, x_2)/\beta} \quad (8)$$

From this it is possible to show that for some noise generator function G and some utility function μ random utility models factorize. In particular, if we let

$$G(u_1, \dots, u_r) = \left(\sum_{i=1}^r u_i^{1/(1-\rho)} \right)^{1-\rho}, \quad \text{for } (u_1, \dots, u_r) \in \mathbb{R}^r, \quad (9)$$

then the noise is a multivariate version of the standard Gumbel noise (Oliveira, 1961). Here $0 \leq \rho < 1$ is a parameter controlling the overall correlation between the noise components N_1, \dots, N_r (i.e., $\rho = 0$) when the noise components are uncorrelated and $\rho = 1$ when they are perfectly correlated. Applying McFadden's Equation 8 to Equation 9 and after some simple derivations, it follows that

$$P(R = k | X_1 = i, X_2 = j) = \frac{e^{\mu_k(x_1, x_2)/(\beta(1-\rho))}}{\sum_{i=1}^r e^{\mu_i(x_1, x_2)/(\beta(1-\rho))}} \quad (10)$$

Thus, if the utility is additively controlled by the different information sources

$$\mu_k(x_1, x_2) = \mu_k^1(x_1) + \mu_k^2(x_2), \quad (11)$$

then response probabilities factorize as required by the Morton-Massaro law (see Figure 5). Moreover, the deterministic utility

values and strength parameters are related as follows

$$\eta_{x_1}(x_1, k) = e^{\mu_k^1(x_1)/(\beta(1-\rho))}, \quad (12)$$

$$\eta_{x_2}(x_2, k) = e^{\mu_k^2(x_2)/(\beta(1-\rho))}. \quad (13)$$

The noise scale parameter β affects the importance of the utility functions on the response probabilities. As β goes to 0, the response with maximum deterministic utility is chosen with probability 1. As β goes to infinity, all responses are chosen with equal probability. In addition, as the correlation parameter ρ goes to 1, the response with maximum deterministic utility is chosen with probability 1. This makes sense because, if the noise components are perfectly correlated, they just add a constant to the utility values of all the alternatives and thus have no effect on the choice process.

Note that in random utility models factorability of information sources corresponds to the decomposition of the overall utility into additive terms selectively controlled by different sources of information. Morton (1969) himself presented the logogen model as a close relative of a random utility model. Each logogen unit produced an activation value (i.e., a utility function), which was the sum of terms selectively controlled by different sources of information (in his case, stimulus and context) plus independent random noise. He assumed that each logogen had a threshold and that if any logogen reached threshold the corresponding word became available as a response (if more than one became available, the first one determined the response). This response selection policy approximates the policy used in random utility models, in which the item with the largest activation is selected.⁴ Unfortunately, Morton's (1969) article did not clarify the relationship between his formulation of the model and the factorized strength model used to describe the data. The relationship between strength models and random utility models appeared in work by Holman and Marley, cited in Luce and Suppes (1965), and was extensively studied by Yellot (1977) and McFadden (1978).

Complexity Analysis

The Morton-Massaro law is the statement that factorized strength models provide good approximations to the choice probabilities obtained in a wide variety of perceptual experiments. One potential explanation for why factorized models fit empirical data so well is that they impose no constraints on response probability matrices (i.e., they can fit any arbitrary probability matrix). For example, it is well known that, given a large number of hidden units, multilayer perceptrons can fit arbitrary input-output functions (Hornik, Stinchcombe, & White, 1989). Thus, it would not be

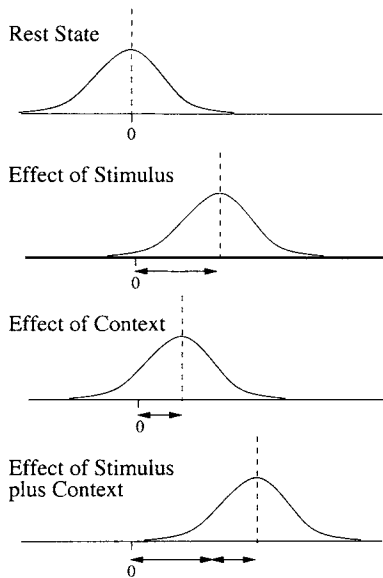


Figure 5. In Morton's logogen model, stimulus and context have additive influence on the state of logogen units. From "The Interaction of Information in Word Recognition," by J. Morton, 1969, *Psychological Review*, 76, p. 167. Copyright 1969 by the American Psychological Association.

³ The regularity conditions are that (1) $G(u_1, \dots, u_r)$ be nonnegative homogeneous of degree one for $u_1 \geq 0, \dots, u_r \geq 0$; (2) $\lim_{u_i \rightarrow \infty} G(u_1, \dots, u_r) = \infty$ for $i = 1, \dots, r$; and (3) for any $k \in N$ and $(i_1, \dots, i_k) \in \{1, \dots, r\}^k$ the partial derivatives $\partial^k G(u_1, \dots, u_r) / \partial u_{i_1} \dots \partial u_{i_k}$ are nonnegative if k is odd and nonpositive if k is even.

⁴ It should be noted that Morton assumed that the threshold of a logogen was reduced by prior use and thus was lower for frequent and recent words. In the corresponding random utility model, this corresponds to assuming that the activation is increased by prior use, so that the item tends to be relatively more active for frequent and recent words.

of interest to observe that response probability matrices are well approximated by such models. In the Appendix we show that this is not a good explanation.

The Appendix shows that in experiments with factors X_1, \dots, X_m and response alternatives $\{1, \dots, r\}$ factorized strength models are fully defined using just $(r - 1)(|\mathcal{X}_1| + \dots + |\mathcal{X}_m| - m + 1)$ parameters, where $|\mathcal{X}_i|$ is the number of treatments for factor X_i . Note that arbitrary response probability matrices have $(r - 1)(|\mathcal{X}_1|) \dots (|\mathcal{X}_m|)$ degrees of freedom. Thus, the complexity⁵ of arbitrary response probability matrices grows exponentially with the number of information factors, whereas the complexity of factorized models grows linearly. The smallest experimental design that allows falsifying factorized models is a 2×2 factorial design with two response alternatives (i.e., $m = 2, r = 2$ and $|\mathcal{X}_1| = |\mathcal{X}_2| = 2$). In this simple case, factorized strength models use $(2 - 1)(2 + 2 - 2 + 1) = 3$ degrees of freedom, whereas arbitrary data matrices can have $(2 - 1)(2)(2) = 4$ degrees of freedom. Table 1 shows an example of a 2×2 probability matrix for which the best approximation achievable with factorized models is grossly inaccurate.

Two Response Alternatives

If there are only two response alternatives (i.e., $r = 2$), then the data in the response matrix are amenable to an analytic technique known as factorial logistic analysis (Myers, 1990). Logistic analysis is based on a simple and quite commonly used statistical model in which treatment variables X_1, \dots, X_m are used to model a dependent variable bounded in the $(0, 1)$ range (e.g., a probability value). To simplify the presentation, we default to the case with two information sources, which we call stimulus (X_1) and context (X_2). The generalization to more than two sources is transparent.

The name “logistic” comes from the fact that the dependent variable is assumed to be related to the treatment variables by the so-called logistic function

$$P(R = 1 | X_1 = x_1, X_2 = x_2) = \frac{1}{1 + \exp(-(\mu + \delta_{x_1, x_2}(x_1, x_2)))}, \quad (14)$$

where $1/(1 + \exp(-(\cdot)))$ is the logistic function, μ is a parameter constant for all the treatments, and $\delta_{x_1, x_2}(x_1, x_2)$ is a parameter known as the logistic joint effect of stimulus x_1 and context x_2 . These parameters are commonly estimated using standard

Table 1
A Nonfactorable Probability Matrix

	$X_2 = 1$	$X_2 = 2$
$X_1 = 1$	0.1	0.9
$X_1 = 2$	0.9	0.1

Note. There are two response alternatives, two stimuli, and two context conditions. The number within each cell represents the probability of the first response alternative given the specific combinations of information sources indicated by the row and column of that cell. In this case, the best mean square approximation achievable with factorized models gives probability 0.5 to each cell, resulting in a root mean square error of 0.4.

maximum-likelihood approaches. As in analysis of variance (ANOVA), the logistic joint effects are constrained to add up to 0 and are decomposed into logistic main effects, $\delta_{x_1}(x_1)$, $\delta_{x_2}(x_2)$, and logistic interaction effects, $\delta_{x_1 \times x_2}(x_1, x_2)$:

$$\delta_{x_1, x_2}(x_1, x_2) = \delta_{x_1}(x_1) + \delta_{x_2}(x_2) + \delta_{x_1 \times x_2}(x_1, x_2), \quad (15)$$

where

$$\delta_{x_1}(x_1) = \frac{1}{|\mathcal{X}_2|} \sum_{x_2 \in \mathcal{X}_2} \delta_{x_1, x_2}(x_1, x_2), \quad (16)$$

$$\delta_{x_2}(x_2) = \frac{1}{|\mathcal{X}_1|} \sum_{x_1 \in \mathcal{X}_1} \delta_{x_1, x_2}(x_1, x_2), \quad (17)$$

$$\delta_{x_1 \times x_2}(x_1, x_2) = \delta_{x_1, x_2}(x_1, x_2) - \delta_{x_1}(x_1) - \delta_{x_2}(x_2). \quad (18)$$

The number of free parameters divides up as follows: one parameter for the constant μ , $|\mathcal{X}_1| - 1$ parameters for the main effects of X_1 , $|\mathcal{X}_2| - 1$ parameters for the main effects of X_2 , and $(|\mathcal{X}_1| - 1) \times (|\mathcal{X}_2| - 1)$ parameters for the interaction effects. As in ANOVA, this partition follows from the fact that the joint effects are constrained to add up to 0.

As pointed out by Massaro and Friedman (1990) and by Crowther, Batchelder, and Hu (1998) in experiments with only two response alternatives, factorized strength models are equivalent to logistic models with no interaction effects. To see why, let

$$p(x_1, x_2) = P(R = 1 | X_1 = x_1, X_2 = x_2), \quad (19)$$

then if factorability holds,

$$p(x_1, x_2) = \frac{\eta_{x_1}(x_1, 1)\eta_{x_2}(x_2, 1)}{\eta_{x_1}(x_1, 1)\eta_{x_2}(x_2, 1) + \eta_{x_1}(x_1, 2)\eta_{x_2}(x_2, 2)}, \quad (20)$$

assuming none of the η_{x_1} and η_{x_2} terms are 0, this can be written as

$$p(x_1, x_2) = \frac{1}{1 + \exp(-(\mu + \delta_{x_1}(x_1) + \delta_{x_2}(x_2)))}, \quad (21)$$

where

$$\mu + \delta_{x_1}(x_1) = \log \frac{\eta_{x_1}(x_1, 1)}{\eta_{x_1}(x_1, 2)}, \quad (22)$$

$$\delta_{x_2}(x_2) = \log \frac{\eta_{x_2}(x_2, 1)}{\eta_{x_2}(x_2, 2)}, \quad (23)$$

showing that the Stimulus \times Context interaction effects are 0, thus saving $(|\mathcal{X}_1| - 1) \times (|\mathcal{X}_2| - 1)$ degrees of freedom.

Modularity and Statistical Interaction

It should be noted that the computational notion of interactivity, as an antithesis to modularity, has very little to do with the

⁵ We informally use the term *complexity* to represent the number of nonredundant parameters in a model. This notion is adequate for our work, because we operate within the framework of Luce strength models. For other notions of model complexity, see Vapnik (1995).

statistical notion of interaction in ANOVA. In particular, factorized strength models can be modular (like the FLMP), yet they can produce some interaction effects. Consider, for example, the case in which there are only two experimental factors and two response alternatives. If the untransformed probability values predicted by factorized models are plotted against the treatments of one factor for each of the treatments of the other factor, this results in a family of curves spanning a region shaped somewhat like an American football or a portion of one. This shape is a characteristic signature of factorized strength models and the Morton-Massaro law (see Massaro, 1989b, p. 107). Such patterns can yield significant statistical interaction effects when their raw response probabilities are subjected to ANOVA.

Although factorized strength models can produce some interaction effects, when there are only two response alternatives, they cannot produce crossover interactions. To see why, fix the context to x_2 and examine the difference in logit probabilities for two different stimuli x_1 and \bar{x}_1 . Using Equation 21,

$$\text{logit } p(x_1, x_2) - \text{logit } p(\bar{x}_1, x_2) = \delta_{x_1}(x_1) - \delta_{x_1}(\bar{x}_1), \quad (24)$$

where the logit function $\text{logit}(p) = \log(p/(1-p))$, is just the inverse of the logistic, so that if $p = \text{logistic}(x)$, then $\text{logit}(p) = x$. Note how the difference of logit probabilities is not dependent on the context x_2 . The transformed data, that is, $\text{logit}(p(x_1, x_2))$ can be plotted against the stimulus indices x_1 for each of the different values of x_2 ; if the Morton-Massaro law holds, the curves should all be parallel if it were not for deviations resulting from sampling error. Moreover, because the logit function is strictly an increasing transformation, it follows that

$$\text{sign}(p(x_1, x_2) - p(\bar{x}_1, x_2)) = \text{sign}(\delta_{x_1}(x_1) - \delta_{x_1}(\bar{x}_1)), \quad (25)$$

which is also independent of the context (x_2). Thus, if $\delta_{x_1}(x_1) - \delta_{x_1}(\bar{x}_1)$ is positive, then switching from stimulus x_1 to stimulus \bar{x}_1 can only increase the probability of response 1 regardless of the context (x_2). Similarly the probability can only decrease if $\delta_{x_1}(x_1) - \delta_{x_1}(\bar{x}_1)$ is negative. In other words, when there are only two response alternatives, factorized models cannot produce crossover interactions on the choice probabilities. This property is useful for quick verification of gross violations of the Morton-Massaro law. If there are only two response alternatives, then significant crossover interactions indicate violations of the Morton-Massaro law. It should be pointed out that when there are more than two response alternatives, crossover interactions do not necessarily violate the law.

Functional Analysis

As mentioned earlier, we do not use the term Morton-Massaro law to imply that it applies in all perceptual problems at all levels of analysis. Indeed, if the law applied at all levels of the perceptual process, the problem of machine perception might have been solved long ago. One aspect that makes machine perception so difficult is that information must sometimes be integrated in a nonfactorable manner. It is thus important to question under what conditions the Morton-Massaro law reflects a sensible way to combine sources of information. Functional questions like this are an essential component of

scientific inquiry in disciplines such as biology. Unfortunately, cognitive psychology has tended to focus on structural questions (e.g., the modularity debate) while paying little attention to functional questions. See Marr (1982), Anderson (1990), and Knill and Richards (1996) for notable exceptions.

In this section, we adopt a Bayesian framework to formulate a notion of optimal perceptual inference and examine under what conditions and in what sense the Morton-Massaro law conforms to this definition of optimality. The goal is to gain an understanding of the function of the Morton-Massaro law and its relationship to the statistics of the environment in which it develops. For other perspectives on the optimality of the Morton-Massaro law, see Massaro and Friedman (1990), Movellan and Chadderdon (1996), Friedman and Massaro (1988), and Massaro (1989b, p. 272).

From a Bayesian perspective, human perception is viewed as a form of probabilistic inference (Knill & Richards, 1996; Clark & Yuille, 1990), which is describable in terms of subjective probabilities. Bayesian approaches do not assume that such probabilities are explicitly represented by the perceptual system (e.g., as they would be in a classic Bayesian expert system). Subjective probabilities can be simply seen as descriptive variables used by Bayesian analysts to help understand the logic of the outcome of perceptual processes. Within this framework, one can typically analyze in what sense and under what conditions these subjective probabilities and the inferences they afford might be seen as optimal.

We consider this question in the problem of choosing from among a set of alternatives (perceptual hypotheses or categories) when given multiple sources of information about those alternatives. The approach will help us understand the ubiquity of the Morton-Massaro law. As in previous sections, we use the random variables X_1, \dots, X_m to represent sources of information, and the random variable R to represent the individual's responses. The variables X_1, \dots, X_m provide information about a variable A , which represents the alternatives (e.g., letters, words, phonemes) the participant is required to consider, in the sense that he or she must designate one of these as the response.

It is assumed that the perceptual system has no direct access to the value of A , and its role is to infer it from the available information sources. At this level of analysis, perceptual systems are just probabilistic mappings whose inputs are the information sources, X_1, \dots, X_m , and whose outputs are probability distributions of responses (i.e., two perceptual systems with identical mappings are considered equivalent regardless of internal mechanisms). We denote these mappings as $p_{R|X_1, \dots, X_m}$, referring to the distribution function of responses conditional on all possible values of the information sources. The notation $p_{R|X_1, \dots, X_m}(k|x_1, \dots, x_m)$ corresponds to the probability that R takes the value k given that X_1 takes the value x_1 and X_m takes the value x_m . In this context, we investigate in what sense and under what conditions perceptual systems (i.e., distribution functions $p_{R|X_1, \dots, X_m}$) that adhere to the Morton-Massaro law are optimal.

At this point, it is important to make a distinction between true environmental probabilities and subjective estimates of these probabilities, which are internal to the mechanism that is making perceptual decisions. The idea is that human observers do not have direct access to the probability values given by the environment. Instead, the human perceptual system must work (explicitly or implicitly) with quantities that are best seen as internal or subjec-

tive estimates of the relevant probabilities. To make this distinction clear, we mark such internal subjective probability estimates with a tilde.

A common criterion for optimality is minimization of the subjective error rate (i.e., maximization of the subjective probability that the response equals the correct alternative). It should be apparent that policies that satisfy the following rule are optimal with respect to that criterion (Duda & Hart, 1973)

$$p_{R|X_1, \dots, X_m}(r|x_1, \dots, x_m) = 0 \text{ if there is } r' \in \mathcal{R} \text{ such that}$$

$$\tilde{p}_{A|X_1, \dots, X_m}(r'|x_1, \dots, x_m) > \tilde{p}_{A|X_1, \dots, X_m}(r|x_1, \dots, x_m). \quad (26)$$

In words, this rule says that we should never choose a response r if some other response r' is estimated to be more likely to correspond to the correct alternative. So if one response is more likely to correspond to the correct alternative than any other, we should always choose that response. If there are two or more responses tied for being the most likely alternative, then all distribution of responses among these alternatives are equally optimal. In the pattern recognition literature, this policy is known as "maximum a posteriori," or MAP, because it maximizes the subjective estimate of the posterior probability that the response is correct, given the data available to the perceiver.

Maximizing the posterior probability is equivalent to maximizing a statistic known as the expected discriminant value (Duda & Hart, 1973), which can be expressed as follows

$$E^P(\log \tilde{p}_{X_1, \dots, X_m|A}(X_1, \dots, X_m, R))$$

$$= \sum_{x_1} \dots \sum_{x_m} \sum_r p_{X_1, \dots, X_m|R}(x_1, \dots, x_m, r)$$

$$\log \tilde{p}_{X_1, \dots, X_m|A}(x_1, \dots, x_m, r). \quad (27)$$

In words, the expected discriminant value is the average over all combinations of inputs and responses of the log of the subjective estimate of the probability that the given input has occurred and the response r actually corresponds to the alternative that generated the input. Large values of the expected discriminant value indicate that the policy is effectively exploiting the subjective information about A contained in X_1, \dots, X_m . MAP policies achieve the largest possible expected discriminant value.

Having introduced the basic ideas of Bayesian decision theory, let us return now to the optimality of the Morton-Massaro law. The first thing that we can note is that the MAP policy is, in general, a deterministic function of the information sources. If there is a single alternative that is subjectively most likely to be correct for a given combination of X_1, \dots, X_m , then the MAP policy would always choose that response when that particular combination of information sources is presented. However, in experiments in which participants believe that one alternative is more probable than the others, they consistently still choose the less probable alternatives a substantial number of times (Friedman & Massaro, 1988).

One way to explain this deviation from MAP is to construe participants as maximizing a combination of entropy and subjective discriminant value

$$E^P(\log \tilde{p}_{X_1, \dots, X_m|A}(X_1, \dots, X_m, R)) + \beta H^P(X_1, \dots, X_m, R), \quad (28)$$

where

$$H^P(X_1, \dots, X_m, R) = - \sum_{x_1} \dots \sum_{x_m} \sum_r p_{X_1, \dots, X_m|R}(x_1, \dots, x_m, r)$$

$$\log p_{X_1, \dots, X_m|R}(x_1, \dots, x_m, r), \quad (29)$$

is the joint entropy of the information sources and the response. The joint entropy is an information theoretic measure of the dispersion of the joint distribution of information sources and responses (Cover & Thomas, 1991). The parameter $\beta > 0$ is fixed and controls the relative importance of the entropy versus the subjective discriminant value. We can now introduce a theorem and a corollary for the Morton-Massaro law that suggest in what sense and under what conditions the law is optimal. The proof for the theorem and its corollary are included in the Appendix. Here we focus on their heuristic interpretation.

Optimality Theorem

A Bayesian system that satisfies the two following conditions conforms to the Morton-Massaro law:

1. It maximizes a weighted sum of subjective discriminant value and entropy

$$E^P(\log \tilde{p}_{X_1, \dots, X_m|A}(X_1, \dots, X_m, R)) + \beta H^P(X_1, \dots, X_m, R), \quad (30)$$

where E^P , H^P , respectively, symbolize expected values and entropy with respect to the probability measure P , and $\beta > 0$ is a fixed parameter.

2. It assumes a distribution of information sources, which factorizes as follows:

$$\tilde{p}_{X_1, \dots, X_m|A}(x_1, \dots, x_m|a)$$

$$= \phi_A(a) \phi_X(x_1, \dots, x_m) \phi_1(x_1, a) \dots \phi_m(x_m, a), \quad (31)$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$. Here ϕ_A is a term selectively influenced by A , not X ; ϕ_X is selectively influenced by X , not A ; and the terms ϕ_i are selectively influenced by the i th element of X and by A .

Corollary

$p_{R|X_1, \dots, X_m}(a|x_1, \dots, x_m)$ satisfies the Morton-Massaro law for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$ if and only if there is a Bayesian system that satisfies Conditions 1 and 2 of the optimality theorem and whose response probabilities equal $p_{R|X_1, \dots, X_m}(a|x_1, \dots, x_m)$ for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$.

Examining Condition 1. The first condition required by the optimality theorem is that the perceptual system should maximize not just the expected value of the subjective probability of being correct, as a pure Bayesian system would do, but a combination of this and response entropy. A deterministic policy would be associated with a single response), and distributing the probability of choice evenly across the available alternatives would yield the greatest degree of dispersion. Maximization of entropy, not just discriminant value, is beneficial when the subjective internal probability estimates are not certain or the environment is nonstation-

ary. In such cases, some entropy in responding helps sampling the environment for the purpose of improving the subjective probability estimates (Dayan & Sejnowski, 1996). This positive value of entropy is well known in the pattern recognition literature, where it is presented in terms of a trade-off between “exploration” and “exploitation” (Sutton & Barto, 1988).

To illustrate this trade-off, consider the following example from a real-life application of Bayesian inference. A World Wide Web server has a set of banners that may be displayed for advertisement purposes. The server gets paid every time a client clicks on the banners. To avoid overwhelming the client with advertisements, only one banner is displayed per page. The banner chosen to be displayed plays a role analogous to the response R in our analysis. Each time a client visits the Web site, the server can obtain partial information about the client (e.g., zip code, age, gender). These sources of information are analogous to X_1, \dots, X_m in our analysis. Sophisticated Web servers typically have a Bayesian model, which provides estimates of the posterior probability that the client clicks on each of the banners given the information known about the client. These estimates are the equivalent of the subjective probabilities $\tilde{p}_{A|X_1, \dots, X_m}$ in our analysis.

The goal of the server is to maximize the number of clicks. If the probability estimates were perfect, then the best strategy would be to display the banner that has the maximum probability of being clicked given the information at hand. However, typically the probability estimates are not perfect. One reason for this is that people’s preferences change (i.e., the environment is nonstationary), and estimates based on past experience may need recalibration. If the system displayed only the most probable banner, it would lose the opportunity to learn that now other banners may be more effective. Thus, it is important for the system to display banners that have low subjective probability of success. At the same time, the goal is to make money; thus, it is also important to give more priority to those banners that have a higher subjective probability of success. This contradiction is at the heart of the exploration-exploitation dilemma.

From a Bayesian point of view, the problem faced by the Web server is not unlike that faced by perceptual systems. Both systems need to balance exploitation of internal knowledge with the need of calibrating that knowledge. Equation 30 provides a way of formalizing a common heuristic used in the pattern recognition literature to balance exploration and exploitation (Sutton & Barto, 1988). The parameter β in Equation 30 can be seen as an “uncertainty parameter” that modulates the importance of exploration versus exploitation. In the limit, as $\beta \rightarrow 0$ the obtained policy is that of a pure Bayesian observer, which bases its decision on the subjective probability measures. As β goes to infinity, all responses are chosen with equal probability, regardless of their subjective value, as might be appropriate if the probability estimation process is viewed as completely uncertain. If $\beta = 1$, then the objective probability of choosing an alternative equals the subjective probability of that alternative (a form of probability matching). The optimal value of β depends on the statistical properties of the environment and the validity of the subjective probability estimates.

The entropy term in Equation 28 can also be viewed as an undesirable but unavoidable force that is influencing processing (i.e., noise). Under this interpretation, the individual maximizes the weighted sum by attempting to maximize the first term, which is

the only one over which he or she actually has control. The noise can still be viewed as nature’s way of forcing a sampling of alternatives, which may lead to new information allowing updating of subjective probabilities.

In our consideration of Condition 1, we have acted as though correctness of the response choice is equally important for all stimuli and all responses. However, it is possible to take into account differences in the benefit of making particular responses when they are correct and in the costs of making particular responses when they are incorrect. The optimal policy can be derived if estimates of these costs and benefits are available, and can be incorporated into the representation of the subjective discriminant value.

Examining Condition 2. The second condition of the optimality theorem amounts to assuming a particular statistical structure of the world. The assumption only needs to hold for the specific values of the information sources and response alternatives probed in the experiments in which the Morton-Massaro law is tested. An important special case that satisfies Condition 2 occurs when the information sources are assumed to be independent within each perceptual class or category. This condition is called class-conditional independence, or just conditional independence. (For a consideration of other cases in which Condition 2 holds, see the Appendix). Figure 6 illustrates the concept of conditional independence (see also Movellan & Chadderdon, 1996, and Massaro, 1989b). Each point in Figure 6 represents an instance of a stimulus-context combination, with the symbol indicating the response category that each combination comes from. On the left, the stimulus and context are roughly independent within perceptual categories (class conditional independence) but are nonindependent when collapsed across categories. For example, knowing that an item is an example of the square category tells us that the value of the context will be intermediate, and knowing the value of the stimulus adds no further information. On the right, stimulus and context are roughly independent when collapsing across categories even though they are nonindependent within categories. For example, in this case, if we know that an item is an example of the square category, it is hard to specify the value of the context, but if we also know that the stimulus value is high we can predict that the context value will be low.

Conditional independence is sufficient to satisfy Condition 2, but it is not necessary. Figure 7 shows distributions consistent and inconsistent with Condition 2. The horizontal and vertical axes




Figure 6. An illustration of the difference between conditional and unconditional independence. Each point represents a stimulus by context combination. There are three different perceptual categories, represented as circles, squares, or triangles. On the left, stimulus and context are class conditionally independent, because the two dimensions are roughly independent within each category. However, overall, stimulus and context are not independent. On the right, stimulus and context are not conditionally independent but are roughly independent overall.

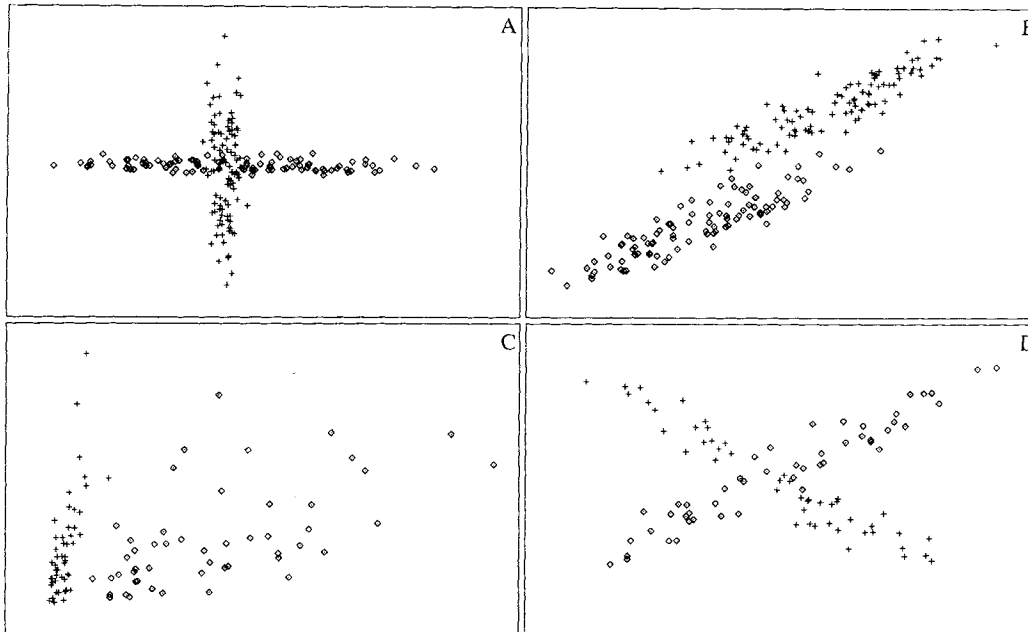


Figure 7. The horizontal and vertical axes represent two sources of information. The type of point (cross vs. diamond) represents the underlying perceptual category. In A, B, and C, the Morton-Massaro law provides an optimal way of combining information. In D the law would not be an adequate way of combining information.

represent two sources of information. The underlying perceptual category is depicted using different point patterns. Each figure has 100 points sampled from a theoretical probability distribution. In A the distribution is class conditionally independent (i.e., independent within each perceptual category). In B the two perceptual categories are Gaussian with equal correlation structure but different means. In C the samples were taken from Gaussian distributions with equal correlation structure and then passed through a monotonic transformation, the exponential function. In D the perceptual categories are Gaussian with different correlation structure. Case D is incompatible with the Morton-Massaro law. In this case, discrimination of the categories requires a crossover interaction of response probabilities which, as shown earlier, is incompatible with the law. In the Appendix, we show that cases A, B, and C satisfy Condition 2; thus, they can be optimally discriminated using the Morton-Massaro law.

Condition 2 is well known in the pattern recognition literature, where it is often called the “naive Bayesian assumption.” This assumption has been shown to work surprisingly well in a very wide variety of pattern recognition problems (Domingos & Paz-zani, 1997). It should be noted that Condition 2 involves subjective probabilities and thus is internal to the perceptual system, not a statistical property of the environment. However, because human perception works quite well under many conditions, it is reasonable to infer that, when the Morton-Massaro law holds, the naive Bayesian assumption provides an adequate approximation to the statistical structure of the environment. Indeed, one might propose a functional principle that the human perceptual system has evolved to conform to the statistical structure of the environment.⁶ This principle suggests that the ubiquity of the Morton-Massaro law may ultimately be due to the fact that, for the type of cases in

which the Morton-Massaro law holds, Condition 2 provides a good working approximation to the actual structure of the environment. This principle further suggests that when Condition 2 does not provide a good approximation to the environment we should expect violations of the Morton-Massaro law. Of course, it is an empirical matter to determine just how adequately the human perceptual system does reflect the statistics of the environment. In the next two sections, we present evidence relevant to this issue. First, we present a statistical analysis of audiovisual speech signals and show that conditional independence of audio and video inputs given (within) words is a reasonable assumption. This helps explain why the Morton-Massaro law has been found to hold so well in experiments involving the integration of auditory and visual speech (Massaro, 1989b). Second, we present a case in which Condition 2 grossly misrepresents the environment. In such case, the optimality theorem suggests that violation of the Morton-Massaro law may occur. An experiment consistent with this prediction is presented.

Statistical Analysis of Audiovisual Speech

Visual information about articulation is an important source of information in face-to-face speech perception. Sensitivity to correspondences in auditory and visual information for speech events has been shown in 4-month-old infants (Kuhl & Meltzoff, 1982). By 6 years of age humans consistently use audiovisual contingencies

⁶ For the moment, it is irrelevant whether this conformity arises through direct specification or through an adaptive process that shapes the perceptual system to the characteristics of its environment through experience.

cies to understand speech (Massaro, 1987b), and by adulthood visual articulation automatically modulates perception of the acoustic signal (McGurk & MacDonald, 1976). Massaro (1989b) investigated how acoustic and visual sources are integrated in speech perception and found that the FLMP does a good job at describing the experiments. According to the corollary of the optimality theorem, this suggests that acoustic and visual speech signals may indeed be treated in a factorized manner without much loss of information.

We investigated the plausibility of this hypothesis by analyzing a database of audiovisual speech signals using statistical machine perception techniques. The main point of the exercise was to illustrate how one may go about testing whether conditional independence provides reasonable approximations to the posterior probabilities of perceptual categories. Our approach involved testing whether assuming conditional independence hinders the performance of automatic audiovisual speech recognition systems, which are trained to extract the statistical structure of audiovisual speech. The idea is to test a general class of speech recognition systems that do not assume conditional independence against a restricted version of that class in which we impose the assumption of conditional independence. These two classes are then optimized using a database of audiovisual speech; if the best recognizer from the constrained class is not worse than the best recognizer from the unconstrained class, we have an indication that conditional independence is a reasonable assumption.

Method

Database. We collected a database of audiovisual speech from 9 male and 3 female undergraduate students from the Cognitive Science Department at the University of California, San Diego (UCSD). Each student was asked to utter the digits 1 through 4 twice. Thus, the total database consists of 96 utterances. The audio sampling rate was 11.1 kHz with 8 bits per sample. The video signal was digitized at 30 frames per second, grey scale, 100×75 -pixel image, 8 bits per pixel. The video signals show only the students' lips. Participants were asked to center and align their lips in the camera during the sampling. The database, which we named "Tulips1" can be found at <http://cogsci.ucsd.edu> following the links to technical reports and software.

Visual front end. Automatic recognition of visual speech is still an emerging field, and a great deal of effort is focused on finding appropriate image processing techniques (Gray, Movellan, & Sejnowski, 1997; Hennecke, Stork, & Ventakesh Prasad, 1996; Luettin, Thacker, & Beet, 1996; Prasad, Stork, & Wolff, 1993). The approach we used here is one of the most successful in the field. First, each image frame was symmetrized by averaging pixel by pixel the left and right side of each image, using the vertical midline as the axis of symmetry. The benefit of this transformation is resistance to changes in illumination along the horizontal axis. For each image frame, an associated delta image was obtained by subtracting the previous image from the current image. Delta images provide resistance to changes in illumination and emphasize moving regions of the lips.

Each image and its associated delta image were low-pass filtered and subsampled at 10×15 equidistant points. The resulting pixel values were soft thresholded and scaled in the (0, 1) range using a logistic function matched to the statistics of the images

$$y = f\left(\frac{\pi}{\sqrt{3}\sigma}(x - \mu)\right), \quad (32)$$

where f is the logistic function, and μ and σ are, respectively, the average and standard deviation of the distribution of image intensities across an

entire word. This transformation is a fast approximation to histogram equalization, a standard image enhancement technique.

Composites of the relevant portions of the low-pass filtered images and their associated delta images were fed to the statistical recognition engine. The number of variables per processed image was 300 (150 pixel values from the original images and 150 from the delta images). Figure 8 shows the effect of the different preprocessing stages.

Acoustic front end. The acoustic signals were preprocessed using LPC/cepstral analysis. This is a standard technique in the speech recognition literature that parameterizes an estimate of the human vocal tract's transfer function (Rabiner & Juang, 1993). First, the high-frequency components of the audio track were enhanced using a first-order linear filter. Then the audio signal was divided into nonoverlapping frames synchronized with the visual frames (i.e., each audio frame lasted $\frac{1}{30}$ of a second). Each frame was multiplied by a Hamming window to reduce artifacts resulting from the blocking procedure (Rabiner & Juang, 1993). LPC analysis was performed on each frame, yielding 8 LPC coefficients per frame. These coefficients were transformed into a log-power coefficient and 12 cepstral coefficients (i.e., Fourier coefficients of the LPC estimates of the log power spectrum; Rabiner & Juang, 1993). The cepstral coefficients were weighted to minimize the effect of the highest and lowest order cepstrals. Then, for each frame, a differential log-power coefficient and 12 differential cepstral coefficients were calculated using polynomial approximations with three frames before and after the desired frame. This resulted on a total of 26 coefficients per audio frame (12 cepstrals, 12 differential cepstral, 1 log-power, and 1 differential log-power). Each of these coefficients was linearly scaled to assume values in the (0, 1) range.

Recognition engine. The recognition engine was a bank of hidden Markov models (HMMs), which are the dominant approach in audio and visual speech recognition. These are stochastic models whose parameters can be tuned to approximate the statistical structure of time-varying sequences (Rabiner & Juang, 1993). Each HMM was trained using the expectation-maximization algorithm (Rabiner & Juang, 1993) to estimate the log-likelihood of the speech signal given a perceptual hypothesis (one HMM was trained with examples of the word "one," another with examples of the word "two," and so on). Because in our database each of the four perceptual alternatives had equal prior probability, the log-likelihood estimates provide the "subjective" discriminant values on which the MAP rule operates. Classification of the speech signals proceeds by calculating, for each model, the log-likelihood of the observed sequence given that model and choosing the model with maximum log-likelihood.

We tested two different types of architectures, one of which was a restricted version of the other. The restricted version consisted of two independent banks of HMMs: one visual and the other auditory (Figure 9). These were trained on their respective feature data sets, and during testing the resulting responses of each bank were integrated to obtain an audiovisual classification. Under the assumption of conditional independence,



Figure 8. Image processing steps: (1) raw image, (2) symmetrized image, (3) delta image, (4) final composite.

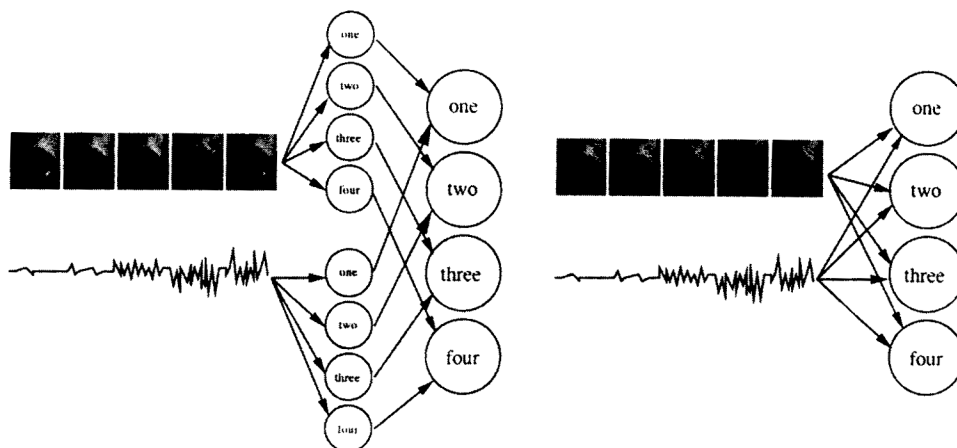


Figure 9. Two approaches to automatic audiovisual speech recognition. The factorized approach (left) uses separate banks of hidden Markov models for the audio and visual signal. This approach saves a large number of parameters but is justified only if the audio and video signals are conditionally independent given the categories to be recognized. The unfactorized approach (right side) is more general and does not assume conditional independence. The first approach can be shown to be a special case of the second approach.

this amounts to summing the auditory and visual log-likelihood responses for each word category and selecting the category that has the highest sum as the winning classification. Hereafter, we refer to this architecture as factorized.

We also tested an unfactorized architecture that did not assume conditional independence. The models used in this architecture consisted of a single bank of audiovisual HMMs. During training, visual and auditory feature data were mixed together into a single observation sequence, in which each time frame contained the concatenation of the visual and auditory feature vectors (i.e., 326 coefficients per frame).

Testing procedure. The two architectures were evaluated in terms of their generalization performance (i.e., their performance when tested with individuals other than those used for training the system). Tukey's jack-knife (Efron, 1982) method was used to obtain generalization estimates: The two architectures were trained, leaving out the utterances of 1 of the 12 participants, and testing was done on the utterances of the excluded participant. This process was repeated 12 times, leaving out a different participant each time. Jack-knife estimates were based on the average generalization obtained with these 12 samples.

HMMs differ in the number of hidden states and Gaussian mixtures per state. For a fixed number of hidden states and mixtures, the parameters trained using the EM algorithm were (a) the hidden state transition probability matrix, (this matrix was constrained to be upper triangular); (b) the centroids of each Gaussian mixture, and (c) the covariance matrix of the mixtures. These covariances were constrained to be of the form σI , where $\sigma > 0$ is an adaptive parameter, and I is the identity matrix. Note that this structure allows one to capture high-order dependencies, because multiple Gaussian mixtures are used. To give a fair chance to the factorized and unfactorized architectures, we optimized them by testing 45 different variations of the two architectures with different number of states (2, 3, 4, 5, 6) and Gaussian centroids per state (2, 3, 4, 5, 6, 7, 8, 9, 10). Pilot work showed that optimal architectures were located within this region. This optimization was very costly computationally but was a crucial part of this exercise. By selecting a particular value of Gaussian centroids and number of states, one could find that the factorized architecture performed better or worse than the unfactorized architecture. The crucial point here is whether the best models within the factorized architecture are as good as the best models within the unfactorized architecture. We chose the best two models within each architecture (factorized and unfactorized) as representative of the entire architecture and report average generalization perfor-

mance with those two best models. For the factorized architecture, the best auditory models had five states/seven Gaussians and five states/eight Gaussians. These models, respectively, had 949 and 1,084 free parameters. The best visual model (which was used with the two best auditory models) had three states/three Gaussians for a total of 2,711 free parameters. For the unconstrained architecture, the best models had three states/eight Gaussians and five states/seven Gaussians; they had 7,850 and 11,449 free parameters, respectively.

Results

When performance is measured in terms of generalization to new data, it is actually possible for restricted versions of a model to perform better than more general versions as long as the restricted version provides a good working approximation to the statistics of the environment. This has to do with a well-known trade-off between bias and variance of statistical models (Geman, Bienenstock, & Doursat, 1992). This is precisely what happened in our study: The best two factorized architectures gave an average generalization performance of 98% correct, slightly better than the average performance of the best two unfactorized architectures (95.5%). Thus, conditional independence provided a close enough approximation to the actual statistical structure of the signals so that the bias introduced by the assumption was well compensated by the reduction in the number of free parameters (the best unfactorized models had about three times as many parameters as the best factorized models).

Discussion

The main point of this exercise was to illustrate how machine perception techniques may be used to test for conditional independence of natural, high dimensional signals. We know that the Morton-Massaro law holds well when humans integrate audio and visual speech (Massaro, 1989b). According to the optimality theorem, this suggests that conditional independence of acoustic and visual signals approximates well the posterior probabilities of

perceptual categories in audiovisual speech. Therefore, assuming conditional independence should not result in a significant loss of information when the assumption is applied to artificial audiovisual speech recognition systems. We tested this prediction by building state-of-the-art audiovisual speech recognition systems with and without the assumption of conditional independence. We found that the assumption of conditional independence did not hinder performance. This suggests that indeed audio and visual speech signals may be treated as though they were conditionally independent without significant loss of information about speech categories.

The experiment we report illustrates an approach to these issues and is certainly consistent with the idea that independent treatment of audio and visual speech is reasonable. The main limitation of our analysis is that our database was relatively small. Thus, our work should be seen only as exploratory research. We had to use a small database because the convincingness of the analysis depended on the fact that the factorized and unfactorized architectures were as good as they could possibly be. This required optimizing over the space of HMMs, which is a very time-consuming process. As computing power increases, it should be possible to perform this type of study with larger databases.

The Conjunctive Context Effect

The optimality theorem suggests that violations of the Morton-Massaro law may occur when objective probability ratios of perceptual categories do not factorize. A particularly obvious violation of factorability occurs when there are only two categories and optimal information integration requires crossover interactions (see earlier discussion). Previous research has shown that individuals have no difficulty learning category structures that require such interactions (Kruschke, 1996). In this section, we present a case of a crossover interaction (i.e., a violation of the Morton-Massaro law) in a letter perception task.

Consider the task of recognizing three-letter words in English. The goal is to choose between two alternatives E and U for the identity of the middle letter in the following contexts: M_N, R_D, M_D, R_N. In this case, we have three sources of information, corresponding to the three letters in the word. The stimulus is whatever letter appears in the second position. The other two letters (which are now part of the context) provide additional information about the second letter because, in written English the strings MEN and RED are high-frequency words but MUN and RUD are not, whereas MUD and RUN are high-frequency words but MED and REN are not. If we look at any context letter (let us say R), the constraint it places on the identity of the target letter reverses as we manipulate the identity of the other letter in the context. Taken by itself, R is consistent with both E or U, because it occurs with E in RED and with U in RUN. When taken together with one of the context letters, however, R unequivocally supports only one of the two alternatives. When the final letter is N, R supports middle U; when the final letter is D, R supports middle E.

Consider how a factorized model may work here. First, three autonomous modules, one per letter location, would selectively evaluate the letter presented in a particular location and would output quantities that would indicate the support of that location for E and U on the second location. Second, the output of these

three modules would be integrated using a multiplication rule. Third, an alternative would be chosen with probability proportional to the product of the three evaluations.

In English, we have that

$$p(R_N|_U_) \approx 1, \quad (33)$$

$$p(R_N|_E_) \approx 0, \quad (34)$$

and thus the context R_N strongly supports the alternative U for the middle letter. However, if we assume conditional independence, an analysis of the Kucera and Francis (1967) corpus restricted to three-letter words tells us that

$$p(R_|_U_)P(_N|_U_) = \frac{244}{9555} \frac{493}{9555} = \frac{132}{100000}, \quad (35)$$

$$p(R_|_E_)P(_N|_E_) = \frac{256}{11178} \frac{1095}{11178} = \frac{224}{100000}. \quad (36)$$

Thus, if we assume conditional independence, R_N would strongly support E instead of U for the middle letter. Although the assumption of conditional independence will not always support the nonword alternative, it is apparent that a model with such an assumption may have problems producing a strong word superiority effect when the context letters combine in a conjunctive manner.

Upper Bound on the Predicted Effect

Consider how the factorized model would be parameterized when the inputs are the letter combinations MEN, MUN, REN, RUN, MED, MUD, RED, RUD, the task is to recognize the middle letter, and the response alternatives are E, and U. In this case, we would need two parameters to describe the support of M__ and R__ for the E alternative, two parameters for the support of _E_ and _U_ for the E alternative, and two parameters for the support of __D and __N for the E alternative. Thus, the factorized model would use a total of six parameters (in fact, as we saw in a previous section two of these parameters would be redundant). Similar models could be used when the task at hand is to recognize the first or the last letters.

If our dependent variable were $\text{logit}(p)$, where p is the probability of correct identification, then this model would predict no difference between the word and nonword conditions. One problem with the logit function is that it is undefined for extreme values of p , and thus it is desirable to work with raw probability estimates rather than their logit transformation. Unfortunately, because the logit transformation is nonlinear, it is possible to find parameter values of the factorized model that produce a small word superiority effect when the dependent variable is the raw probability of correct identification. To take care of this problem, we specified the largest possible word superiority effect achievable by the factorized model. First, we performed a grid search looking for combinations of parameter values that maximized the difference in error rate between word and nonword stimuli for groups of three letters of the type described previously. The search was constrained by the following conditions:

1. A parameter set was not used if, for any of the words in the word set, both the stimulus and context support the incorrect alternatives. For example, consider the case of the item MEN

with the E in the second letter position taken to be the stimulus. A parameter set would not be used if for this item (a) the stimulus E supports the E response less than it supports the U response, and (b) at the same time the context M_N supports the E response less than it supports the U response.

2. The expected error rate for all conditions could not be smaller than 1%. This constraint is realistic in the experiment we describe later, which produced an overall error rate of 22%.

The grid search was followed by Powell's derivative-free optimization algorithm (Press, Flannery, Teukolsky, & Vetterling, 1989) using the best condition found by the grid search as its starting point. The result of this extensive search was that the maximum average word superiority effect achievable by the factorized model was 2.26%. Thus, word superiority effects stronger than this are inconsistent with the factorized model.

Method and Materials

Using a computer program, we searched for cases in the Kucera and Francis (1967) corpus for sets of three-letter words organized in a conjunctive manner. The condition was that, for the three positions, switching one letter with the letter in the same position of another word in the set should produce a nonword. For example, in the set MEN, RED, MUD, RUN, letter switching generates a matched set of nonwords MUN, RUD, MED, REN. We chose 10 sets of words that had this conjunctive property and that maximized median frequency of word occurrence while minimizing the difference in frequency of occurrence within the words in the same set. The 10 sets of words, matched nonwords, and the frequency of occurrence of each in Kucera and Francis (1967) is displayed in Table 2. We treated membership in the lexicon as a discrete binary variable: Words are letter strings that occur in the Kucera-Francis word list, and all other letter strings are treated as nonwords. However, in our stimuli some of the "nonwords" are, in fact, low-frequency words, and others, such as MED, are used as abbreviations. It is likely that such items act as words at least to some individuals, but they are generally much lower in frequency than the stimuli categorized as words (as attested by their frequency of 0 in the Kucera-Francis word count).

The 10 sets of words have the characteristic that any of the three letter positions can serve as target in each set. For instance, if the letter to be reported is the last letter instead of the middle letter, the conjunctive effect is still there. In the MEN-RED-MUD-RUN example, middle E favors N when the first letter is M, but it favors D when the first letter is R. Therefore, from 10 sets of eight three-letter stimuli, there are 240 possible test items, in which each three-letter stimulus occurs three times, with each letter serving as the target once. The stimuli were presented using a

486-Dx33-based computer with a 15-in. NEC4FGx color monitor running at 72-Hz refresh rate and standard VGA resolution (640 × 480 pixels, 16 colors). The display parameters were adjusted using a pilot study so that the average error rate was about 25%.

Participants. The participants were 30 UCSD undergraduate students. They received partial course credit in introductory cognitive science courses for participating in the experiment. All were tested individually.

Procedure. Each participant received the following instructions: "You will be presented letters for brief periods of time. Your goal is to identify these letters. There will be 24 practice trials followed by 240 testing trials. Press the space bar to get a new set of letters. Press the up or down arrow key to choose the correct alternative; 50% of the trials show letters that make a word; 50% of the trials show letters that do not make a word. The experiment is self-paced. You can take as much time as you want between stimulus presentations."

After the instructions, each participant received four demonstration trials and was given the opportunity to ask questions. After answering any questions, the experimenter left the room. The demonstration trials were followed by 24 practice trials chosen from an additional set of stimuli organized in a conjunctive fashion. After 24 practice trials, each participant received the 240 test trials. The sequence of test trials was randomized. Each trial was initiated by the participant pressing the space bar key on the computer keyboard, which triggered presentation of a fixation point for 1 s. The fixation point was immediately followed by the three-letter display, which in turn was followed by a mask made of three # characters. The stimulus-onset asynchrony between the stimulus display and the mask was 28 ms (two refresh cycles). After 200 ms of mask presentation, a probe appeared indicating the target position and the two possible response alternatives.

Each participant was run under one of three possible feedback conditions. These conditions were included to explore the generality of our results across situations that could potentially create different response biases. In feedback Condition I, the participant received no feedback. In feedback Condition II, the response was followed for 1 s with a presentation of the correct letter in the appropriate position. In addition, a beep occurred if the incorrect choice had been made. Feedback Condition III was the same except that now both the stimulus and context were displayed after the response was made. Participants in Conditions II and III were told at the outset that after the response the computer would display the correct alternative and beep if the response was incorrect.

Design. There were three independent factors: word versus nonword (W); position (P) of the tested letter (first, second, third); and feedback condition (F). Factors W and P were within subjects, and F was between subjects. The subject's identity (S) was treated as a random factor. The dependent variable was the probability of correct response averaged across the 10 stimulus sets. An additional analysis treating stimulus set as a factor to check on the consistency of the results across stimuli is described later.

Table 2
Conjunctive Word Sets and Frequency of Occurrence

Words								Nonwords							
had	5,133	war	464	her	3,038	wed	2	har	0	wad	0	hed	0	wer	0
aid	130	oil	93	all	3,001	old	660	ail	0	oid	0	ald	0	oll	0
few	601	hex	4	fox	13	how	834	fex	0	hew	0	fow	0	hox	0
men	763	red	197	mud	32	run	212	med	0	ren	0	mun	0	rud	0
age	227	ugh	1	ash	11	use	589	agh	0	uge	0	ase	0	ush	0
nae	1	war	464	nor	195	woe	5	nar	0	wae	0	noe	0	wor	0
bad	142	gap	17	bop	3	god	318	bap	0	gad	0	bod	0	gop	0
gal	5	sat	150	got	482	sol	3	gat	0	sai	0	gol	0	sot	0
dry	68	gre	1	due	142	guy	51	dre	0	gry	0	duy	0	gue	0
sam	79	tar	12	sir	95	tim	25	sar	0	tam	0	sim	0	tir	0

Results

The ANOVA indicated a significant word superiority effect. No other main effect or interaction was significant. Table 3 shows the analysis results using Keppel's (1993) standard notation.

The difference in accuracy between the word and nonword conditions was 4.472%,

$$P(\text{correct}|\text{word}) = 0.799, \quad P(\text{correct}|\text{nonword}) = 0.754,$$

$$F(1, 27) = 29.82, p < 0.001. \quad (37)$$

The word superiority effect is somewhat smaller than the difference found in the only other experiment we know of in which only one of the forced-choice alternatives makes a word with the context (Estes, 1975). The relatively small difference in our case may be explained in part by the fact our stimuli are only three letters and our "nonwords" were all pronounceable pseudo-words, whereas those used by Estes were not necessarily word-like.

An additional analysis using the 10 word sets as an additional factor showed a significant effect of the stimulus set, $F(9, 243) = 3.29, p < .01$, but no Set \times Word interaction, $F(9, 243) = 1.14, p \geq .05$. Thus, some sets were easier than others over all, but there is no evidence that the word advantage differed across sets.

The obtained word superiority effect was significantly larger than 2.26%, the upper bound for factorized models: $t(27) = (4.47 - 2.26)/\sqrt{(2)(30.17)/90}, p < 0.01$, where the pooled estimate of the standard deviation of the effect is derived from the $w \times s/f$ error term of the ANOVA. Thus, simple factorized models can be rejected, indicating a violation of the Morton-Massaro law. Equivalent results were obtained using the multivariate ANOVA approach to within subjects analysis.

Discussion

It should be noted that, as in the prototypical experiments described in previous sections, our experiment does not distinguish whether the effect of context is perceptual or postperceptual. In particular, it is possible that participants' perceptions of the stimulus were not affected by context but their choices were. The

quantitative formulations of information integration that are considered here (e.g., the logogen model and the FLMP) were developed to account for data obtained in experiments like ours, in which it is not possible to distinguish between perceptual and postperceptual processes. What is an issue in our experiment is the nature of the proper characterization of the mappings between information sources and responses, not whether these mappings were perceptual or postperceptual. In some ways, the experiment may seem an obvious one in that it predicts a word advantage over nonwords in a situation in which a response bias could account for the effect. Many experiments have already reported such findings (Estes, 1975). What makes this experiment different is that it rules out the possibility that the word advantage is due to guessing based on pairwise letter constraints and requires that it be based on whole-word constraints.

A model with three autonomous modules, each selectively evaluating the support of a letter position for each of the different response categories is in direct contradiction with our experimental results. However, our results are not incompatible with all forms of factorization. In particular, the effect of the stimulus may factorize from the effect of context as a whole. For example, our results are compatible with an architecture that uses a module for the stimulus and a separate module for the context as a whole. In fact, this is the approach used by Massaro and Cohen (1991) to account for the word superiority effect (they used a version of the FLMP with factors selectively influenced by the stimulus and by the context as a whole). This points out an important issue: The pattern of results predicted by the Morton-Massaro law depends on how we choose to group our information sources. In our experiment, the Morton-Massaro law does not hold if each letter is considered as an experimental factor, but the law still holds if the two context letters are jointly considered as a single factor. Unfortunately, such a version of the Morton-Massaro law does not tell us how the two sources that make up the context are combined.

Thus, our results indicate the need for further development of models such as the current version of the FLMP. It may be useful to consider what such a model might look like, because this will help bring a second important lesson of the experiment. To begin, let us imagine a case in which the displays are three letters as in our experiment, but only one display position (e.g., the first) is designated as the target position; the other letter positions serve only to provide context. In this case, we could construct an FLMP model to account for the data for first position target letters. To do so, we would need a pathway for evaluation of the support for each alternative provided by the features of the target letter, and a separate pathway for evaluation of the support for each alternative provided by the context. This latter pathway would have to combine the two context letters conjunctively, while maintaining mutual independence with the pathway that processes the features of the target letter; the outputs of these two pathways would then be combined at the integration stage, in accordance with the FLMP equation.

This approach offers a straightforward method for accounting for the results obtained when the target position is known in advance. However, it does not provide a complete account of performance in our actual experiment, because the participants did not know which position would contain the target until after the stimulus had been followed by a mask. The task was sometimes to identify the first letter, sometimes the middle letter, and sometimes

Table 3
Analysis of Variance Table of Conjunctive Context Experiment

Source	df	M	F	p
w	1	900.0347	29.829	<.001
w \times s/f	27	30.1736		
p	2	39.6181	0.478	$\geq .05$
p \times s/f	54	82.9398		
w \times p	2	30.4514	1.269	$\geq .05$
w \times p \times s/f	54	24.0046		
f	2	313.3681	0.333	$\geq .05$
s/f	27	942.2569		
w \times f	2	22.3264	0.740	$\geq .05$
w \times s/f	27	30.1736		
p \times f	4	35.5035	0.428	$\geq .05$
p \times s/f	54	82.9398		
w \times p \times f	4	16.9618	0.707	$\geq .05$
w \times p \times s/f	54	24.0046		

the last letter, and participants were not informed in advance of each trial which letter position would be tested. To deal with the possibility of targets in the second position, the stimulus and context pathways would have to be configured differently, and a third configuration would be needed for third position targets. To cope with the fact that participants have no prior information about which letter position will be tested, it would appear that the three different configurations would have to be used in parallel on each trial; otherwise, the perceptual data would have to be buffered for postprocessing after the alternatives have been interpreted and the appropriate network chosen or configured.

Massaro et al. did not take a stand on these issues, and for us this reinforces our suggestion that the FLMP model (even as elaborated by Massaro & Cohen, 1991, to address the word superiority effect) may be viewed as a descriptive and functional framework for capturing aspects of experimental data (see Jacobs & Grainger, 1994, for a similar way of viewing the FLMP). However, Massaro clearly viewed the FLMP as describing structural aspects of processing itself (e.g., its feed-forward nature) not just function. Our suggestion is that the descriptive, functional, and structural aspects of the FLMP should be considered separately. We suggest that the FLMP equations provide a good description of certain aspects of the experimental facts, that the fuzzy logic framework provides a reasonable framework for the functionality of the equations, but that it is necessary to perform further analyses to understand whether they provide guidance about the structure of the perceptual system.

There are models that offer detailed mechanisms explaining how information from individual context letters is combined. Our results are incompatible with models that only keep information about pairwise letter statistics, like Golden's (1986) word recognition model. However, it is likely that our findings are compatible with models that use information about frequencies of higher order conjunctions of letters in English, not just pairwise letter statistics. These models include Rumelhart and Siple's (1974) model, McClelland and Rumelhart's (1981) interactive activation model, Paap, Newsome, McDonald, and Schvaneveldt's (1982) activation verification model, and Grainger and Jacobs' (1994) dual read-out model. In all of these models, the co-occurrences of three or four letters together to form words are used in determining the likelihood of a letter in a given context. Thus, these models take conjunctions of context letters into account, thereby making it likely that they could account for the conjunctive context effect.

Although these models may account for the conjunctive context effect, it is unclear whether they can also account for Massaro and Cohen's (1991) findings (i.e., that the Morton-Massaro law holds when the sources of information are grouped into stimulus and context as a whole). In the next section, we consider this issue for the case of interactive models of perception. In particular, we show that a stochastic version of McClelland and Rumelhart's (1981) interactive activation model can (a) account for the conjunctive context effect and (b) account for the fact that the effects of the target letter and of the context as a whole factorize.

Structural Analysis

The Morton-Massaro law is a statement about an empirical regularity: the fact that the joint effects of information sources on response probabilities often factorize. As such, it can be used as a

descriptive characterization of data without reference to underlying processing mechanisms. The optimality theorem helps us understand the possible rational basis of the law, but it does not tell us anything about internal processing mechanisms. Are there implications of the Morton-Massaro law for the structure of the underlying processing mechanisms? In this section we consider such questions.

In particular, we address the potential incompatibility of the Morton-Massaro law with interactive models of perception. As mentioned, the Morton-Massaro law tells us that the joint support of information sources for different response alternatives can be factorized into terms selectively influenced by each information source. This fact is consistent with modular approaches, like the FLMP or Morton's logogen model, in which information processing is decomposed into autonomous modules. However, it is not clear whether the law is consistent with interactive architectures in which feedback and lateral connections allow every part of the system to be affected by every other part. Massaro (1989c) noted that the probabilities of responses generated by the original version of the interactive activation model did not factorize, and he conjectured that this was an inescapable consequence of the interactive character of the model. He thus attempted to conclude that adherence to the law proved that the mechanism is feed-forward.

McClelland (1991) refuted Massaro's conjecture by showing that stochastic versions of the interactive activation model can be factorized if their connectivity is constrained in certain ways. The analysis presented there assumed networks of units with discrete binary states and local representations. Simulations suggested that interactive networks with continuous states may also exhibit factorability, but no mathematical analysis was presented for the continuous case. Massaro and Cohen (1991) reported simulations in which they could not find good fits to actual empirical data using the stochastic interactive activation model, but Movellan and McClelland (1995) showed that their failure was due to inappropriate limitations on parameters. By removing these constraints they obtained excellent fits.

In this section, we expand the mathematical analysis of interactive models to networks that work in continuous time with continuous states. These networks relate to diffusion models, which have a rich tradition in the human information-processing literature (Ashby, 1989; Ratcliff, 1978; Townsend & Ashby, 1983) and take us a step closer to understanding how perceptual models may be implemented in electric circuits and neural hardware.

First, we formally introduce the class of models we are concerned with and present the concept of stochastic equilibrium, which plays a central role in the analysis. Second, we introduce the concept of channel separability and show its relationship to the Morton-Massaro law.

Diffusion Networks

McClelland (1993) emphasized the principles of graded, random, and interactive propagation of information to explain human cognition and to bridge gaps between the study of the mind and the brain. Following that idea here, we focus our analysis on network models that are interactive (i.e., in the general case they may have lateral and feedback connections), graded (i.e., operate in continuous time with continuous representational states), and intrinsically stochastic (i.e., they are defined by stochastic differential

equations instead of ordinary differential equations). We call these models diffusion networks, because they are stochastic diffusion processes defined by adding Brownian motion to the standard connectionist network dynamics (Mineiro et al., 1998; Movellan, 1994, 1998; Movellan & McClelland, 1993). Diffusion networks are very similar to the continuous stochastic interactive activation networks presented in McClelland (1991). We modified the networks to give them a useful physical interpretation as models of low-power electronic circuits and to facilitate their analysis using tools from the theory of stochastic differential equations and diffusion processes (Karatzas & Shreve, 1988; Kloeden & Platen, 1992; Oksendal, 1992).

Diffusion networks consist of a set of n units whose activation values change through time because of influences from other units and from external inputs. In addition, the network dynamics are probabilistic because of an intrinsic noise component. Diffusion networks are defined by stochastic differential equations of the following form:⁷

$$dY_i(t) = \mu_i(Y(t), X)dt + \sigma dB_i(t), \quad \text{for } i \in \{1, \dots, n\}, \quad (38)$$

where $Y_i(t)$ is a random variable representing the internal state at time t of the i th unit, $Y(t) = (Y_1(t), \dots, Y_n(t))'$ is the vector of activation states at time t , $X = (X_1, \dots, X_m)'$ is the external input, and B_i is Brownian motion, a mathematical model of the random noise present in many natural systems. The constant $\sigma > 0$, known as the dispersion, controls the amount of noise injected into each unit. The function μ_i , known as the drift, determines the average instantaneous change of activation. The drift is modulated by a matrix w of connections between units, and a matrix v that controls the influence of the external inputs onto each unit,

$$\mu_i(Y(t), X) = \frac{1}{\kappa_i(Y_i(t))} \left(\bar{Y}_i(t) - \frac{1}{\rho_i} Y_i(t) \right) - \gamma Y_i(t), \quad (39)$$

for $i \in \{1, \dots, n\}$,

where κ_i is a positive function that will be defined in Equation 43, $\rho_i > 0$ is a constant parameter,

$$Y_i(t) = \sum_j w_{ij} Z_j(t) + \sum_k v_{ik} X_k \quad \text{for } i \in \{1, \dots, n\}, \quad (40)$$

is the net input of unit i , w_{ij} is the weight from unit j into unit i , v_{ik} is a parameter modulating the effect of input X_k on unit i , and

$$Z_j(t) = \varphi_j(Y_j(t)) = \varphi(\alpha_j Y_j(t)), \quad (41)$$

is the activation of unit j . The term φ represents the activation function, which we assume to be differentiable and bounded. In practice, we use the logistic function, but other options are also possible:

$$Z_j(t) = \varphi(\alpha_j Y_j(t)) = \frac{1}{1 + e^{\alpha_j Y_j(t)}}. \quad (42)$$

The term $\gamma > 0$ is a decay parameter, and the $\alpha_j \geq 0$ terms are gain parameters that control the sharpness of the activation functions. For large values of α_j , the activation function of unit j converges to a step function (Figure 10). Finally, the function κ_i is defined as follows:




Figure 10. Logistic activation function for several values of the gain parameter. As the gain goes to infinity, the logistic function converges to a step function.

$$\frac{1}{\kappa_i(Y(t))} = \frac{d\varphi_i(Y_i(t))}{dY_i(t)}. \quad (43)$$

This particular form of κ_i was chosen to facilitate the mathematical analysis of the equilibrium behavior of these networks (Movellan, 1998). In practice, if the activation function is logistic,

$$\frac{1}{\kappa_i(Y(t))} = \frac{d\varphi_i(Y_i(t))}{dY_i(t)} = \alpha_i \varphi(Y_i(t))(1 - \varphi(Y_i(t))). \quad (44)$$

Intuition for Equation 39 can be achieved by thinking of it as a the limit of a discrete time difference equation, in such case

$$Y_i(t + \Delta t) = Y_i(t) + \mu_i(Y(t), X)\Delta t + \sigma \sqrt{\Delta t} N_i(t), \quad (45)$$

where the $N_i(t)$ are independent standard Gaussian random variables. For a fixed state at time t , there are two forces controlling the change of state: the drift, which is deterministic, and the dispersion, which is stochastic (Figure 11). This results in a distribution of states at time $t + \Delta t$. As Δt goes to zero, the solution to the difference Equation 45 converges in distribution to the diffusion process defined in Equation 38.

Diffusion networks can be seen as a generalization of the linear diffusion model that has been used by Ratcliff's (1978) to model reaction time in memory experiments and other paradigms. In Ratcliff's model, the drift is constant within an experimental trial, and the units are not coupled. In our diffusion networks, the units are coupled by connection weights, and the drift is a function of the state of activation. Diffusion networks are closely related to the coupled diffusion processes investigated by Ashby (in press), but our approach is more general because we allow nonlinear activation functions and feedback, not just lateral connections between processing units. Diffusion networks have a physical interpretation as models of low-power electronic circuits in which the random motion of the electrons plays a nonnegligible role. In addition, diffusion networks are reasonable models of biological neural

⁷ We do not include the initial distribution because it is irrelevant for the equilibrium solution, which is the focus of our article.




Figure 11. The drift and the dispersion combine to produce a distribution of states at time $t + dt$, given the state at time t .

networks with an intrinsic noise component. In such case, each unit represents a point neuron, the weight values represent synaptic conductances, κ_i is a capacitance function, ρ_i a transmembrane resistance, the internal states Y are presynaptic potentials, the elements of \bar{Y} are net input currents, and the elements of Z are postsynaptic short-term firing rates (see Hopfield, 1984, p. 3089).

Stochastic Equilibrium

Figure 12 shows an example of a diffusion network with two units responding to a transient input. The two units were coupled to each other by bidirectional excitatory connections, the external input had an excitatory effect on Unit 1. The “noisy curves” show the evolution of activation in a particular run of the network. Consider an ensemble of identical clones of the same diffusion network, what is sometimes known as a canonical ensemble. We can think of members of the canonical ensemble as replicas of the same network running in parallel or as distinct runs of the same network running at different times. For each unit and for each time step, we could compute expected activation values by averaging across the entire ensemble. Figure 12 displays the evolution of such averages (the two smooth curves). Note how some time after

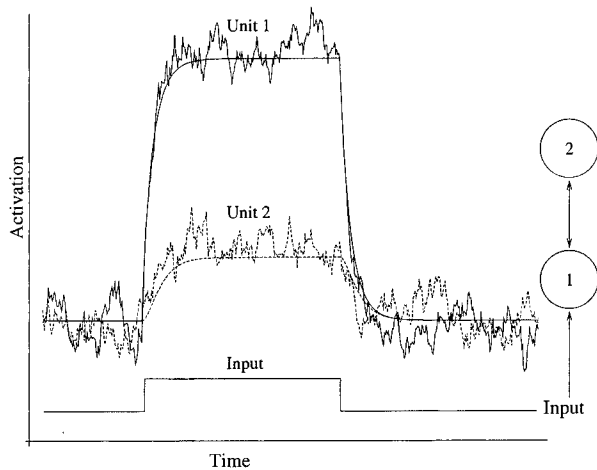


Figure 12. Simulation of a two-unit diffusion network, illustrated at right. The input was a rectangular signal (bottom function). A sample activation path, one per unit in the network, and the average activation path are shown.

the presentation of the stimulus these averages equilibrate. We say that a network is at stochastic equilibrium when the probability distribution of states across the ensemble equilibrates; that is, at stochastic equilibrium all the ensemble statistics (e.g., mean, variance, covariance, kurtosis) equilibrate. Note that because of the effects of noise the activations obtained in individual runs never equilibrate. What equilibrates is the distribution of states across an entire ensemble of runs. Figure 13 shows a simulation of a single-unit diffusion network. The network was run 50 million times with internal states starting at 0. The activation histogram was collected after 10, 40, and 600 cycles. Figure 13 displays the theoretical equilibrium distribution, which was almost perfectly achieved after 600 cycles. Simulation details are in the Appendix.

A system like this can be used to do statistical inference in the following way: An input x is introduced that disturbs the established equilibrium. Eventually, a new equilibrium distribution is reached. This stable distribution represents the response of the system to the input x .

Consider now an experiment involving repeated trials, each displaying particular combinations of information sources. Each such trial with a particular combination of information sources can be thought of as another sample from the canonical ensemble. On each trial, we let the network run for a period of time sufficient to approximate stochastic equilibrium and count the proportion of trials that the state of the network is in a particular region of the state space. As the number of trials increases, this proportion converges into the equilibrium probability of being in the region. Each region is associated with one of the possible responses, so that the probability of being in a region corresponds to the probability of making a particular response. This is precisely the approach we use in the next section to analyze the Morton-Massaro law in diffusion networks.

Channel Separability

To simplify the analysis, we assume the network connections are symmetric (i.e., the connection $w_{j,i}$ from unit j to unit i equals the connection $w_{i,j}$ from unit i to unit j). Later we discuss more general cases in which this assumption is relaxed. Symmetric




Figure 13. The empirical activation histograms obtained by simulating a one-unit network 50 million times. Histograms were collected after 10, 40, and 600 cycles. The theoretical equilibrium distribution is displayed with a continuous line. Simulation details are in the Appendix.

weights allow for each state of the network to have an associated goodness value. Here goodness represents how well the activation of the units in the network conforms to the constraints imposed by the external inputs and the network architecture. The goodness function has the following form

$$G(z|x) = \frac{1}{2} \sum_i \sum_j z_i w_{ij} z_j + \sum_i \sum_k z_i v_{ik} x_k - \sum_i S_i(z_i), \quad (46)$$

where $G(z|x)$ is the goodness of activation state z given the external input x . The state of activation z is an n -dimensional vector, z_i is the i th element of z (i.e., the activation of the i th unit), x is an input vector containing the state of the external sources of information, and x_j is the j th component of that vector. The expression $\sum_i \sum_j z_i w_{ij} z_j$ is a measure of the consistency of the pattern of activation z in the network with the matrix of connection weights w . One can see that if w_{ij} is positive, then goodness increases monotonically with either z_i or z_j as long as they have the same sign. When w_{ij} is negative, goodness increases monotonically with either z_i or z_j as long as they have opposite signs. The next part of the expression, $\sum_i \sum_k z_i v_{ik} x_k$ is a measure of the consistency of the pattern of activation z with the external input x . The first two terms in Equation 46 favor states with extreme values. The function S_i in the final term penalizes extreme activations of the units. The exact form of this function is found in Equation 87 of the Appendix, but its most important characteristic is that it rapidly increases as the activations approach extremes. Thus, overall, the goodness of a state represents a balance between a tendency to maximize activations of units with positive weights from other active elements (units or inputs) while avoiding extreme activations. A fuller discussion of goodness, and how networks tend to settle toward good states, may be found in Rumelhart et al. (1986). In Movellan (1998), it is shown that if the weights are symmetric, then at equilibrium the probability of an activation state z given an external input x is an exponential function of the goodness

$$p(z|x) = \frac{1}{K(x)} \exp((2/\sigma^2)G(z|x)), \quad (47)$$

where $K(x)$ is a term that makes probability densities integrate to 1. Note that the “better” a state is (i.e., the better it satisfies the network constraints), the more probable it is to visit it at equilibrium.

In the Appendix we show that for networks of this type there is a condition, which we named channel separability, that is sufficient to exhibit the Morton-Massaro law. Our formalization of channel separability entails separating units into the following sets:

1. Response specification units: A unit belongs to the response specification set, if, when the state of all the other units in the network is fixed, changing the state of this unit can affect the probability distribution of external responses.
2. Information channels: A unit belongs to the channel C_i of information source X_i if (a) it is not a response specification unit, and (b) when the state of the response units is fixed, the probability distribution of the activations of this unit can be affected by X_i .

Channel separability. We say that two channels are separable if they are disjoint (i.e., if they have no units in common).

Units within a channel can be sensors or relays:

1. Sensors: A unit is a sensor for source X_i if, when the state of all the other units in the network is fixed, the probability distribution of activations of this unit is affected by X_i . If a unit is a sensor for more than one source, we say that it is a mixed sensor.
2. Relays: A unit is a relay if it is not a sensor and it is not a response unit.

Channel separability requires that there be no mixed sensors and that there be no direct connectivity between relays belonging to different channels, as illustrated in Figure 14. Separability would be violated if, for example, the external stimulus is visual, the external context is auditory, and the auditory stimulation is so loud that it has a physical effect on the retina strong enough to influence neural activity there. In such an extreme case, the retina would be a mixed sensor, thus violating separability.

Separability implies that the influences of an information source on the channel of another information source should be mediated by the response specification units. Because the connections in the network may be bidirectional throughout, such indirect influences are possible. Direct connections between primary visual cortex and primary auditory cortex would violate channel separability. However, if responses are read out of a downstream cortical area, bidirectional connections linking the downstream area from which responses are read out to each of the two primary cortices would not violate the separability constraint.

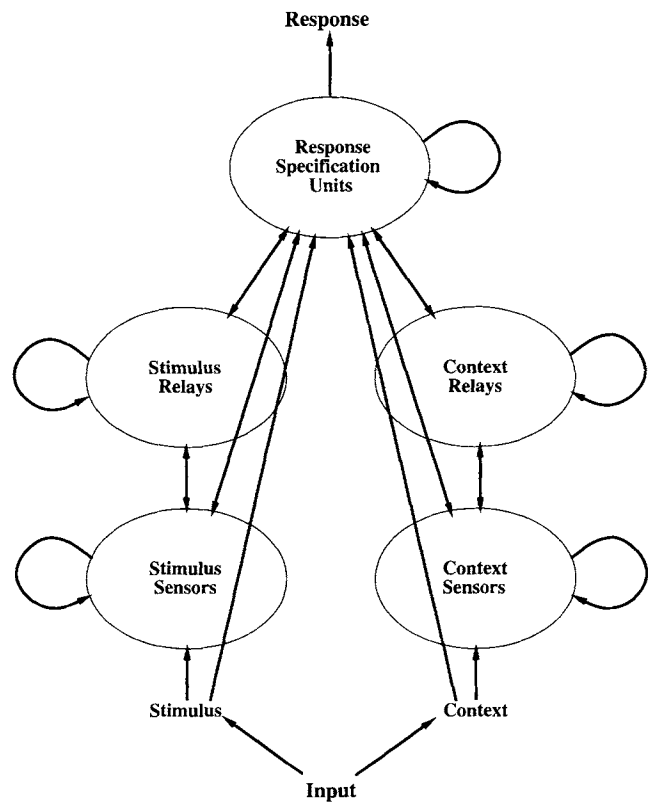


Figure 14. A network with separable context and stimulus processing channels. The stimulus sensor and stimulus relay units make up the stimulus channel units, and the context sensor and context relay units make up the context channel units.

To analyze the channels in a diffusion network, first we isolate the response specification units and disconnect them from the rest of the units in the network (i.e., we create a virtual network with all the connections in and out the response specification units set to 0). The channel C_i for information source X_i is the set of units in the virtual network whose activation is modulated by X_i . Two channels C_i and C_j are separable if they have no units in common.

To see how this works in practice, consider the simplified version of McClelland and Rumelhart's (1981) word reading model shown in Figure 15. Suppose the task is to identify the second letter of a three-letter word, and the external response is determined by the states of the letter units in the second location. Thus, there are two response specification units: the units labeled E and U (see Figure 15). The stimulus is X_2 , and the context is composed of two sources: X_1 and X_3 . When the response specification units are disconnected from the rest of the network, there are no units other than the response specification units, modulated by the stimulus X_2 . Thus, in this case the stimulus channel is empty, $C_2 = \emptyset$. When the response specification units are disconnected, the set of units affected by X_1 is as follows

$$C_1 = \{\text{MEN, MUD, RED, RUN, M, R, D, N}\}. \quad (48)$$

Finally, $C_3 = C_1$, because all the units affected by X_1 are also affected by X_3 . We can conclude that (a) the stimulus channel is separable from the context channel as a whole, because $C_2 \cap (C_1 \cup C_3) = \emptyset$; and (b) the two context channels are not separable, because $C_1 \cap C_3 \neq \emptyset$. Note how our definition of information channel depends on the choice of response specification units. If the external responses were based on the letter units in the left-most position, then the units assigned to each information channel would be different (i.e., we would get $C_1 = \emptyset$, and $C_2 = C_3 = \{\text{MEN, MUD, RED, RUN, E, U, D, N}\}$). We now state a theorem relating channel separability and response factorability in diffusion networks.

Separability theorem. Consider a diffusion network as defined in Equation 38 through Equation 43 with symmetric weights and external input $X = (X_1, \dots, X_m)'$. Let C_j be the channel for X_j , $j = 1, \dots, m$ as defined previously. If $C_i \cap C_j = \emptyset$ for $j = 1, \dots, m, j \neq i$, then the equilibrium density of responses




Figure 15. A simplified version of McClelland and Rumelhart's (1981) word reading model. Negative connections end with a circle. Not all the negative connections are displayed.

factorizes into a term controlled by X_i and a term jointly controlled by all the other inputs.

The theorem is proved in the Appendix. The theorem tells us that if we can find a source X_i with a channel separable from all the other sources, then response probability ratios factorize into a term controlled by source X_i and a term jointly controlled by all the other sources. Although the details of the proof are technical, the reason why separable networks factorize is easy to understand. Call X_i the stimulus and all the other input variables the context. Separability of the stimulus and context channels means that there are no direct connections between the stimulus and context units and thus there are no terms in the goodness function (Equation 46) that depend jointly on the stimulus and context. Because of this, the goodness can be separated into two additive terms: one that depends on the stimulus and one that depends on the context. Because probability densities are proportional to the exponential of the goodness (Equation 47), it follows that if the stimulus and context channels are separable, then the probability density of the response specification units factorizes into a term selectively influenced by the stimulus and a term selectively influenced by the context (see Appendix). We now study how the separability theorem can be used to analyze experiments with continuous and discrete responses.

Continuous Responses

In some cases, it is of interest to model continuous responses rather than discrete categorical judgments. This situation might arise, for example, in choosing a number or the position of a dial that represents the perceived size of an object. Such cases can be modeled with diffusion networks using one-to-one mappings between the states of a response unit and the external responses. The separability theorem tells us that if the weights are symmetric and the stimulus and context channels are separable, then response probability densities factorize into terms selectively controlled by the stimulus and by the context as a whole.

Discrete Responses

McClelland (1991) showed that, for networks of units that take on only discrete binary states, channel separability is sufficient to exhibit the Morton-Massaro law. The situation is more complex with diffusion networks because we need to map continuous states into discrete responses. To do so, we associate overt responses with entire regions of the response specification space rather than with individual points in that region. To get the probability of a particular response, we need to integrate the response probability density over the region associated with that response.

The problem is that integrals of probability densities do not necessarily factorize even though the densities factorize at every point. Figure 16 illustrates a case in which probability densities factorize, yet careful choice of response regions results in a violation of the Morton-Massaro law. Figure 16 shows the response distributions of a simple diffusion network to four different combinations of stimulus and context. It also shows a particular mapping from continuous activations into two discrete responses. This mapping produces a violation of factorability of the discrete response probabilities, even though the underlying distribution of continuous responses factorizes. Note how when Context = 1, the




Figure 16. The four figures show the distribution of activation states in response to four different combinations of stimulus and context. The distribution of continuous activation states factorizes. However, for the particular mapping between continuous states and discrete responses used in this figure, the distribution of discrete responses does not factorize.

effect of the stimulus is to increase the probability of Response 2. However, when Context = 0, the stimulus decreases the probability of Response 2. Because there are only two response alternatives, this crossover interaction indicates a clear violation of factorability (see *Modularity and Statistical Interaction*).

Fortunately, there are two important cases for which factorability holds, at least as a good approximation. The first case is when the response regions are relatively small, and thus we can approximate the integral over that region by the density at a point times the volume of the region. In such a case, the ratio of the integrals can be approximated by the ratio of the probability densities of those two points times the ratio of the two volumes. Because probability densities of individual states factorize, probability ratios of small regions should approximately factorize, with the approximation getting arbitrarily better as we make the response regions smaller.

The second case applies to models, like McClelland and Rumelhart's (1981) word reading model, in which each response is associated with a distinct response unit. These models typically have negative connections among the response units so that at equilibrium one unit tends to be active while the others are inactive. In such a case, a common response policy chooses the response corresponding to the most active unit. It turns out that when the response units inhibit each other, this response policy can be closely approximated by another policy, the active-wins policy, which is easier to analyze mathematically (Figure 17). In the active-wins approach, a response is chosen when a unit is active beyond a given threshold and all the other units are below threshold. A similar policy has been used by Grainger and Jacobs (1994, 1996) in several models of visual word identification. In the Appendix, we show that if the network has separable channels, such a policy can approximate the Morton-Massaro law to an arbitrary level of precision as the gain parameter of the response

units is increased (see Figure 10). In practice, we have observed that excellent approximations to the Morton-Massaro law can be achieved with small values of the gain parameters (see details of the simulation displayed in Figure 18).

Summary and Discussion

- In discrete networks with binary states, symmetric weights, separable channels, and an active-wins response rule, the response probabilities factorize (McClelland, 1991).
- In diffusion networks with symmetric weights and separable channels, response probability densities factorize.
- Diffusion networks with symmetric weights, separable channels, and an active-wins response rule can arbitrarily approximate factorability by increasing the sharpness of the response units.

The main lesson of the analysis is that feedback is not incompatible with the Morton-Massaro law. Reasons in favor of or against interactive models must be found outside of this law. Even though stimulus and context units are interdependent in diffusion networks with feedback connections, it is still possible to factorize their effect on response probabilities. The analysis presented in the Appendix takes us a long way toward understanding mathematically why this is the case.

The separability theorem can be applied to a wide range of parallel-distributed processing models. For example, consider the schematic structure of the word perception model shown in Figure 15. As we saw earlier, the stimulus channel C_2 is separable from the context channel as a whole, that is, $C_2 \cap (C_1 \cup C_3) = \emptyset$. Thus, in such a model, response probabilities should factorize into a term controlled by the stimulus and a term controlled by the context as a whole. However, note that the channel for each context letter is not separable from the other two channels; thus, it is possible for this model to produce a conjunctive context effect.




Figure 17. The equilibrium distribution of two response specification units. Because of inhibitory connections between these units, the activations are negatively correlated. The diagonal line separates the response regions produced by a response policy that chooses the most active unit. These regions can be closely approximated using two thresholds (the horizontal and vertical lines) and an active-wins response rule.

Moreover, the same property holds when the response units are in the first or third position, not just the second, so that responses based on activations of the letter units in any of the positions should factorize into a term controlled by the stimulus and another term controlled by the context as a whole.

To test this prediction, we implemented a diffusion network version of the word perception model shown in Figure 15. There were six stimulus conditions, corresponding to different levels of external input to the unit E in the second position. There were four context conditions: M_N, R_D, M_D, R_N. For each of these 24 conditions, 40,000 trials were simulated. On each trial, stimulus and context were presented and after a settling period, the activations of the E and U units were collected. The largest activation determined the external response. The weight parameters were set by hand to approximate roughly the results we obtained in our experiment on the conjunctive context effect (i.e., when the correct identification of letters in the context of nonwords was about 75%, the correct identification of letters in the context of words was about 80%). The specific parameters used in this simulation are presented in the Appendix. Figure 19 shows a single trial in which the network was presented with a stimulus ambiguous between E and U in the context of R_N. Figure 18 shows the simulation results over the entire ensemble of trials. The ordinate shows the probability of response E using a logit scale, that is, $\log(p/(1-p))$. As shown previously, factorized effects on probability ratios are additive on the logit transform of these probabilities. As predicted, the stimulus and context as a whole factorized. This can be seen in Figure 18 by the fact that all the lines are parallel. Moreover, the two context letters were processed in a conjunctive manner. This can be seen by the fact that the line corresponding to

the context R_N is above the line corresponding to the context R_D, whereas the line for the context M_N is below the line for the context M_D. Thus, in the model, the support of the letter presented in the third position for the letter in the second position depends on the letter presented in the first position.

As we previously noted, interactive models are not the only ones that can produce this pattern of results. For example, the FLMP can describe these results by assigning separate parameters to the stimulus and to all the possible context combinations. However, such a model would not provide a mechanism to explain how the two letters that make up the context are combined. Models consistent with the conjunctive context effect, such as Rumelhart and Siple's (1974) model, Paap et al.'s (1982) activation verification model, and Grainger and Jacobs's (1994) dual read-out model, still need to be evaluated for their consistency with factorability of stimulus and context as a whole. Other approaches certainly cannot be ruled out, but our work shows that a diffusion version of the interactive activation architecture nicely solves the problem of (a) satisfying factorability of stimulus and context as a whole while (b) exhibiting conjunctive context effects and (c) providing a mechanism explaining how such contextual conjunctions may emerge.

The optimality theorem assumes symmetric weights. The non-symmetric case is difficult to analyze mathematically; fortunately, however, results can be found in some special cases. For example, it is possible to show that in diffusion networks with separable channels, if the noise is low, or if the activation functions φ are linear, then response probability densities factorize regardless of weight symmetry (Movellan & McClelland, 1995). These results suggest that weight symmetry is not a key property for factorabil-

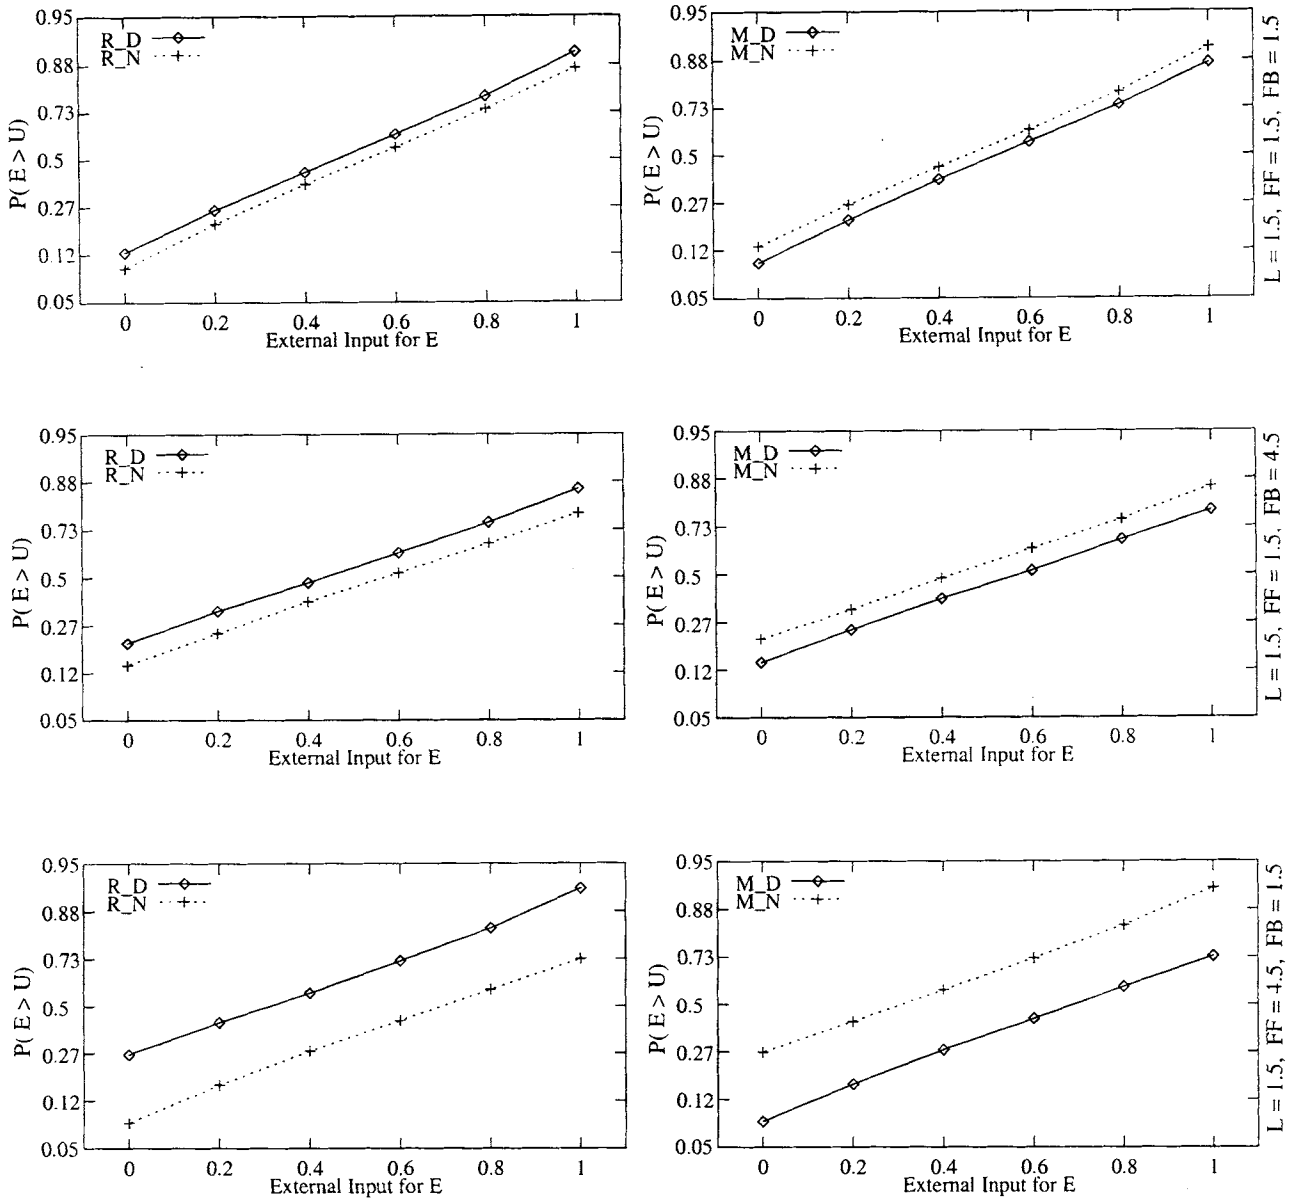


Figure 18. Simulation of the conjunctive context effect with the network displayed in Figure 15. The ordinate shows response probabilities using a logit scale. The first row shows results for a network with lateral weights (L), feed-forward weights (FF), and feedback weights (FB) set to 1.5. The second row shows results with the same network when the FB weights were set to 4.5. The last row shows results when the FF weights were set to 4.5. In all cases the network showed a conjunctive context effect while adhering well to the Morton-Massaro law for the stimulus and context as a whole.

ity to hold. Figure 18 shows simulations in which the feed-forward connections were three times larger or three times smaller than the feedback connections. The effect of this weight asymmetry was to change the intercept of the lines, but the overall pattern of results was the same regardless of weight asymmetry: The stimulus factorized from the context as a whole, but the two context letters combined in a conjunctive manner.

When the activation values are interpreted as spiking rates, diffusion networks are reasonable models of physical neural net-

works. Thus, the analysis presented here makes novel predictions at the neural systems level. For example, it is now known that the response of neurons in primary visual cortex is modulated by the background stimuli outside their classical receptive fields. These contextual modulations are due to the effect of lateral and feedback connections (Figure 20). Suppose we present a neuron in primary visual cortex with combinations of inputs inside its classical receptive field (the stimulus) and outside its classical receptive field (the context). For each combination of stimulus and context, we

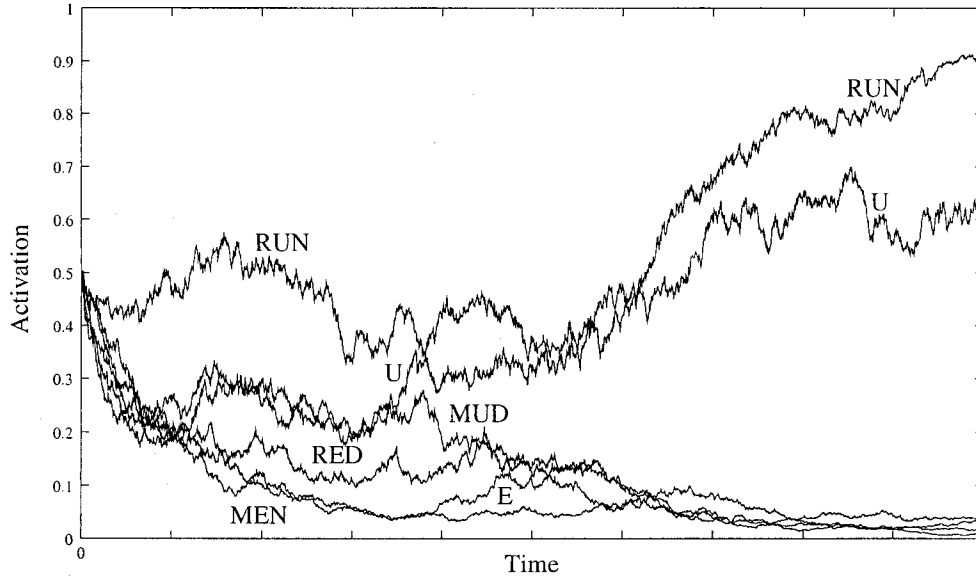


Figure 19. A simulation trial of the diffusion network displayed in Figure 15. In this trial an ambiguous stimulus in the second position was presented in the context of R_N.

can collect the distribution of spike rates of the neuron under consideration. As Figure 20 makes clear, the stimulus channel is separable from the background channel as a whole. Thus, the prediction is made that in this case the Morton-Massaro law should hold (i.e., at equilibrium the spike rate histograms should factorize into terms selectively controlled by the stimulus and the background). The analysis presented here may also be used to infer the

functional connectivity of the brain. We may excite two different regions and observe the activity of a third region. If the observed distribution of spike rates in the third region does not follow the Morton-Massaro law, then it is likely that the channels from the two input areas to the response area are not separable.

Relation to Other Notions of Separability and Independence

The notion of perceptual independence has played a central role in the cognitive psychology literature. Unfortunately, this concept has many interpretations leading to a large and confusing proliferation of terms, for example, orthogonal channels (Green, Weber, & Duncan, 1977), stimulus separability (Shepard, 1964), sampling independence (Townsend, Hu, & Ashby, 1981), response parity (Garner & Morton, 1969), and uncorrelated pathways, among others.

Ashby and Townsend's (1986) general recognition theory (GRT) is arguably the most thorough and successful attempt to formalize and unify the different notions of independence present in the psychological literature. In this section, we develop a framework inspired on an interpretation of GRT from the point of view of information theory. This framework helps clarify that our notion of channel separability in diffusion networks is unrelated to previous notions of independence and separability in the psychological literature. To begin the analysis, it is helpful to differentiate four types of variables relevant to the description of perceptual processes:

1. External sources of information: (X_1, \dots, X_m) , the external inputs.
2. Hidden variables: (H_1, \dots, H_m) , preperceptual variables internal to the subject that are influenced by the external sources and that have an influence on the perceptual experience (e.g., the state of the ganglion cells in the retina). Here H_i is a vector of hidden variables influenced by X_i .

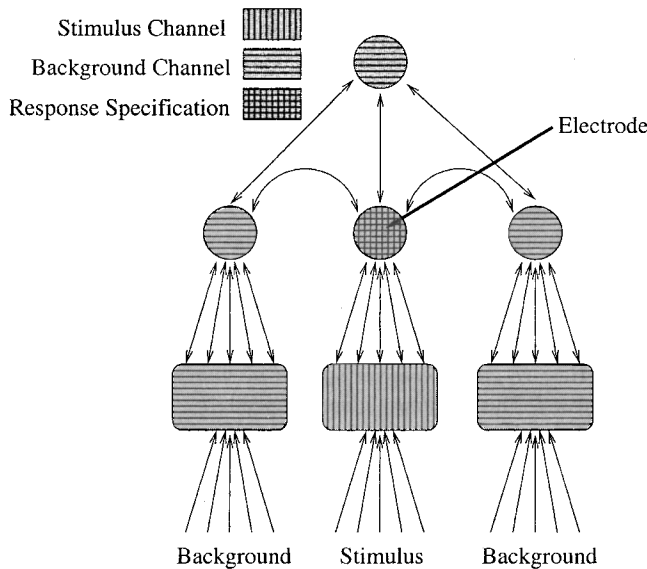


Figure 20. The arrows connecting the stimulus to the unit in the center represent the classical receptive field of that unit. External inputs affecting the classical receptive field are called "stimuli," and all the other inputs are called "background." The center unit determines the responses under consideration (e.g., spiking rates of that neuron). In this preparation the stimulus and background channels are separable.

3. Perceptual variables: (P_1, \dots, P_m) , internal variables that represent the perceptual representations of sources X_1, \dots, X_m .
4. Output variables: (O_1, \dots, O_m) , variables that represent observable responses to inquiries about the sources X_1, \dots, X_m .

Different models emphasize different variables. For example, in our analysis of diffusion networks, we gloss over the difference between internal perceptions and overt responses (i.e., because the focus of the Morton-Massaro law is the relationship between external inputs and responses, we treat perceptual variables as relays between inputs and responses). Other models, like the GRT, focus on the relationship between inputs, perceptions, and responses without using hidden variables (i.e., preperceptual states are ignored).

Most terms found in the perceptual independence literature can be reduced to two different notions of stochastic independence applied at different levels of the perceptual process. We call these two notions classic independence and classic separability.

Classic perceptual independence (Ashby & Townsend, 1986) is the lack of contingency between the perceptual interpretation of two sources once we partial out the effect of these sources. For example, we would say that X_i and X_j are perceptually independent if

$$I(P_i, P_j | X_i, X_j) = 0, \quad (49)$$

where $I(P_i, P_j | X_i, X_j)$ stands for mutual information between the perceptual representation of X_i and X_j , when the effects of these two external sources are partialled out.⁸ For a better understanding of this notion, consider the following simple linear model in which

$$P_i = \sum_j \alpha_{i,j} X_j + \beta_{i,j} N_j, \quad (50)$$

where N_1, \dots, N_m are noise components independent of the external sources and of each other. The $\alpha_{i,j} \geq 0$ terms represent the effect of input source X_j on the perceptual interpretation of X_i . The $\beta_{i,j} \geq 0$ terms represent the effect of noise source N_j on the perceptual representation of X_i . In this case, perceptual independence would be satisfied if there is no cross-talk between the noise terms (i.e., $\beta_{i,j} = 0$ when $i \neq j$). Note how it is possible for X_i and X_j to be perceptually independent even though X_j may have an effect on the perception of X_i by virtue of having $\alpha_{i,j} > 0$.

The same mathematical notion of independence can be applied between external inputs and outputs

$$I(O_i, O_j | X_i, X_j) = 0, \quad (51)$$

and appears in the literature with names such as "sampling independence" (Townsend et al., 1981, p. 159).

The classic notion of separability relates to a different form of stochastic independence. This notion formalizes the idea that processes are selectively influenced by different sources. Garner and Morton (1969) were first to formalize this notion mathematically as follows

$$I(O_i, \bar{X}_i | X_i) = 0, \quad (52)$$

where \bar{X}_i represents all the external sources with the exclusion of X_i . This simply says that once we know the source X_i , then the other sources give no further information about the response O_i .

The same notion of separability can be applied at other levels of the perceptual process. For example, when applied between external inputs and perception,

$$I(P_i, \bar{X}_i | X_i) = 0, \quad (53)$$

it appears with the name "perceptual separability" (Ashby & Townsend, 1986). In the simple linear model (Equation 50), this notion of separability would be satisfied if there are no cross-terms between the external sources (i.e., $\alpha_{i,j} = 0$ when $i \neq j$). Note how X_i and X_k may be perceptually separable even though their perceptual representations may not be independent by virtue of having $\beta_{i,j} > 0$ for $i \neq j$.

When this notion is applied between perceptions and responses

$$I(O_i, \bar{P}_i | P_i) = 0, \quad (54)$$

it is known as "response separability" (Ashby & Townsend, 1986). Here \bar{P}_i represents the perceptual representations of all external sources, except for X_i . When applied between external inputs and hidden preperceptual processes such as the "the output of auditory detectors"

$$I(H_i, \bar{X}_i | X_i) = 0, \quad (55)$$

it receives names such as "independent detection" and "orthogonality" (Green et al., 1977).

Now that we have a framework in which to classify notions of independence and separability, we will address where our notion of channel separability falls within this framework.

First, note that in the experimental paradigm of the Morton-Massaro law we assume that more than one source of information will have an influence on the external responses. What is at stake in these experiments is whether the external sources influence the external responses in a manner consistent with the Morton-Massaro law. Morton pointed out that "in a logogen, evidence relevant to a particular response is collected indifferently with respect to the source of information" (Garner & Morton, 1969). Massaro emphasized that a central aspect of the FLMP is the fact that different sources of information are "independently" evaluated before their influences are integrated. Morton's idea of indifferent evaluation and Massaro's notion of independent evaluation is, in fact, a case of classic separability that can be formulated as follows

$$I(H_i, \bar{X}_i | X_i) = 0, \quad (56)$$

where $H_i = (H_{i,1}, \dots, H_{i,m})$ is a vector of preperceptual evaluation terms. Here $H_{i,j}$ represents the support of source X_i for response alternative j . This is simply classic separability applied between external sources and preperceptual terms controlled by these sources (e.g., the evaluation values in the FLMP). Both the logogen model and the FLMP exhibit this classic notion of separability.

Ashby (in press) developed a continuous time version of the GRT, which is closely related to our work. The model is defined by a stochastic differential equation similar to the ones used in diffusion networks. The main difference between his work and

⁸ Mutual information is a generalized measure of contingency between random variables. Two random variables are independent if and only if their mutual information is zero.

ours is that we allow feedback connections and nonlinear activation functions, whereas the networks he studied are linear and do not have feedback connections. Ashby (in press) showed how the classic notions of separability and independence relate to architectural constraints in feed-forward linear networks. The situation is more complex in our work because of the presence of nonlinearities and feedback connections. As far as we can tell, diffusion networks with separable channels and feedback connections do not exhibit any of the classic notions of independence and separability presented in this section, but yet they exhibit the Morton-Massaro law.

Summary and Conclusions

First, we summarize the main arguments and implications of this article in the form of a set of questions and answers:⁹

1. What is the Morton-Massaro law? The Morton-Massaro law is a frequently observed relationship between inputs and responses in tasks that involve the integration of multiple sources of information. A prototypical example is the combination of acoustic and visual sources in speech perception.
2. What behavior defines the Morton-Massaro law? Ratios of response probabilities factorize into components selectively influenced by a single information source.
3. Is the law determined by the environment? No, it is determined by the behavior of the subject.
4. Can humans violate the law? Yes. Many violations have been reported. It is also possible to train participants to learn category structures that require violations of the Morton-Massaro law.
5. Why is this relationship called a law? It is called a law because it provides a good description of results in a very wide range of perceptual experiments and because it informs us about important characteristics of the perceptual process. A good example of this kind of law in the physical sciences is Ohm's law. Knowing whether Ohm's law holds or does not hold in particular circumstances informs us about the characteristics of the material under investigation. Similarly, knowing whether the Morton-Massaro law holds or does not hold in particular experiments informs us about functional and structural characteristics of the perceptual process.
6. Why is it called the Morton-Massaro law? Morton was first to point out this behavior in the context of word and letter perception. Massaro et al. have shown that it holds in many experiments in a wide range of domains.
7. What can cases of adherence (or failure to adhere) to the Morton-Massaro law tell us about structural issues? Failures of adherence suggest that the information sources in question influence responses by nonseparable channels. The law tells us nothing about feed-forward versus interactive mechanisms. Both feed-forward and interactive architectures can give rise to performance that adheres or does not adhere to the law.
8. What does the Morton-Massaro law tell us about functional issues? If observers are indeed approximately optimal, adherence to the law suggests that the objective class likelihood ratios should factorize. By the same reasoning, if observers do not follow the law, this suggests that the objective class likelihood ratios do not factorize.
9. Why is performance that adheres to the Morton-Massaro law observed so often? The behavior is optimal for environments in which objective class likelihood ratios factorize. This condition appears to occur often in nature.
10. How can we test whether objective class likelihood ratios factorize? We can test this by building pattern recognition systems that do not assume factorized class likelihoods and comparing them with restricted versions that make the assumption. If the restricted versions perform in generalization tasks as well as the unrestricted versions, then the assumption is a reasonable one. This approach is more convincing the larger the training/testing database.
11. What does this work teach us about the study of cognition? It suggests that structural questions, which focus on internal processing mechanisms, should be complemented by functional questions, which focus on the statistics of the environment, what is optimal in that environment, and what processing mechanisms achieve optimality in that environment.

Many researchers have emphasized the importance of studying perception and cognition at multiple levels of analysis, and several different taxonomies of levels have been proposed (Anderson, 1990; Marr, 1982; Newell, 1982). Here we distinguish among descriptive, functional, and structural levels, and within structural levels we further distinguish between information-processing mechanisms and physical implementations. At a descriptive level, the goal is to find succinct ways of describing patterns of results. The functional level favored by "rational" and Bayesian paradigms (Anderson, 1990; Knill & Richards, 1996) deals with the optimality of such patterns. Many psychologists adopt a structural approach at the information-processing level of analysis for characterizing mechanisms that may account for the obtained patterns of results. At this level of analysis, no effort is made to map information-processing mechanisms to actual physical processes. Neuroscientists, on the other hand, look for the actual underlying physical mechanisms. Although it is useful to treat these different levels of analysis somewhat independently, it is also important to show how they inform and constrain each other. As cognitive neuroscientists, we tend to seek models that can be viewed as characterizing the mechanisms at an information-processing level while suggesting possible physical implementations that might be realizable in the brain.

At the descriptive level, the Morton-Massaro law is a mathematical expression that describes a relationship between external sources of information and the response probabilities produced by perceptual systems in response to these sources. As we have noted repeatedly, this is a law that is not always observed in data but that very often does provide at least a good approximate account. Such an account is often of great value because of its succinctness (Jacobs & Grainger, 1994): While the number of parameters required to describe arbitrary response probability matrices grows exponentially with the number of sources, the number of parameters required to describe the Morton-Massaro law grows linearly,

⁹ We thank Ken Paap for providing an earlier version of this summary, which we have adapted and expanded, with his permission.

thus providing a succinct representation of the response probability matrices typically found in perceptual experiments.

At the functional level, we showed that the Morton-Massaro law reflects an “assumption” made by the perceptual system about the statistical structure of the environment: factorization of likelihood ratios given (within) response categories. This assumption is adequate when it provides good approximations to the actual posterior probabilities of the perceptual categories under consideration (e.g., letters). Thus, one would expect the law to hold when objective likelihood ratios factorize in the world and to be violated when they do not. We analyzed a database of audiovisual speech signals and showed evidence suggesting that conditional independence of audio and video signals is indeed a reasonable assumption. This helps explain why the Morton-Massaro law has been found to hold so well in perceptual experiments in which humans combine information from acoustic and visual speech sources (Massaro, 1989b). We then presented a case in which factorized objective likelihoods was not a good assumption and observed that in such a case individuals violated the Morton-Massaro law.

At the functional level of analysis, the Morton-Massaro law is a relationship of great value to tell us whether perceptual systems behave consistently with the factorability assumption. Ultimately, the ubiquity of the Morton-Massaro law may be explained by the fact that the assumption provides good approximations to the statistical structure of the environment in a great number of cases; see Domingos and Pazzani (1997) for a review of pattern recognition problems for which the assumption works very well. However, violations should be expected. This kind of law is not unheard of in the physical sciences. For example, Ohm’s law states that the relationship between voltage and current is linear. Although Ohm’s law was initially thought to be a universal principle, we now know that many materials violate it. Ultimately, Ohm’s law is a criterion to classify materials into ohmic (e.g., conductors) and nonohmic (e.g., ionized gas and junction diodes). (See Figure 21.) This criterion is treated as a law because of its tremendous practical value. We believe the Morton-Massaro law of information integration holds a status similar to that of Ohm’s law in physics.

At an information-processing level, the Morton-Massaro law is also useful, even though it has different implications for our understanding of the underlying processes than some (e.g., Massaro, 1989c) originally thought. It is consistent with a series of models, including factorized strength models like the FLMP and

random utility models like Morton’s logogen model. Such models can be viewed either as descriptive or as mechanistic; in the latter case, a common characteristic of these models is that signals are selectively evaluated by separate modules, each of which processes only one information source (i.e., one of the factors in Massaro’s factorial design) and is oblivious to the other sources. These autonomous modules map low-level inputs into representations that reflect the degree of support of these inputs for high-level perceptual categories.

What we have established here is that the Morton-Massaro law is also consistent with other classes of models that do not make the modular, feed-forward assumption characteristic of the mechanistic interpretation of FLMP. We presented a class of interactive models called diffusion networks, and we showed that in these models the Morton-Massaro law does not rule out either feedback or lateral connections. Instead, it is related to an structural constraint that we called channel separability. We can think of channel separability as a way to implement in interactive systems the assumption of factorability of likelihood ratios. The concept of channel separability, a structural concept, finds its mirror image in the assumption of factorability, a functional concept. At a functional level, separable channels in an interactive model are similar to distinct, independent modules in feed-forward models, although there is a big difference: In models with separable channels, influences between the channels can occur; they are simply constrained to occur in accordance with the Morton-Massaro law.

Our work illustrates some interesting points about the issue of modularity of cognitive architectures. Suppose we are given a black-box with some input sensors (e.g., a video camera and a microphone) and some output devices (e.g., a series of lights that indicate which of n words has been uttered). We test the system using Massaro’s experimental paradigm and find that response probabilities factorize into components selectively controlled by audio and video signals. According to the optimality theorem, we would be justified to say that this black box “assumes” a world in which audio and video signals are conditionally independent. At this level of analysis, we would also be justified to say that the black box treats audio and video in a modular manner. However, the work presented here shows that it is possible for the black box to appear modular at this level of analysis (in that it treats sources as independent) while being completely interactive at other levels of analysis. For example, if given access to the internal circuitry, we could very well find that all the circuit elements are bimodal (i.e., that they are modulated by audio and visual information). This possibility is, in fact, consistent with findings emerging from the neurosciences: Although acoustic and visual speech are perceived and integrated in a manner consistent with the Morton-Massaro law, visual information from lip movements has been shown to modify the activity of the human auditory cortex (Sam et al., 1991).

In our view, an important property of interactive hardware is that it affords a very flexible form of computation in which many different tasks can be solved with the same circuitry. In interactive networks, what counts as input and output is not at all fixed. Depending on the task at hand, the same unit can act as stimulus, as context, or as a relay for influences in other portions of the system. One can send information through each unit and distribute it throughout the entire net, or one can probe a specific unit to see what the net “thinks” about the specific hypothesis it represents.




Figure 21. Ohm’s law states that the electrical current flowing through a material is proportional to the voltage drop. The law does not reflect a universal principle but a relationship that applies to some devices.

This allows interactive models to propagate new top-level information into a low-level computation on the fly in a very efficient manner.

Despite these potential advantages, it is fair to say that at this point the psychological evidence against and in favor of interactive models is inconclusive. It is, in fact, possible that traditional behavioral experiments based on reaction time and percentage correct may never provide enough evidence to distinguish these two hypotheses. Fortunately, diffusion networks can be given physical interpretations as models of electronic circuits or biological neural networks. That these models have physical interpretations may help bridge gaps between the information-processing level of interest to psychologists and the physical mechanisms of interest to neuroscientists. In this context, it is perhaps worth taking note that physiological evidence presents a picture of the brain in which bidirectional connections and influences between areas are the norm. It is our experience that, although some cognitive psychologists ask "why use interactive models if we can explain our data using feed-forward models," most neuroscientists ask "why do you give us feed-forward models if all we see in the brain is interactive circuitry?" At a minimum, this leaves us with the task of trying to understand how psychological laws (e.g., the Morton-Massaro law) may be implemented in the noisy and interactive circuitry used by the brain. Diffusion networks represent a small step in that direction.

Significant work remains to be done, and we have identified the following lines of research:

1. It would be of interest to investigate whether the Morton-Massaro law is also found at the neural level. One could, in principle, probe single neurons or neural pools and test whether their responses to combined stimuli adhere to the Morton-Massaro law. Diffusion networks are reasonable models of biological neural networks; thus, our analysis provides a good starting point for how to proceed in this direction.

2. Our analysis centered on equilibrium behavior. It would be of interest to investigate whether diffusion networks can also simulate extensions of the FLMP, which describe the time course of information propagation not just equilibrium behavior.

3. More research needs to be done to understand better the computational advantages and disadvantages of interactive architectures.

4. The statistical analysis of audiovisual signals should be extended to other domains and larger databases to explore whether factorized likelihood models provide good approximations to the statistical structure of natural signals as ubiquitously as suggested by the Morton-Massaro law.

Our work illustrates an approach to the study of cognition and perception that emphasizes mathematical traceability of statements, simultaneous pursuit of functional and structural questions, and exploration of neural basis of cognition. Ultimately, the constraints on theory imposed by the convergence of biological, computational, and psychological investigations will provide the strongest guide in the effort to bridge the gaps between the study of the brain, the mind, and the computer.

References

Allman, J., Miezin, F., & McGuinness, E. (1985). Stimulus specific responses from beyond the classical receptive field: Neurophysiological

- mechanisms for local-global comparisons in visual neurons. *Annual Review Neuroscience*, 8, 407–430.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Ashby, F. G. (1989). Stochastic general recognition theory. In D. Vickers & P. L. Smith (Eds.), *Human information processing: Measures, mechanisms, and models* (pp. 435–457). Amsterdam: Elsevier Science.
- Ashby, F. G. (in press). A stochastic version of general recognition theory. *Journal of Mathematical Psychology*.
- Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, 93(2), 154–179.
- Bülthoff, H. H., & Yuille, A. L. (1996). A Bayesian framework for the integration of visual modules. In T. Inui & J. L. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication* (pp. 49–70). Cambridge, MA: MIT Press.
- Clark, J. J., & Yuille, A. L. (1990). *Data fusion for sensory information processing systems*. Boston: Kluwer Academic.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. New York: Wiley.
- Crowther, C. S., Batchelder, W. H., & Hu, X. (1998). A measurement-theoretic analysis of the fuzzy logic model of perception. *Psychological Review*, 102, 396–408.
- Dayan, P., & Sejnowski, T. (1996). Exploration bonuses and dual control. *Machine Learning*, 25, 5–22.
- Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, 93(3), 283–321.
- Domingos, P., & Pazzani, M. (1997). On the optimality of the simple Bayesian classifier under zero-one loss. *Journal of Machine Learning*, 29, 103–130.
- Duda, R. O., & Hart, P. E. (1973). *Pattern classification and scene analysis*. New York: Wiley.
- Efron, B. (1982). *The jackknife, the bootstrap, and other resampling plans*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Estes, W. K. (1975). The locus of inferential and perceptual processes in letter identification. *Journal of Experimental Psychology: General*, 104, 122–145.
- Felleman, D. J., & Essen, D. C. V. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Friedman, D., & Massaro, D. (1988). Understanding variability in binary and continuous choice. *Psychonomic Bulletin and Review*, 5(3), 370–389.
- Garner, W. R., & Morton, J. (1969). Perceptual independence: Definitions, models, and experimental paradigms. *Psychological Bulletin*, 72, 233–259.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Ghazanfar, A. A., & Nicolelis, A. L. (1997). Nonlinear processing of tactile information in the thalamocortical loop. *Journal of Neurophysiology*, 78, 506–510.
- Gidas, B. (1986). Metropolis-type Monte Carlo simulation algorithms and simulated annealing. In J. L. Snell (Ed.), *Topics in contemporary probability and its applications* (pp. 159–232). Boca Raton: CRC Press.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1–67.
- Golden, R. M. (1986). A developmental neural model of visual word recognition. *Cognitive Science*, 10, 241–276.
- Grainger, J., & Jacobs, A. M. (1994). A dual read-out model of word context effects in letter perception: Further investigations of the word superiority effect. *Journal of Experimental Psychology: Human Perception and Performance*, 20(6), 1158–1176.
- Grainger, J., & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, 103, 518–565.
- Gray, M. S., Movellan, J. R., & Sejnowski, T. (1997). Dynamic features for

- visual speechreading: A systematic comparison. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in neural information processing systems* (pp. 751–759). Vol. 9, Cambridge, MA: MIT Press.
- Green, D. M., Weber, D. L., & Duncan, J. E. (1977). Detection and recognition of pure tones in noise. *Journal of the Acoustical Society of America*, 62(4), 948–954.
- Hennecke, M. E., Stork, D. G., & Ventakesh Prasad, K. (1996). Visionary speech: Looking ahead to practical speech reading systems. In D. G. Stork & M. E. Hennecke (Eds.), *Speechreading by humans and machines: Models, systems, and applications* (pp. 331–349). New York: NATO/Springer-Verlag.
- Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, 98(1), 74–95.
- Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Science*, 81, 3088–3092.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2, 259–366.
- Hubel, D. H., & Weisel, T. N. (1962). Receptive fields, binocular orientation and functional architecture in the cat's visual cortex. *Journal of Physiology*, 166, 106–154.
- Hubel, D. H., & Weisel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–243.
- Inui, T., & McClelland, J. L. (Eds.). (1996). *Attention and performance XVI: Information integration in perception and communication*. Cambridge, MA: MIT Press.
- Jacobs, A. M., & Grainger, J. (1994). Models of visual word recognition—Sampling the state of the art. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 1311–1334.
- Karatzas, I., & Shreve, S. E. (1988). *Brownian motion and stochastic calculus*. New York: Springer-Verlag.
- Keppel, G. (1993). *Design and analysis: A researcher's handbook*. Englewood Cliffs, NJ: Prentice Hall.
- Kloeden, P. E., & Platen, E. (1992). *Numerical solutions to stochastic differential equations*. Berlin: Springer.
- Knill, D., & Richards, W. (1996). *Perception as Bayesian inference*. New York: Cambridge University Press.
- Kruschke, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science*, 8, 201–223.
- Kucera, H., & Francis, W. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Kuhl, P. K., & Meltzoff, A. M. (1982). The bimodal perception of speech in infancy. *Science*, 218, 1138–1141.
- Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience*, 15(2), 1605–1615.
- Lee, T., Mumford, D., Romero, R., & Lamme, V. (1998). The role of primary visual cortex in higher level vision. *Vision Research*, 38, 2429–2454.
- Lorente De No, R. (1922). La corteza cerebral del raton (Primera contribucion. La corteza acustica) [The cerebral cortex of the mouse (A first contribution—The acoustic cortex)]. *Trabajos del Laboratorio de Investigaciones Biologicas de la Universidad de Madrid*, 20, 41–78.
- Luce, R. D. (1959). *Individual choice behavior*. New York: Wiley.
- Luce, R. D., & Suppes, P. (1965). Preference, utility and subjective probability. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 3, pp. 249–410). New York: Wiley.
- Luettin, J., Thacker, N., & Beet, S. (1996). Statistical lip modelling for visual speech recognition. In *Proceeding of the VIII European Signal Processing Conference: Vol. I* (pp. 137–140). Trieste, Italy.
- Marr, D. (1982). *Vision*. New York: Freeman.
- Massaro, D. W. (1987a). Categorical partition: A fuzzy logical model of categorization behavior. In S. Harnad (Ed.), *Categorical perception*. Cambridge, England: Cambridge University Press.
- Massaro, D. W. (1987b). *Speech perception by ear and eye: A paradigm for psychological inquiry*. Hillsdale, NJ: Erlbaum.
- Massaro, D. W. (1989a). *Experimental psychology: An information processing approach*. San Diego, CA: Harcourt Brace Jovanovich.
- Massaro, D. W. (1989b). *Perceiving talking faces*. Cambridge, MA: MIT Press.
- Massaro, D. W. (1989c). Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology*, 21, 398–421.
- Massaro, D. W. (1996). Integration of multiple sources of information in language processing. In T. Inui & J. L. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication* (pp. 397–432). Cambridge, MA: MIT Press.
- Massaro, D. W., & Cohen, M. M. (1983). Phonological constraints in speech perception. *Perception and Psychophysics*, 34, 338–348.
- Massaro, D. W., & Cohen, M. M. (1991). Integration versus interactive activation: The joint influence of stimulus and context in perception. *Cognitive Psychology*, 23, 558–614.
- Massaro, D. W., & Friedman, D. (1990). Models of integration given multiple sources of information. *Psychological Review*, 97, 225–252.
- McClelland, J. L. (1981). Retrieving general and specific information from stored knowledge of specifics. In *Proceedings of the Third Annual Conference of the Cognitive Science Society* (pp. 170–172). Berkeley, CA: Cognitive Science Society.
- McClelland, J. L. (1991). Stochastic interactive activation and the effect of context on perception. *Cognitive Psychology*, 23, 1–44.
- McClelland, J. L. (1993). Toward a theory of information processing in graded, random, and interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 655–688). Cambridge, MA: MIT Press.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1–86.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375–407.
- McFadden, D. (1978). Modeling the choice of residential location. In L. Lundqvist & J. Weibull (Eds.), *Spatial interaction and residential location* (pp. 75–96). Amsterdam: North-Holland.
- Mineiro, P., Movellan, J. R., & Williams, R. J. (1998). Learning path distributions using nonequilibrium diffusion networks. In M. Kearns (Ed.), *Advances in neural information processing systems* (pp. 597–599). Cambridge, MA: MIT Press.
- Morton, J. (1969). The interaction of information in word recognition. *Psychological Review*, 76, 165–178.
- Movellan, J. R. (1994). A reinforcement algorithm to learn trajectories with stochastic neural networks. In J. Cowan, G. Tesauro, & J. Alsppector (Eds.), *Advances in neural information processing systems: Vol. 6* (pp. 83–87). San Mateo, CA: Morgan Kaufman.
- Movellan, J. R. (1998). A learning theorem for networks at detailed stochastic equilibrium. *Neural Computation*, 10(5), 1157–1178.
- Movellan, J. R., & Chadderdon, G. (1996). Channel separability in the audio visual integration of speech: A Bayesian approach. In D. G. Stork & M. E. Hennecke (Eds.), *Speechreading by humans and machines: Models, systems, and applications* (pp. 473–487). New York: NATO/Springer-Verlag.
- Movellan, J. R., & McClelland, J. L. (1993). Learning continuous probability distributions with symmetric diffusion networks. *Cognitive Science*, 17, 463–496.
- Movellan, J. R., & McClelland, J. L. (1995). *Stochastic interactive processing, channel separability and optimal perceptual inference: An examination of Morton's law* (Technical Report PDP.CNS.95.4) Pittsburgh, PA: Carnegie Mellon University.

- Mumford, D. (1992). On the computational architecture of the neocortex. *Biological Cybernetics*, 66, 241–251.
- Myers, R. H. (1990). *Classical and modern regression with applications*. Belmont, MA: Duxbury Press.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18, 87–127.
- Oden, G. C., & Massaro, D. W. (1978). Integration of featural information in speech perception. *Psychological Review*, 85, 172–191.
- Oksendal, B. (1992). *Stochastic differential equations*. Berlin: Springer-Verlag.
- Oliveira, J. T. (1961). La representation des distributions extremes bivariées [The representation of extreme bivariate distributions]. *Bulletin of the International Statistical Institute*, 33, 477–480.
- Paap, K. R., Newsome, S. L., McDonald, J. E., & Schvaneveldt, R. W. (1982). An activation-verification model of letter and word recognition: The word superiority effect. *Psychological Review*, 89, 573–594.
- Partan, S., & Marler, P. (1999). Communication goes multimodal. *Science*, 283, 1272–1273.
- Peon, R. H. (1961). Reticular mechanisms of sensory control. In W. Rosenblith (Ed.), *Sensory communication*. Cambridge, MA: MIT Press.
- Phaf, R. H., van der Heijden, A. H. C., & Hudson, P. T. W. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology*, 22, 273–341.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.
- Prasad, K. V., Stork, D. G., & Wolff, G. J. (1993). *Preprocessing video images for neural learning of lipreading* (Technical Report No. TR93-26). Menlo Park, CA: Ricoh California Research Center.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1989). *Numerical recipes in C: The art of scientific computing*. Cambridge, England: Cambridge University Press.
- Rabiner, L. R., & Juang, B.-H. (1993). *Fundamentals of speech recognition*. Englewood Cliffs, NJ: Prentice Hall.
- Ramon y Cajal, S. (1892). *Nuevo concepto de la histologia de los centros nerviosos* [New concept for the histology of nervous centers]. Barcelona: Imprenta de Henrich y Ca.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–107.
- Reif, F. (1967). *Statistical physics*. New York: McGraw-Hill.
- Repp, B. H., Manuel, F. Y., Lieberman, A. M., & Studdert-Kennedy, M. (1983). Exploring the McGurk effect. *Bulletin of the Psychonomic Society*, 21, 366.
- Rumelhart, D. E., & Siple, P. (1974). Process of recognizing tachistoscopically presented words. *Psychological Review*, 81(2), 99–117.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 7–57). Cambridge, MA: MIT Press.
- Sam, N., Aulanko, R., Hari, M., Lounassmaa, O., Lu, S., & Simola, J. (1991). Seeing speech: Visual information from lip movements modifies activity in the human auditory cortex. *Neuroscience Letters*, 127, 141–145.
- Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, 110, 474–494.
- Samuel, A. G. (1996). Does lexical information influence the phonemic restoration of phonemes? *Journal of Experimental Psychology: General*, 125, 28–51.
- Shepard, R. N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology*, 1, 54–87.
- Shiffrin, R. M., & Raaijmakers, J. (1992). The SAM retrieval model: A retrospective and prospective. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *Essays in honor of William K. Estes* (pp. 69–86). Hillsdale, NJ: Erlbaum.
- Sperry, R. W. (1969). A modified concept of consciousness. *Psychological Review*, 76, 532–536.
- Stein, B. E., & Meredith, M. A. (1993). *The merging of the senses*. Cambridge, MA: MIT Press.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donder's method. *Acta Psychologica*, 30, 276–315.
- Sutton, R. S., & Barto, A. G. (1988). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Szentagothai, J., & Arbib, M. A. (1974). *Conceptual models of neural organization*. Cambridge, MA: MIT Press.
- Townsend, J. T., & Ashby, G. (1983). *Stochastic modeling of elementary psychological processes*. Cambridge, England: Cambridge University Press.
- Townsend, J. T., Hu, G. G., & Ashby, F. G. (1981). Perceptual sampling of orthogonal straight line features. *Psychological Research*, 43, 259–275.
- Vapnik, V. N. (1995). *The nature of statistical learning theory*. New York: Springer-Verlag.
- Yellot, J. I. (1977). The relationship between Luce's choice axiom, Thurstone theory of comparative judgement, and the double exponential distribution. *Journal of Mathematical Psychology*, 15, 109–144.
- Zadeh, L. A. (1988). *Fuzzy logic* (Technical Report CSLI-88-116). Stanford, CA: Stanford University, Center for the Study of Language and Information.

(Appendix follows)

Appendix

Complexity of Factorized Strength Models

Let X_1, \dots, X_m be the treatment factors (sources of information) and R the participant's response to these treatments. The elements of $\mathcal{R} = \{1, \dots, r\}$ are the response alternatives, and the elements of \mathcal{X}_i for $i = 1, \dots, m$ are the treatment levels of factor X_i . For any set \mathcal{S} , the term $|\mathcal{S}|$ represents the number of elements in that set. Let $\tilde{\eta}_{X_i}(x_i, k)$ be the parameters of a factorized strength model, that is, according to this model

$$P(R = k | X_1 = x_1, \dots, X_m = x_m) = \frac{\prod_{i=1}^m \tilde{\eta}_{X_i}(x_i, k)}{\sum_{l=1}^r \prod_{i=1}^m \tilde{\eta}_{X_i}(x_i, l)}, \quad (\text{A1})$$

for $x_i \in \mathcal{X}_i$ and $k \in \mathcal{R}$. Thus, this model appears to have $(r)(|\mathcal{X}_1| + \dots + |\mathcal{X}_m|)$ parameters. We can reparameterize the model to make clear that, in fact, some of these parameters are redundant. In particular, let

$$\eta_R(k) = \prod_{i=1}^m \left(\prod_{x_i \in \mathcal{X}_i} \tilde{\eta}_{X_i}(x_i, k) / \tilde{\eta}_{X_i}(x_i, 1) \right)^{1/|\mathcal{X}_i|}, \quad \text{for } k \in \mathcal{R}, \quad (\text{A2})$$

$$\eta_{X_i}(x_i, k) = \frac{\tilde{\eta}_{X_i}(x_i, k)}{\tilde{\eta}_{X_i}(x_i, 1)} \left(\prod_{u \in \mathcal{X}_i} \tilde{\eta}_{X_i}(u, 1) / \tilde{\eta}_{X_i}(u, k) \right)^{1/|\mathcal{X}_i|} \quad (\text{A3})$$

for $i = 1, \dots, m$; $x_i \in \mathcal{X}_i$, and $k \in \mathcal{R}$. Note that

$$\frac{\prod_{i=1}^m \tilde{\eta}_{X_i}(x_i, k)}{\prod_{i=1}^m \prod_{l=1}^r \tilde{\eta}_{X_i}(x_i, l)} = \frac{\eta_R(k) \prod_{i=1}^m \eta_{X_i}(x_i, k)}{\sum_{l=1}^r \eta_R(l) \prod_{i=1}^m \eta_{X_i}(x_i, l)}, \quad (\text{A4})$$

for $i = 1, \dots, m$; $x_i \in \mathcal{X}_i$ and $k \in \mathcal{R}$. Thus, the model with parameters η is indistinguishable from the original model. Moreover, the η parameters have the following constraints

$$\eta_R(1) = 1, \quad (\text{A5})$$

$$\eta_{X_i}(x_i, 1) = 1, \quad \text{for } i = 1, \dots, m; x_i \in \mathcal{X}_i, \quad (\text{A6})$$

$$\prod_{x_i \in \mathcal{X}_i} \eta_{X_i}(x_i, k) = 1, \quad \text{for } i = 1, \dots, m; k \in \mathcal{R}. \quad (\text{A7})$$

Thus, to reproduce an arbitrary factorized model with parameters $\tilde{\eta}$, we need the following parameters: $r - 1$ parameters η_R , and $(r - 1)(|\mathcal{X}_i| - 1)$ parameters η_{X_i} per treatment factor. This gives us a total of $(r - 1)(|\mathcal{X}_1| + \dots + |\mathcal{X}_m| - m + 1)$ free parameters.

Functional Analysis

Let (Ω, \mathcal{F}) be the measurable space in which our random variables are defined. For simplicity, we work with discrete random variables, because generalization to continuous random variables is transparent. The random variables X_1, \dots, X_m represent external sources of information. The random variable A represents a property of these sources that the perceptual system needs to infer (e.g., the correct response alternative). The random variable R represents the participant's response to X_1, \dots, X_m . The set of values taken by X_i with nonzero probability is symbolized \mathcal{X}_i . The set of values taken by R or A with nonzero probability is symbolized \mathcal{R} . The participant's task is to guess A based on X_1, \dots, X_m . We let the joint probability mass function $p_{X_1, \dots, X_m, A, R}$ factorize as follows:

$$\begin{aligned} p_{X_1, \dots, X_m, A, R}(x_1, x_2, a, r) \\ = p_{X_1, \dots, X_m}(x_1, \dots, x_m) p_{A|X_1, \dots, X_m}(a|x_1, \dots, x_m) p_{R|X_1, \dots, X_m, A}(r|x_1, \dots, x_m, a), \end{aligned} \quad (\text{A8})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ and $a, r \in \mathcal{R}$. This factorization corresponds to the assumption that the participant has no direct access to A and must guess it from the information contained in (X_1, \dots, X_m) . The participant's knowledge is captured by the conditional probability mass function $\tilde{p}_{A|X_1, \dots, X_m}$. From this, we define the joint probability mass function $\tilde{p}_{X_1, \dots, X_m, A, R}$ as follows

$$\begin{aligned} \tilde{p}_{X_1, \dots, X_m, A, R}(x_1, \dots, x_m, a, r) \\ = p_{X_1, \dots, X_m}(x_1, \dots, x_m) \tilde{p}_{A|X_1, \dots, X_m}(a|x_1, \dots, x_m) p_{R|X_1, \dots, X_m, A}(r|x_1, \dots, x_m, a), \end{aligned} \quad (\text{A9})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ and $a, r \in \mathcal{R}$. Let P and \tilde{P} be measures on (Ω, \mathcal{F}) consistent with $p_{X_1, \dots, X_m, A, R}$ and $\tilde{p}_{X_1, \dots, X_m, A, R}$. We refer to them as the objective (P) and subjective (\tilde{P}) probability measures.^{A1} We shall assume that $\tilde{p}_{X_1, \dots, X_m, A}(x_1, \dots, x_m, k) \neq 0$, for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $k \in \mathcal{R}$.

Optimality Theorem

A Bayesian system that satisfies the two following conditions conforms to the Morton-Massaro law:

1. It maximizes a weighted sum of subjective discriminant value and entropy

$$E^p(\log \tilde{p}_{X_1, \dots, X_m, A}(X_1, \dots, X_m, R)) + \beta H^p(X_1, \dots, X_m, R), \quad (\text{A10})$$

where E^p and H^p , respectively, symbolize expected values and entropy with respect to the probability measure P , and $\beta > 0$ is a fixed parameter.

2. It assumes a distribution of information sources that factorizes as follows:

$$\begin{aligned} \tilde{p}_{X_1, \dots, X_m, A}(x_1, \dots, x_m, a) \\ = \phi_A(a) \phi_X(x_1, \dots, x_m) \phi_1(x_1, a) \dots \phi_m(x_m, a), \end{aligned} \quad (\text{A11})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$. Here ϕ_A is a term selectively influenced by A , not X , ϕ_X is selectively influenced by X , not A , and the term ϕ_i is selectively influenced by the i th element of X and by A .

Corollary

The equation $p_{R|X_1, \dots, X_m}(a|x_1, \dots, x_m)$ satisfies the Morton-Massaro law for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$, if and only if there is a Bayesian system that satisfies Conditions 1 and 2 of the optimality theorem and whose response probabilities equal $p_{R|X_1, \dots, X_m}(a|x_1, \dots, x_m)$ for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$.

Proof: We seek a probability mass function $p_{R|X_1, \dots, X_m}$ that maximizes

$$E^p(\log \tilde{p}_{X_1, \dots, X_m, A}(X_1, \dots, X_m, R)) + \beta H^p(X_1, \dots, X_m, R). \quad (\text{A12})$$

First, we decompose Equation A12 into terms controlled by the environment and terms controlled by the perceptual system

$$\begin{aligned} \beta H^p(X_1, \dots, X_m) + \sum_{x_1 \in \mathcal{X}_1} \dots \sum_{x_m \in \mathcal{X}_m} p_{X_1, \dots, X_m}(x_1, \dots, x_m) \\ (E^p(\log \tilde{p}_{X_1, \dots, X_m, A}(X_1, \dots, X_m, R | X_1 = x_1, \dots, X_m = x_m)) \\ + \beta H^p(R | X_1 = x_1, \dots, X_m = x_m)), \end{aligned} \quad (\text{A13})$$

where $H^p(X_1, \dots, X_m)$, the joint entropy of X_1, \dots, X_m , and p_{X_1, \dots, X_m} are controlled by the environment. The terms affected by $p_{R|X_1, \dots, X_m}$ are (a)

^{A1} Note that the only subjective term in \tilde{P} is due to $\tilde{p}_{A|X_1, \dots, X_m}$.

$H^p(R|X_1 = x_1, \dots, X_m = x_m)$, the response entropy for fixed values of X_1, \dots, X_m ; and (b) $E^p(\log \bar{p}_{X_1, \dots, X_m A}(X_1, \dots, X_m, R)|X_1 = x_1, \dots, X_m = x_m)$, the subjective discriminant value for fixed values of X_1, \dots, X_m . Thus, for each $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ an optimal policy should maximize

$$E^p[\log \bar{p}_{X_1, \dots, X_m A}(X_1, \dots, X_m, R)|X_1 = x_1, \dots, X_m = x_m] + \beta H^p(R|X_1 = x_1, \dots, X_m = x_m). \quad (\text{A14})$$

Equation A14 is a Helmholtz energy function, with β playing the role of thermal temperature; thus, it is uniquely optimized by a Boltzmann distribution (Reif, 1967)

$$p_{R|X_1, \dots, X_m}(r|x_1, \dots, x_m) = \frac{1}{Z(x_1, \dots, x_m)} \exp\left(\frac{1}{\beta} \log \bar{p}_{X_1, \dots, X_m A}(x_1, \dots, x_m, r)\right), \quad (\text{A15})$$

where

$$Z(x_1, \dots, x_m) = \sum_{k \in \mathcal{R}} \exp\left(\frac{1}{\beta} \log \bar{p}_{X_1, \dots, X_m A}(x_1, \dots, x_m, k)\right), \quad (\text{A16})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $r \in \mathcal{R}$. This expression can be written as

$$p_{R|X_1, \dots, X_m}(r|x_1, \dots, x_m) = \left(\frac{\bar{p}_{X_1, \dots, X_m A}(x_1, \dots, x_m, r)}{\sum_{k \in \mathcal{R}} \bar{p}_{X_1, \dots, X_m A}(x_1, \dots, x_m, k)} \right)^{1/\beta}. \quad (\text{A17})$$

Using Condition 2 we find that

$$p_{R|X_1, \dots, X_m}(r|x_1, \dots, x_m) = \left(\frac{\phi_A(r) \phi_1(x_1, r) \dots \phi_m(x_m, r)}{\sum_{k \in \mathcal{R}} \phi_A(k) \phi_1(x_1, k) \dots \phi_m(x_m, k)} \right)^{1/\beta}. \quad (\text{A18})$$

We now take probability ratios between two response alternatives and find that they factorize

$$\frac{p_{R|X_1, \dots, X_m}(r|x_1, \dots, x_m)}{p_{R|X_1, \dots, X_m}(k|x_1, \dots, x_m)} = \left(\frac{\phi_A(r) \phi_1(x_1, r)}{\phi_A(k) \phi_1(x_1, k)} \right)^{1/\beta} \dots \left(\frac{\phi_m(x_m, r)}{\phi_m(x_m, k)} \right)^{1/\beta}. \quad (\text{A19})$$

The ‘‘if’’ part of the corollary is obvious. The ‘‘only if’’ part can be proven by construction. Because $p_{R|X_1, \dots, X_m}$ follows the Morton-Massaro law, there are $\eta_{X_i}(\cdot, \cdot) > 0$ terms such that

$$p_{R|X_1, \dots, X_m}(k|x_1, \dots, x_m) = \frac{\eta_{X_1}(x_1, k) \dots \eta_{X_m}(x_m, k)}{\sum_{l=1}^r \eta_{X_1}(x_1, l) \dots \eta_{X_m}(x_m, l)}, \quad (\text{A20})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$. Consider now a system with subjective likelihood functions such that

$$\bar{p}_{A|X_1, \dots, X_m}(x_1, \dots, x_m|k) = \frac{1}{Z(k, \beta)} (\eta_{X_1}(x_1, k) \dots \eta_{X_m}(x_m, k))^\beta, \quad (\text{A21})$$

for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $a \in \mathcal{R}$. Here β is a positive constant and Z is a proportionality function needed to make sure probabilities add up to 1. Such a system assumes conditional independence, thus satisfying Condition 2 of the optimality theorem. Moreover, if we use uniform subjective priors, that is, $\bar{p}_A(k) = 1/r$ for $k \in \mathcal{R}$, then from Equation A18 it follows that a system maximizing

$$E^p(\log \bar{p}_{X_1, \dots, X_m A}(X_1, \dots, X_m, R)) + \beta H^p(X_1, \dots, X_m, R), \quad (\text{A22})$$

exhibits the desired distribution of external responses.

Condition 2 of the Optimality Theorem

In this section, we analyze useful cases under which adherence to Condition 2 of the optimality theorem can be proven.

Case 1: conditionally independent sources. In this case

$$p_{X_1, \dots, X_m A}(x_1, \dots, x_m|r) = p_{X_1|A}(x_1|r) \dots p_{X_m|A}(x_m|r), \quad (\text{A23})$$

for all for $(x_1, \dots, x_m) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, $r \in \mathcal{R}$. Thus, in this case, likelihood ratios factorize, satisfying Condition 2 of the optimality theorem. An example of this case is shown in Figure 7A.

Case 2: Gaussian categories with equal covariance within classes. In this case,

$$p_{X|A}(x|k) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2} (x - \mu_k)' \Sigma^{-1} (x - \mu_k)\right), \quad (\text{A24})$$

for $x \in \mathbb{R}^n$, $k \in \mathcal{R}$. Here Σ is the within-class covariance matrix and μ_k the within-class expected value. After a few steps, likelihood ratios can be expressed as follows:

$$\frac{p_{X|A}(x|k)}{p_{X|A}(x|l)} = \exp(x' \Sigma^{-1} (\mu_l - \mu_k)) \frac{\exp(-1/2 \mu_k' \Sigma^{-1} \mu_k)}{\exp(-1/2 \mu_l' \Sigma^{-1} \mu_l)}, \quad (\text{A25})$$

factorizing as prescribed by Condition 2 of the optimality theorem. An example of this case is shown in Figure 7B.

Case 3: monotonic transformations of sources with factorized likelihood ratios. Let Z_1, \dots, Z_n be discrete random variables with factorized likelihood ratios. Let $X_i = h_i(Z_i)$ for $i = 1, \dots, n$, where h_i is a strictly monotonic function. It follows that

$$p_{X|A}(x|k) = p_{Z|A}(h_i^{-1}(x)|k), \quad (\text{A26})$$

for $x \in \mathbb{R}^n$, $k \in \mathcal{R}$. Because $p_{Z|A}$ factorizes, so does $p_{X|A}$. A similar argument can be made for continuous random variables. An example of this case is shown in Figure 7C.

Separability Theorem

Proposition: Consider a diffusion network as defined in Equations 38 through 43 with symmetric weights and external input $X = (X_1, \dots, X_m)'$. Let C_j be the channel for X_j , $j = 1, \dots, m$ as defined earlier. If $C_i \cap C_j = \emptyset$ for $j = 1, \dots, m$, $j \neq i$, then the equilibrium density of responses factorizes into a term controlled by X_i and a term jointly controlled by all the other inputs.

Proof: Some of the symbols used in this proof are defined in Equations 38 through 43. Without loss of generality, we let $\rho_i = 1$ in Equation 39 for $i = 1, \dots, n$. Let $\alpha = (\alpha_1, \dots, \alpha_n)$ be a vector of gain parameters. Let $y \in \mathbb{R}^n$ represent the internal state of a diffusion network and z the activation vector corresponding to state y , that is, $z_i = \varphi(\alpha_i y_i)$ for $i = 1, \dots, n$. We assume the activation function is bounded monotonic and differentiable (e.g., the logistic function). This satisfies the conditions for existence of stochastic equilibrium specified in Gidas (1986). Call X_i the stimulus and the rest of the external inputs the context. Let z^s , z^c , and z^r represent the components of z for the units in the stimulus channel, context channel, and response specification module. Let $x \in \mathbb{R}^m$ be a vector

representing an external input and let x^s, x^c be the components of x for the external stimulus and context. Let $Z_\alpha(t)$ be a random vector representing the activation at time t of a diffusion network with gain vector α . Let $Z_\alpha = \lim_{t \rightarrow \infty} Z_\alpha(t)$, represent the activations at stochastic equilibrium. In Movellan (1998), it is shown that if the weights are symmetric (i.e., $w = w'$), then the probability density of Y is as follows:

$$p_{Z_\alpha|X}(z|x^s, x^c) = \frac{1}{K_\alpha(x_s, x_c)} \exp((2/\sigma^2)G_\alpha(z|x_s, x_c)), \quad (\text{A27})$$

where

$$K_\alpha(x_s, x_c) = \int \exp((2/\sigma^2)G_\alpha(z|x_s, x_c)) dz, \quad (\text{A28})$$

$$G_\alpha(z|x) = H(z|x) - \sum_{i=1}^n S_\alpha(z_i), \quad (\text{A29})$$

$$H(z|x) = z'wz/2 + z'vx, \quad (\text{A30})$$

$$S_\alpha(z_i) = \frac{1}{\alpha_i} \int_{\varphi(0)}^{z_i} \varphi^{-1}(u) du + \frac{\sigma^2}{2} \log \frac{dz_i}{dy_i} + \frac{\gamma}{2} y_i^2, \quad (\text{A31})$$

and φ^{-1} is the inverse function of φ . The H function is sometimes known as the "harmony" and S_i as the "stress." The equilibrium density is derived from the fact that the drift is the gradient of G with respect to the internal states and that the activations are a monotonic function of the internal states (Movellan, 1998).

Without loss of generality, hereafter we set $\sigma^2 = 2$. When there are no direct connections between the stimulus and context units, there are no terms in the goodness function in which x^s or x^c occur jointly with x^c or x^s . Because of this, the goodness can be separated into three additive terms that depend on x^s, x^c , and a third term that depends on the response units:

$$G_\alpha(z^s, z^c, z'|x^s, x^c) = G_\alpha^s(z^s, z'|x^s) + G_\alpha^c(z^c, z'|x^c) + G_\alpha^r(z^r), \quad (\text{A32})$$

where

$$G_\alpha^s(z^s, z'|x^s) = (z^s)'w_{s,s}z^s/2 + (z^s)'w_{s,r}z^r + (z^s)'v_{s,s}x^s + (z^s)'v_{r,s}x^s - \sum_i S_\alpha(z_i^s), \quad (\text{A33})$$

$$G_\alpha^c(z^c, z'|x^c) = (z^c)'w_{c,c}z^c/2 + (z^c)'w_{c,r}z^r + (z^c)'v_{c,c}x^c + (z^c)'v_{r,c}x^c - \sum_i S_\alpha(z_i^c), \quad (\text{A34})$$

$$G_\alpha^r(z^r) = (z^r)'w_{r,r}z^r/2 - \sum_i S_\alpha(z_i^r). \quad (\text{A35})$$

Here $w_{s,r}$ is a submatrix of w connecting the stimulus and response units. Similar notation is used for the other submatrices of w and u . It follows that we can write the ratio of the joint probability density of any two activation states z, \bar{z} as follows:

$$\frac{p_{Z_\alpha|X}(z^s, z^c, z'|x^s, x^c)}{p_{Z_\alpha|X}(\bar{z}^s, \bar{z}^c, \bar{z}'|x^s, x^c)} = \frac{\exp(G_\alpha^s(z^s, z'|x_s) + G_\alpha^c(z^c, z'|x^c) + G_\alpha^r(z^r))}{\exp(G_\alpha^s(\bar{z}^s, \bar{z}'|x_s) + G_\alpha^c(\bar{z}^c, \bar{z}'|x^c) + G_\alpha^r(\bar{z}^r))}, \quad (\text{A36})$$

which factorizes as desired. To obtain probability densities for the response units, we integrate over the states of all the other units

$$p_{Z_\alpha|X}(z^r|x^s, x^c) = \iint p_{Z_\alpha|X}(z^s, z^c, z'|x^s, x^c) dz^s dz^c, \quad (\text{A37})$$

and after rearranging terms

$$p_{Z_\alpha|X}(z^r|x^s, x^c) = \frac{1}{K_\alpha(x^s, x^c)} \left(\int \exp(G_\alpha^s(z^s, z'|x^s) + G_\alpha^r(z^r)) dz^s \right) \times \left(\int \exp(G_\alpha^c(z^c, z'|x^c)) dz^c \right), \quad (\text{A38})$$

which also factorizes.

Application of the Separability Theorem to Discrete Responses

Here we show that an active-wins policy results in external response probability ratios, which factorize as we increase the gain parameter of the response units. Besides the previous assumptions on the activation function φ , in this section we further assume that it takes values on $[0, 1]$ and its derivative is an even function (e.g., φ could be the logistic function). Let $z^{(1)} = (1, 0, 0, \dots, 0)'$, $z^{(2)} = (0, 1, 0, \dots, 0)'$ be two r -dimensional vectors representing two activation states of the response specification units. Let $\mathbf{1} = (1, 1, 1, \dots, 1)'$. For $i \in \{1, 2\}$ and $\Delta \in (0, 0.5)$ let

$$z_\Delta^{(i)} = (\mathbf{1} - z^{(i)})\Delta + (z^{(i)})(1 - \Delta), \quad (\text{A39})$$

$$R_\Delta^{(i)} = \{x \in \mathbb{R}^r: x_j \in ((1 - \Delta)z_j^{(i)}, \Delta + (1 - \Delta)z_j^{(i)}), \text{ for } j = 1, \dots, r\}. \quad (\text{A40})$$

Consider an active-wins policy in which we choose external response i when the activation of response unit i is larger than $1 - \Delta$ and the activation of all the other response units is smaller than Δ . In this case, the sets $R_\Delta^{(1)}$ and $R_\Delta^{(2)}$ are regions of the $[0, 1]^r$ space, which map into the External Responses 1 and 2. We now investigate the convergence of the probability ratio of these two external responses as we let $\Delta \rightarrow 0$, that is, as the response regions collapse into corners of $[0, 1]^r$,

$$\lim_{\Delta \rightarrow 0} \frac{P(Z_\alpha^r \in R_\Delta^{(2)}|X = x)}{P(Z_\alpha^r \in R_\Delta^{(1)}|X = x)} = \lim_{\Delta \rightarrow 0} \frac{\int_{R_\Delta^{(2)}} p_{Z_\alpha|X}(u|x) du}{\int_{R_\Delta^{(1)}} p_{Z_\alpha|X}(u|x) du} = \quad (\text{A41})$$

$$\lim_{\Delta \rightarrow 0} \frac{\Delta^r p_{Z_\alpha|X}(z_\Delta^{(2)}|x)}{\Delta^r p_{Z_\alpha|X}(z_\Delta^{(1)}|x)} = \lim_{\Delta \rightarrow 0} \frac{\iint e^{G_\alpha(z_\Delta^{(2)}, z', z'|x)} dz^s dz^c}{\iint e^{G_\alpha(z_\Delta^{(1)}, z', z'|x)} dz^s dz^c}. \quad (\text{A42})$$

Now note that

$$G_\alpha(z_\Delta^{(1)}, z^c, z'|x) = H(z_\Delta^{(1)}, z^c, z'|x) - \sum_{i=1}^r S_\alpha(z_{\Delta,i}^{(1)}) - \sum_i S_\alpha(z_i^c) - \sum_j S_\alpha(z_j^r), \quad (\text{A43})$$

where $Z_\Delta^{(1)}$ represents the i th component of the vector $Z_\Delta^{(1)}$. Because φ' is even, it can be shown that $S_\alpha(x) = S_\alpha(1 - x)$ for $x \in (0, 1)$, $i = 1, \dots, r$. Thus,

$$\sum_{i=1}^r S_\alpha(z_{\Delta,i}^{(1)}) = \sum_{i=1}^r S_\alpha(z_{\Delta,i}^{(2)}), \quad (\text{A44})$$

and, therefore,

$$\lim_{\Delta \rightarrow 0} \frac{P(Z_\alpha^r \in R_\Delta^{(2)}|X = x)}{P(Z_\alpha^r \in R_\Delta^{(1)}|X = x)} = \frac{\iint e^{H(z_\Delta^{(2)}, z', z'|x) - \sum_i S_\alpha(z_i^c) - \sum_j S_\alpha(z_j^r)} dz^s dz^c}{\iint e^{H(z_\Delta^{(1)}, z', z'|x) - \sum_i S_\alpha(z_i^c) - \sum_j S_\alpha(z_j^r)} dz^s dz^c}. \quad (\text{A45})$$

It is easy to show that this ratio factorizes. Moreover, for $\Delta > 0$, if we let $\alpha_1 = \dots = \alpha_r = \alpha$, where $\alpha > 0$ then $\lim_{\alpha \rightarrow \infty} P(Z_\alpha^r \in [\Delta, 1 - \Delta]^r) = 0$, because as the gain of the response units increases (see Figure 10) the distribution of Z_α^r converges to a distribution with probability mass at the corner of the $[0, 1]^r$ hypercube and with factorized probability ratios as expressed on Equation A45. This tells us that in the limit, as we increase

the gain, the response probabilities of the external responses associated with Units 1 and 2 factorize. Because the indexing of the response units was arbitrary, this argument holds for all the external responses.

Simulating Diffusion Networks

We simulated diffusion networks using a forward Euler approach, that is stochastic differentials were replaced by discrete time differences of the form

$$Y_i(t + \Delta t) = Y_i(t) + \mu_i(Y(t), X)\Delta t + \sigma \sqrt{\Delta t} N_i(t), \quad (\text{A46})$$

for $i = 1, \dots, n$, where the $N_1(t), \dots, N_n(t)$ are independent standard Gaussian random variables (i.e., with mean zero and variance 1). As $\Delta t \rightarrow 0$, the process defined in Equation 102 converges in distribution to the continuous time process defined by the stochastic differential equation

$$dY_i(t) = \mu_i(Y(t), X)dt + \sigma dB_i(t). \quad (\text{A47})$$

The results presented here hold for the underlying continuous time process and are expected to hold in discrete time simulations provided Δt is small. How small Δt should be is an empirical question whose answer may depend on the problem. In our simulations we found that for $\Delta t < 0.01$ the simulations closely matched the expected theoretical results. The network used for the simulation displayed in Figure 13 had a single unit with no external input. The core of the program looked as follows:

```
sigma=1; alpha=1; gamma= 0.08; dt=0.1; /* Network
parameters */
/* Obtain 50 million sample paths */
for(trial=0;trial < 50000000;trial++){
  y=0.0; /* Initial internal state */
  /* We cycle 600 times */
  for(cycle=0;cycle<600; cycle++){
    z = 1.0/(1.0+exp(-alpha*y)); /* Activation */
    ybar=0; /* In this simple network the net input is
always zero */
    drift= alpha*z*(1.0-z)*(ybar - y) - gamma*y;
    y += drift*dt + sigma*noise()*sqrt(dt); /*
    noise generates standard gaussian samples with
    zero mean and unit variance */
  }
}
```

The theoretical probability density displayed in Figure 13 was obtained as follows:

```
for(z=0.0025;z<=1.0;z+=0.005){
  y = log(z/(1.0 -z))/alpha;
  /* The goodness function */
  g = - (z* log(z) + (1.0-z)*log(1.0-z))/
  alpha - gamma*y*y/2.0;
  pdensity = (exp(2.0*g/sigma/sigma))/z/(1.0-z);
  pdensity=pdensity/15.488505; /* The number 15.488505
was computed using a
Riemann sum approxi-
mation. It makes the
density integrate
to 1.
*/
}
```

For the simulation displayed in Figure 18, the network consisted of 10 units: 6 letter units and 4 word units (see Figure 15). The four word units were fully interconnected with inhibitory connections. Letter units had negative connection to and from other letter units within the same position (e.g., M and R had bidirectional negative connections). Letter units had bidirectional positive connections with consistent word units and bidirectional negative connections with inconsistent word units (e.g., The letter unit E had positive connections to and from word unit RED and negative connections to and from MUD). In the symmetric case, the magnitude of all the weights in the network was 1.5. We also tried a case in which the magnitude of the feed-forward letter-to-word weights was 4.5 and a case in which the magnitude of the feedback word-to-letter weights was 4.5, three times larger than all the other weights. There were four context conditions representing: M_N, R_D, M_D, R_N. On each context condition, the external input to the appropriate context letters was set to 0.5 and the external input to the other context letters was set to 0. There were six stimulus conditions in which the external inputs to the E and U units were as follows: $\{(0, 0.5), (0.1, 0.4), (0.2, 0, 3), (0.3, 0.2), (0.4, 0.1), (0.5, 0)\}$. Thus, there was a total of 24 different input patterns. For each pattern we ran 40,000 trials. On each trial the internal states Y were initialized to 0 and Equation 102 was sequentially applied 1,000 times. At that point the activation of the two response units was collected, and the response alternative with largest activation value was chosen. The core of the program looked as follows:

```
dt = .02,sigma = 0.2, alpha=1.0, gamma= 0.0005,
bias= 1.0, nunits=10;
/* Set the 24 input patterns */
for(i=0.0;i<1.1;i+=1.0){
  for(j=0.0;j<1.1;j+=1.0){
    for(k=0.0;k<1.1;k+=0.2){
      x[pat][0] = i*0.5; x[pat][1] = (1.0- i)*0.5;
      x[pat][2] = j*0.5; x[pat][3] = (1.0- j)*0.5;
      x[pat][4] = k*0.5; x[pat][5] = (1.0- k)*0.5;
      pat++;
    }
  }
}
/* Start the experiment */
for(pat=0;pat<npats;pat++){
  probE[pat]=0.0;
  for(trial=1;trial<=40000;trial++){
    for(i=0;i<nunits;i++) y[i]=0.0; /* Initial
states */
    for(cycle=1;cycle<1000;cycle++){
      for(i=0;i<nunits;i++){
        /* The units are ordered as follows: M,R,E,U,D,
N,MEN,MUD,RED,RUN */
        /* M is unit 0, E and U are units 2 and 3, the
response units */
        ybar[i] = x[pat][i]-bias/alpha; /* Add external
input and
bias to net
input of each
unit
*/
      }
```

(Appendix continues)

```

z[i] = 1.0/(1.0+exp(-alpha*y[i])); /* Compute
activation values */
}
/* Add the incoming activations to the net input
of each unit */
for(j=0;j<nunits;j++){
  ybar[i] += z[j]*w[i][j];
}
}
for(i=0;i<nunits;i++){
  adel = (gain[i]*(z[i])*(1-z[i])*( ybar[i] - y[i])
- gamma*y[i]*dt + sigma*sqrt(dt)*gnoise());
/* gnoise generates standard
Gaussian samples with zero
mean and unit variance.
*/
  y[i] += adel; /* y is the internal state */
}
if(z[2]>z[3]) probE[pat] +=1.0; /* Response deter-
mined by the
unit with larg-
est activation
*/
}
}

```

Received September 17, 1998
Revision received May 16, 2000
Accepted June 6, 2000 ■