

Bimodal speech perception: an examination across languages

**Dominic W. Massaro, Michael M. Cohen,
Antoinette Gesi and Roberto Heredia**

*Program in Experimental Psychology, University of California, Santa Cruz, Santa Cruz,
CA 95064, U.S.A.*

Minoru Tsuzaki

*ATR Auditory & Visual Perception Research Laboratories, Seika-cho, Kyoto, 619-02,
Japan*

Received 28th October 1991, and in revised form 11th November 1992

We examined whether language and culture influence speech perception in face-to-face communication. Native speakers of Japanese, Spanish and English identified the same synthetic unimodal and bimodal speech syllables. Five-step /ba/-/da/ continua were synthesized along auditory and visual dimensions, by varying properties of the syllable at its onset. In the first experiment, the three language groups identified the test syllables as /ba/ or /da/; in the second, Japanese and English speakers were given an open-ended set of response alternatives. For all language groups, identification of the speech segments was influenced by both auditory and visual sources of information. Given the results, we were able to reject an auditory dominance model (ADM) which assumes that the contribution of visible speech is dependent on poor-quality audible speech. The results also falsified a categorical model of perception (CMP) in which the auditory and visual sources are categorized before they are combined. The fuzzy logical model of perception (FLMP) provided a good description of performance supporting the claim that multiple sources of continuous information are evaluated and integrated in speech perception. No differences in the nature of processing across language groups suggests that the underlying mechanisms for speech perception are similar across language and culture.

1. Introduction

Speech perception has been studied extensively in the last decade. We have learned that people use many sources of information in perceiving and understanding speech. One interesting observation is that people manage to communicate under the most adverse conditions imaginable. In one series of investigations, researchers have examined the important contribution of visible information in face-to-face

Please address all correspondence to: Dr Dominic W. Massaro, Program in Experimental Psychology, Clark Kerr Hall, University of California, Santa Cruz, Santa Cruz, CA 95064, U.S.A.

communication. These experiments have shown that visible speech is particularly helpful when the auditory speech is degraded due to noise, bandwidth filtering or hearing-impairment (Summerfield, 1979, 1983; Breeuwer & Plomp, 1984; Massaro, 1987). Although the influence of visible speech is substantial when auditory speech is degraded, visible speech also contributes to performance even when paired with intelligible speech sounds. The importance of visible speech is most directly observed when conflicting visible speech is presented with intelligible auditory speech. As an example, the auditory syllable /ba/ might be dubbed onto a videotape of a talker saying /da/ (McGurk & MacDonald, 1976). A strong effect of the visible speech is observed because the subject will often report perceiving (or even hearing) the syllable /ɔ̃a/. Thus, the strong influence of visible speech is not limited to situations with degraded auditory input, but it also appears to have an important influence even when paired with perfectly intelligible speech sounds.

Although the study of bimodal speech perception has been primarily carried out with English talkers, it offers a valuable domain for the study of cross-linguistic and cross-cultural differences and similarities. It is important to know to what extent the results to date are dependent on language and culture. In addition, cross-linguistic and cross-cultural differences offer a powerful paradigm for broadening the domain for inquiry (Massaro, 1992). Our empirical findings, theories and models often tend to be limited to highly specific situations. Cross-linguistic studies allow us to determine the degree to which we can generalize our conclusions across language and culture.

Our task manipulates synthetic auditory and visual speech in an expanded factorial design, as shown in Fig. 1. Five levels of audible speech varying between

		Visual					
		/ba/	2	3	4	/da/	None
Auditory	/ba/						
	2						
	3						
	4						
	/da/						
	None						

Figure 1. Expanded factorial design used in the current experiments to include both bimodal speech and auditory and visual conditions presented alone. The five levels along the auditory and visible continua represent auditory and visible speech syllables varying in equal physical steps between /ba/ and /da/. For the auditory continuum, /ba/ corresponds to rising F_2 and F_3 transitions and /da/ corresponds to falling F_2 and F_3 transitions. For the visual continuum, /ba/ corresponds to closed lips at the onset of the syllable and /da/ corresponds to open lips at onset.

/ba/ and /da/ are crossed with five levels of visible speech varying between the same alternatives. The onsets of the second and third formants are varied to give an auditory continuum between the syllables /ba/ and /da/. In analogous fashion, parameters of an animated face are varied to give a continuum between visual /ba/ and /da/. This design allows us to address the question of how the identification of a bimodal syllable occurs as a function of the unimodal syllables that compose it. The design is more powerful than a simple factorial design for testing different models (Massaro & Friedman, 1990).

2. Models of bimodal speech perception

We adhere to a falsification and strong-inference strategy of inquiry (Platt, 1964; Massaro, 1987, 1989a). Results are informative only to the degree that they distinguish among alternative theories. Thus, the experimental task, data analysis and model testing are devised specifically to reject some theoretical alternatives. A fuzzy logical model of perception (FLMP), an auditory dominance model (ADM), and a categorical model of speech perception (CMP) are formalized and tested against the results. The FLMP has been the most successful model to date (Massaro, 1987, 1989b, 1990; Massaro & Friedman, 1990) and we begin with the description of this model.

2.1. Fuzzy logical model of perception

The results from a wide variety of experiments have been described within the framework of the FLMP. Within the present framework, speech perception is robust because there are usually multiple sources of information that the perceiver evaluates and integrates to achieve perceptual recognition. The assumptions central to the model are: (1) each source of information is evaluated to give the degree to which that source specifies the relevant alternatives; (2) the sources of information are evaluated independently of one another; (3) the sources are integrated to provide an overall degree of support for each alternative; and (4) perceptual identification follows the relative degree of support among the alternatives.

According to the FLMP, well-learned patterns are recognized in accordance with a general algorithm, regardless of the modality or particular nature of the patterns. Three operations assumed by the model are illustrated in Fig. 2. Continuously-valued features are evaluated, integrated, and matched against prototype descriptions in memory, and an identification decision is made on the basis of the relative goodness of match of the stimulus information with the relevant prototype descriptions.




Figure 2. Schematic representation of the three operations involved in perceptual recognition. The evaluation of an auditory source of information A_i produces a truth value a_i , indicating the degree of support for alternative R . The visual source V_j is evaluated similarly to give v_j . Integration of the truth values gives an overall goodness of match p_{ij} . The response R_{ij} is equal to the value p_{ij} relative to the goodness of match of all response alternatives.

Applying the FLMP to the bimodal speech perception task, both sources are assumed to provide continuous and independent evidence for each of the prototype alternatives. Defining the onsets of the second (F_2) and third (F_3) formants as the important auditory feature and the degree of initial opening of the lips as the important visual feature, the prototype for /da/ might be something like:

/da/ Slightly falling F_2 - F_3 and Open lips.

The prototype for /ba/ would be defined in an analogous fashion,

/ba/ Rising F_2 - F_3 and Closed lips

and so on for the other prototypes.

Given a prototype's independent specifications for the auditory and visual sources, the value of one source cannot change the value of the other source. The integration of the features defining each prototype is evaluated according to the product of the feature values. We let a_{Di} represent the degree to which the auditory stimulus A_i supports the alternative /da/, that is, has Slightly falling F_2 - F_3 . Similarly, v_{Dj} represents the degree to which the visual stimulus V_j supports the alternative /da/, that is, has Open lips. It is assumed that the outcome of prototype matching for /da/ would be a multiplicative contribution of the auditory and visual support:

$$S(/da/ | A_i \text{ and } V_j) = a_{Di} \times v_{Dj} \quad (1)$$

where $S(/da/ | A_i \text{ and } V_j)$ is the support for the prototype /da/ given auditory and visible speech, and the subscripts i and j index the levels of the auditory and visual modalities, respectively. Analogously, if a_{Bi} represents the degree to which the auditory stimulus A_i has Rising F_2 - F_3 and v_{Bj} represents the degree to which the visual stimulus V_j has Closed lips, the outcome of prototype matching for /ba/ would be:

$$S(/ba/ | A_i \text{ and } V_j) = a_{Bi} \times v_{Bj} \quad (2)$$

and so on for the other prototypes.

The decision operation determines the support for one alternative relative to the sum of the support for each of the relevant alternatives. With only a single source of information, such as the auditory one A_i , the probability of a /da/ response, $P(/da/)$, is predicted to be:

$$P(/da/ | A_i) = \frac{a_{Di}}{\sum_k a_{ki}} \quad (3)$$

where the denominator is equal to the sum of support for all relevant (k) alternatives. Similarly,

$$P(/da/ | V_j) = \frac{v_{Dj}}{\sum_k v_{kj}} \quad (4)$$

Given two sources of information A_i and V_j , $P(/da/)$ is predicted to be:

$$P(/da/ | A_i \text{ and } V_j) = \frac{a_{Di} \times v_{Dj}}{\sum_k (a_{ki} \times v_{kj})} \quad (5)$$

As can be seen in Equations (1) and (2), the absolute support for a given prototype will be less for two sources of information than just one. However, the identification judgement is a function of the relative degree of support as shown in Equations (3), (4) and (5). Thus, it is possible that a given identification will be more likely given two sources of information than given just one (Massaro, 1987, Chapter 7).

One important assumption of the FLMP is that the auditory source supports each alternative to some degree and analogously for the visual source. Each alternative is defined by ideal values of the auditory and visual information. The degree of support is given by how much the source matches the corresponding ideal value. Because we cannot predict the degree to which a particular auditory or visible syllable supports a response alternative, a free parameter is necessary for each unique syllable for each unique response. An auditory parameter is forced to remain invariant across variation in the different visual conditions and, analogously, for a visual parameter. Given five levels of auditory and visual speech, the FLMP requires five free parameters for the visual feature values and five for the auditory feature values for each response alternative. (The procedure for estimating the free parameters for the fit of the models is given in Section 4.2.3.)

2.2. Auditory dominance model

A second potential explanation of bimodal speech perception is that an effect of visible speech occurs only when the auditory speech is not completely intelligible (Sekiyama & Tohkura, 1991). Sekiyama & Tohkura tested four labial and six non-labial consonants in the context /a/, under auditory and auditory-visual conditions. The auditory speech was presented either in quiet or in noise. As expected, identification of the auditory speech was very good in quiet and poor in noise. The influence of visible speech in the bimodal condition depended on the quality of the auditory speech. There was very little visual influence with good-quality auditory stimuli and substantial influence with poor-quality auditory speech. For many alternatives, visible speech had an influence for only those auditory stimuli that were not perfectly identified in the auditory condition. However, there were exceptions to this general trend. The auditory syllable /ma/ was perfectly identified in the auditory condition, but was identified as non-labial about 6% of the time when it was paired with a non-labial visible articulation.

The hypothesis that auditory intelligibility determines whether or not visible speech will have an effect is difficult to test, primarily because intelligibility is not easily defined. Perfect identification in one test might not mean perfect intelligibility. Even given these limitations in the measure of intelligibility, we can formulate one version of an intelligibility model, called the auditory dominance model (ADM). The central assumption of the ADM is that the influence of visible speech given a bimodal stimulus is solely a function of whether or not the auditory speech is identified correctly. This model appears to capture Sekiyama & Tohkura's (1991, p. 1804) conclusion that "human beings may depend on eyes in the presence of auditory uncertainty". Similarly, Vroomen (1992) describes (but does not defend) the possibility of lip-reading as a backup device. In this case, the visual information "is relied on whenever the auditory signal is ambiguous". (Vroomen, 1992, p. 9). These views lend themselves to the current instantiation of the ADM in which




Figure 3. Decision trees for ADM for auditory alone, visual alone, and bimodal trials. See text for explanation.

visible speech has a possible influence *only* when the auditory speech is not identified. It should be noted that the all-or-none assumption about auditory identification in the ADM is *not* inconsistent with the assumption that intelligibility is a continuous measure. Intelligibility is determined from a set of identification trials. Even though identification is all-or-none on any given trial, the proportion of identifications over a set of trials would give a continuous measure of intelligibility.

According to the ADM, the probability of a response can be considered to arise from two types of trials given a speech stimulus. Consider first an auditory alone trial. As shown in the top panel of Fig. 3, the auditory speech is identified as one of the response alternatives r or not. When the subject identifies the auditory stimulus as a given alternative r , he or she responds with that alternative. In the case that no identification is made the subject responds with a given alternative with some bias probability w_r . Therefore, the predicted probability of a response on auditory alone trials is equal to

$$P(r|A) = a_r + \left(1 - \sum_r a_r\right)w_r, \quad (6)$$

where a_r is the probability of identifying the auditory source as response r , $\sum_r a_r$ is

the probability of identifying the auditory source as any of the response alternatives, and the term $\left(1 - \sum_r a_r\right)$ is the probability of not identifying the auditory source.

For visual alone trials the situation is analogous. As shown in the middle panel of Fig. 3, the visual speech is identified as one of the response alternatives r or not. When the subject identifies the visual stimulus as a given alternative r , he or she responds with that alternative. In the case that no identification is made the subject responds with a given alternative with the bias probability w_r . Therefore, the predicted probability of a response on visual alone trials is equal to

$$P(r | V) = v_r + \left(1 - \sum_r v_r\right)w_r, \quad (7)$$

where v_r is the probability of identifying the visual source as response r , $\sum_r v_r$ is the probability of identifying the visual source as any of the response alternatives, and the term $\left(1 - \sum_r v_r\right)$ is the probability of not identifying the visual source.

Finally, we consider the bimodal case, shown in the bottom panel of Fig. 3. For these trials the auditory speech is identified as one of the response alternatives r or not. When the subject identifies the auditory stimulus as a given alternative r , he or she responds with that alternative. In the case that no identification is made the subject responds according to the visual information as described above. Therefore, the predicted probability of a response on bimodal trials is equal to

$$P(r | A \text{ and } V) = a_r + \left(1 - \sum_r a_r\right)\left(v_r + \left(1 - \sum_r v_r\right)w_r\right). \quad (8)$$

Equation (8) represents the theory that the auditory stimulus is either identified or else the subject bases his or her decision on the visual information. The visible speech has an influence only when the auditory speech is not identified as one of the alternatives in the task. The model requires an a_r , v_r and w_r for each response alternative. Relative to the FLMP, this model has an additional five parameters for each response alternative.

If speakers of a given language use visible speech only when the auditory speech is *not* identified correctly, then this model should give a better description of the results than the FLMP. This model has the potential of accounting for a small use of visible speech by speakers of a given language.

Finally, one might wonder why an ADM is necessary because auditory dominance could be built into the FLMP and other models. However, the central assumption of the ADM is qualitatively different from the FLMP. In the FLMP, the influence of visible speech in bimodal speech perception is a direct function of its influence in the identification of visible speech in isolation. A good lip-reader will necessarily show some effect of visible speech in bimodal perception. In the ADM, a subject might be a good lip-reader given just visible speech and show very little influence of visible speech in bimodal perception.

2.3. Categorical model of perception

In the categorical model of perception (CMP), it is assumed that only categorical information is available from the auditory and visual sources and that the response is

TABLE I. The probabilities of the four possible outcomes of the two unimodal categorizations of a bimodal speech stimulus for the CMP

Auditory	Visual	
	/b/	not /b/
/b/	$a_{Bi}v_{Bj}$	$a_{Bi}(1 - v_{Bj})$
not /b/	$(1 - a_{Bi})v_{Bj}$	$(1 - a_{Bi})(1 - v_{Bj})$

based on separate categorizations of the auditory and visual sources. The four possible cases are shown in Table I. If the two categorizations to a given speech event agree, the single possible identification response can be based on either source. When the two categorizations disagree, it is assumed that the subject will respond with the categorization to the auditory source on some proportion p of the trials, and with the categorization to the visual source on the remainder $(1 - p)$ of the trials. The weight p reflects the relative dominance of the auditory source. Considering a /ba/ response, the visual and auditory categorizations could be /ba/-/ba/, /ba/-not /ba/, not /ba/-/ba/ or not /ba/-not /ba/.

The probability of a /ba/ identification response given a bimodal speech event is predicted to be:

$$P(/ba/ | A_i \text{ and } V_j) = (1 - p)a_{Bi}v_{Bj} + p a_{Bi}(1 - v_{Bj}) + (1 - p)(1 - a_{Bi})v_{Bj} + (0)(1 - a_{Bi})(1 - v_{Bj}), \quad (9)$$

where i and j index the levels of the auditory and visual modalities, respectively. The a_{Bj} value represents the probability of a /ba/ categorization given the auditory level i , and v_{Bi} is the probability of a /ba/ categorization given the visual level j . The value p reflects the amount of bias to respond with the categorization of the auditory source. Each of the four terms in Equation (9) represents the likelihood of one of the four possible outcomes multiplied by the probability of a /ba/ identification response given that outcome. Note that Equation (9) reduces to:

$$P(/ba/ | A_i \text{ and } V_j) = (p)(a_{Bi}) + (1 - p)v_{Bj}. \quad (10)$$

For each response alternative, the CMP requires five free parameters for the auditory source, five for the visual. A single bias value p is also a necessary free parameter.

It should be noted that the CMP is mathematically equivalent to both a single channel model in which the subject attends to just one modality on bimodal trials (Thompson & Massaro, 1989) and a weighted averaging model in which the subject simply performs a weighted averaging of the two modalities (Massaro, 1987).

3. Previous results and extension to other languages

Experiments using synthetic auditory and visual speech have been carried out with native English-speaking Americans as subjects (Massaro & Cohen, 1990). These subjects give a variety of responses when they are given a range of response alternatives. Both audible and visible speech has a strong influence on performance. In addition, the contribution of one source is larger to the extent the other source is

ambiguous. The details of these judgements were nicely captured in the predictions of the FLMP.

The goal of the present research is to determine if bimodal speech is processed in the same manner across three languages. The FLMP has provided the best description of previous results with English speakers (Massaro, 1989*a,b*, 1990). The question, thus, reduces to asking whether it will also give a superior description of the results from Japanese and Spanish speakers. We can thus assess how language and culture influence unimodal and bimodal speech perception.

We can speculate about what results might be expected from Japanese and Spanish speakers relative to English. All three languages have /b/ and /d/ segments (Maddieson, 1984). Bilingual speakers of any pair of these languages usually claim that these segments are roughly equivalent across the two languages. However, we can be sure that the ideal auditory and visual speech will not be equivalent for these segments across the three languages. The /d/ is more dental for Spanish speakers, for example. The vowel /a/ is shorter in Japanese than in English and Spanish, but the Japanese might interpret the vowel as their long-vowel /ba:/ and /da:/. Differences in phonetic realizations across the languages should have some influence on performance in our task.

The phonological inventories of these three languages also differ from one another. Unlike English, Japanese does not have the phonemes /ð/ or /v/, and American Spanish also does not have the phoneme /v/. These differences have important consequences for the outcome of bimodal speech perception. The syllables /va/ and /ða/ are frequent response alternatives when auditory and visual speech are varied along a /ba/ to /da/ continuum. These alternatives are reasonable because of the auditory and visible properties of these segments. Our research has shown that perceivers respond with alternatives that have the best fit with both the auditory and visual information. Presented with an auditory /ba/ paired with a visible /da/, we might expect the perceiver to respond with one of these two alternatives. However, this is not the case because there is a complete mismatch on one of the two sources of information. On the other hand, the response alternative /ða/ is reasonable. Auditory /ba/ is more similar to auditory /ða/ than /da/, and visible /da/ is also more similar to visible /ða/ than /ba/. Thus, with open-ended response alternatives we would expect that English speakers would respond /ða/ given an auditory /ba/ paired with a visible /da/.

Following this logic, subjects whose language does not have the segment /ð/ should behave differently in this task with open-ended alternatives. We might even expect somewhat different results from subjects who have only learned English as a second language. Mills & Theim (1980) tested native German speakers who had learned English as a foreign language. These subjects identified English bimodal CV syllables consisting of conflicting auditory and visual information. The 15 syllables represent distinctive phonetic categories. The phoneme /ð/ does not occur in German but was considered to be familiar enough to the subjects, who were native speakers of German but had learned English as a foreign language. Both the auditory and visual components had strong effects on identifying what the speaker had said. With respect to identification of the phoneme /ð/, a visual /v/ paired with auditory /ð/ never produced the identification of /ð/. A visual /ð/ paired with auditory /v/ gave 33% /ð/ responses. This result contrasts with the result for English subjects in which a visual /v/ plus auditory /ð/ gave 17% /ð/ responses and

a visual /ð/ plus auditory /v/ gave over 80% /ð/ responses (Massaro, 1987, Chapter 2, Fig. 6). These differences illustrate how the influence of both audible and visible speech is modulated by the perceiver's native language. Analogously, with open-ended response alternatives we would expect fewer, if any, /ða/ responses for the Japanese and Spanish perceivers. Given different phoneme inventories, we would expect different performance in unimodal and bimodal speech perception. Relative to English, the Japanese and Spanish phoneme space can be considered to be less cluttered around /ba/ and /da/. Given an auditory /ba/ paired with a visual /da/, Japanese and Spanish speakers would be more likely to respond with one of these two alternatives.

Another frequent response for English speakers is the consonant cluster /bda/ when the stimulus is a visible /ba/ paired with an auditory /da/. This perceptual judgement is reasonable if subjects choose a response that is the most consistent with both the auditory and visual information. Visible /bda/ is similar to visible /ba/ because both alternatives begin with closure of the lips followed by a mouth opening. Audible /da/ is also similar to audible /bda/ because both segments begin and end with the same formant transitions. Other alternatives given a visual /ba/ and auditory /da/ are less plausible. The alternative /dba/, for example, is not reasonable because of the huge mismatch of visible /dba/ with visible /ba/. The syllable /ba/ begins with a closing of the lips whereas the alternative /dba/ begins with a mouth opening followed by closure of the lips.

In addition to the different phoneme inventories, we would expect cross-linguistic differences because of the differences in the sequential occurrence of consonant segments in these languages. Consonant clusters, such as /bl/, occur in initial position in English. The /bda/ cluster occurs in simple and compound words and across word boundaries. English has words like *abdicate* and *subdue* and word sequences like *crab dish*. Considering the general case of a cluster consisting of labial followed by an alveolar or velar consonant, we found in a word list composed of the Oxford Unabridged Dictionary plus the standard UNIX word list 4 648 cases in 296 666 words or 1.567%. In most of these cases, the clusters occur at syllable boundaries, i.e., a syllable-final consonant followed by a syllable-initial consonant, with some exceptions occurring for the /pt/ cluster, e.g., "apt" and "crypt". In our task, however, the cluster responses are given as syllable initial. Given that other consonant clusters do occur initially in English, /bda/ might be perceived initially even though it only occurs non-initially.

Consonant clusters do not occur initially in the Kansai dialect of Japanese spoken by our subjects. In addition, stop consonant clusters do not occur across word boundaries in Japanese because all words end in vowels or nasals. In a word list of 115 600 Japanese words there were only three occurrences of labial followed by alveolar or velars (0.003%). Compared to English, there are fewer consonant clusters in Spanish. This is confirmed with a count from a word list of 86 061 Spanish words in which there were 557 occurrences of labial followed by alveolar or velar consonant (0.647%). In addition, consonant clusters are also less likely to occur across word boundaries in Spanish than English because a greater proportion of Spanish words end in a vowel. Thus, even if subjects are given an open-ended set of response alternatives, we might expect that Japanese and Spanish speakers would be less likely to respond /bda/ given a visible /ba/ and an auditory /da/. In this case, they might show less of an influence from the visible speech and be more likely to respond /da/.

4. Experiment 1: two alternatives

We therefore face the problem of comparing the processing of speech across languages when these languages differ in their phoneme inventories. It can be argued that the phoneme inventories should play less of a role if subjects are limited to just two responses, /ba/ and /da/. In this case, there should be no difference across languages with respect to the number of prototypes that are functional in the task. English, Spanish and Japanese languages have /ba/ and /da/ syllables. In this case, the additional /va/, /ɔ̃a/, and /bda/ prototypes for the English speakers should have no influence.

According to the FLMP, performance according to the relative goodness rule (RGR) at decision should be a function of only the relevant alternatives in the task (in this case, only /ba/ and /da/). Consider a hypothetical situation of an auditory /da/ paired with a visual /ba/. Even though there would be significant support for different prototypes in the different languages, the probability of a /da/ judgment for all speakers is predicted to be

$$P(/da/ | A_i \text{ and } V_j) = \frac{a_{Di} \times v_{Dj}}{a_{Di} \times v_{Dj} + (1 - a_{Di}) \times (1 - v_{Dj})}, \quad (11)$$

where a_{Di} is the auditory support for /da/ and v_{Dj} is the visual support for /da/. With just two alternatives, it is sufficient to assume in the FLMP that the support for /da/ is given by one minus the support for /ba/ (Massaro, 1987). In terms of Equations (1)–(5) it can be shown that the model makes equivalent predictions if a_{Bi} is assumed to be equal to $(1 - a_{Di})$ and v_{Bj} is assumed to be equal to $(1 - v_{Dj})$ (Massaro, 1989a). Given the success of the FLMP with English speakers, the hypothesis that all speakers process speech in the same manner predicts that the equation will give an equally good description of Japanese and Spanish speakers.

In our previous research, it has been important to distinguish between information and information processing. Information refers to just the output of the evaluation operation in the FLMP (see Fig. 2). Information processing refers to the nature of the evaluation, integration and decision operations, not the input to or output from these operations. Our study primarily addresses differences in information processing across the three languages. Although perceivers of different languages might process speech in the manner described by the FLMP, a given level of auditory or visual information will not necessarily have equivalent effects across the different languages. Given the phonetic differences in the segments /ba/ and /da/ and the phonological differences across the languages, it is unlikely that a given speech stimulus will be identified equivalently. The hypothesis of no differences in information processing predicts only that the FLMP describe the results for speakers of all three languages.

An alternative hypothesis predicts that the FLMP will fail to describe differences across the three language groups. For example, suppose that one language group is less influenced by visible speech in bimodal speech perception—as has been proposed for Japanese perceivers relative to English speakers (Sekiyama & Tohkura, 1993). If this hypothesis is correct, then Equation (11) should fail. Japanese should be about as accurate as English speakers in identifying visible speech when it is presented alone (without the auditory signal), but should use this information less in the identification of bimodal speech. If the Japanese process speech in this manner, the FLMP should give a poor description of their results

because the FLMP cannot predict this type of selective weighting of one of the two sources of information.

Although we have concentrated on potential linguistic differences, cultural differences might also contribute to performance differences among the three language groups. Speakers within one culture might interact differently with one another in face-to-face communication, relative to speakers within another culture. Speakers of one language might tend to avoid face-to-face contact and therefore, not be influenced by visible speech. As a consequence, these speakers would be expected to learn less about visible speech and to be poorer lip-readers. In this case, the FLMP should still predict the results even though there would be less of an influence of visible speech. Although cultural differences are confounded with language differences in the present study, it is still important to acknowledge both of these potential contributions.

The discussion concerning the FLMP and linguistic differences also applies equally to the ADM and CMP. These models also predict the information processing underlying speech perception, not the information available to a given speaker of a given language. For each speaker, the model only specifies how the information is processed given unimodal and bimodal speech.

4.1. Method

4.1.1. Subjects

Three different subject populations were sampled for this 2 h experiment. The subjects spoke American English, Japanese or American Spanish as their first language. The English speakers were 21 students from the University of California, Santa Cruz. The Japanese subjects were 21 students from Dohshisha University. Their ages ranged from 18 to 20-years-old. These subjects spoke the Kansai dialect and their study of English began at age 13 with the "standard" English lessons at junior high school. The average duration of English instruction was eight years. The focus of their English classes was on reading and writing, but not on listening or speaking. None of the students had lived abroad nor taken any lessons in English outside of school. The 20 native Spanish speakers were also from the University of California, Santa Cruz community. The majority of the subjects were Mexican-Americans and three were Puerto Rican-Americans. The Spanish speakers spoke English to various degrees of proficiency. Seventy-two percent of the subjects reported learning English in elementary school (K-5th grade), and 16% reported learning English in ESL (English as a second language) classes in secondary school. Twelve percent reported learning both Spanish and English simultaneously. The Spanish speakers had an average of 14.5 (SD = 4.9) years speaking English.

Some of the English speakers were recruited from the subject pool who were then given 2 h credit for participating. The rest of the English speakers and the Spanish-speaking subjects were recruited by posted advertisements on the campus, through mutual friends, and word of mouth. The English and Spanish speakers were paid \$10.00. The Japanese were paid 4000 yen or the equivalent of roughly \$30 for this experiment and another 1 h experiment having to do with discrimination of auditory speech. Each of the three groups had a different experimenter who was a native speaker of the subject's native language, and only the native language was used during the experiment.

4.1.2. Apparatus and materials

The stimuli were presented by the means of auditory and visible synthetic speech. Using an auditory speech synthesizer, we created a continuum of five sounds that varied between a good /ba/ and a good /da/. The first sound was a good /ba/. The fifth sound was a good /da/. The middle sound was halfway between /ba/ and /da/. The second sound was somewhat more /ba/-like and the fourth sound was somewhat more /da/-like. In an exactly analogous manner using computer animation, we synthesized a face saying /ba/ and /da/ and also saying three syllables intermediate between them. Thus, a five-step continuum going from /ba/ to /da/ was created.

Synthetic audible speech. Tokens of the first author's /ba/ and /da/ were analyzed using linear prediction to derive a set of parameters for driving a software formant serial resonator speech synthesizer (Klatt, 1980). By altering the parametric information specifying the first 80 ms of the consonant-vowel syllable, a set of five 400 ms syllables covering the range from /ba/ to /da/ was created. Figures 4(b) and 4(c) show how some of the acoustic synthesis parameters changed over time for the most /ba/-like and /da/-like of the five auditory syllables. During the first 80 ms, the first formant (F_1) went from 250 Hz to 700 Hz following a negatively accelerated path. The F_2 followed a negatively accelerated path to 1199 Hz, beginning with one of five values equally spaced between 1187 and 1437 Hz from most /ba/-like to most /da/-like, respectively. The F_3 followed a linear transition to 2729 Hz from one of five values equally spaced between 2387 and 2637 Hz. All other stimulus charac-




Figure 4. Visual and auditory parameter values over time for visual /ba/ and /da/ stimuli and auditory /ba/ and /da/ stimuli. (a) Jaw rotation and lip protrusion; (b) voicing amplitude; (c) formants F_1 , F_2 and F_3 . See text for details.




Figure 5. Spectrograms for the five levels of auditory speech between /ba/ and /da/.




Figure 6. Framework (left) and Gouraud shaded (right) renderings of polygon facial model.

teristics were identical for the five auditory syllables. Figure 5 gives the spectrograms of the five syllables along the continuum.

Synthetic visible speech. As employed in Parke (1974), we used a parametrically controlled polygon topology to generate a fairly realistic animation facial display (Cohen & Massaro, 1990). The animation display was created by modeling the facial surface as a polyhedral object composed of about 900 small surfaces arranged in 3D, joined together at the edges (Parke, 1974, 1975, 1982). The left panel of Fig. 6 shows a framework rendering of this model. To achieve a natural appearance, the surface was smooth shaded using Gouraud's (1971) method (shown in the right panel of Fig. 6). The face was animated by altering the location of various points in the grid under the control of 50 parameters, 11 of which were used for speech animation. Control parameters used for several demonstration sentences were selected and refined by the investigator by studying his own articulation frame by frame and estimating the control parameters values (Parke, 1974). Each phoneme is defined in a table according to target values for segment duration, segment type (stop, vowel, liquid, etc) and 11 control parameters. The parameters that are used are jaw rotation, mouth x scale, mouth z offset, lip corner x width, mouth corner z




Figure 7. The facial model at the onset of the syllable for each of the five levels of visible speech between /ba/ and /da/.

TABLE II. Visual synthesis parameters for the five stops, default position and /a/

Parameter	Default	/b/	2	3	4	/d/	/a/
Jaw rotation	3.00	0.00	0.45	0.90	1.35	1.80	10.00
Mouth <i>x</i> scale	1.00	1.00	1.03	1.06	1.09	1.12	1.00
Mouth <i>z</i> offset	0.00	-1.00	-0.85	-0.70	-0.55	-0.40	2.00
Lip corner <i>x</i> width	0.00	0.00	0.75	1.50	2.25	3.00	20.00
Mouth corner <i>z</i> offset	0.00	-15.00	-15.00	-15.00	-15.00	-15.00	0.00
Mouth corner <i>x</i> offset	0.00	2.00	3.50	5.00	6.50	8.00	0.00
Mouth corner <i>y</i> offset	0.00	0.00	0.45	0.90	1.35	1.80	-5.00
Lower lip 'f' tuck	0.00	-5.00	-5.00	-5.00	-5.00	-5.00	0.00
Upper lip raise	0.00	2.00	3.65	5.30	6.95	8.60	2.00

offset, mouth corner *x* offset, mouth corner *y* offset, lower lip 'f' tuck, upper lip raise, and *x* and *z* teeth offset.

Parke's software, revised by Pearce, Wyvill, Wyvill & Hill (1986) and ourselves (Cohen & Massaro, 1990) was implemented on a Silicon Graphics Inc. IRIS 3030 computer. We adapted the software to allow new intermediate test phonemes. To create an animation sequence, each frame was recorded using a broadcast quality BETACAM video recorder under control of the IRIS.

Figure 7 gives pictures of the facial model at the time of maximum stop closure for each of the five levels between /ba/ and /da/. Table II gives the parameter target values used in the visual synthesis for the consonant portion of each visual stimulus, the default resting parameter values, and the values for the vowel /a/. Figure 4(a) shows how the visual synthesis parameters changed over time for the first (/ba/) and last (/da/) visual levels. For clarity, only two of the visual parameters are shown—jaw rotation (larger parameter value means more open), and lip protrusion ("mouth *z* offset" in Table II, smaller value means more protrusion). Not shown in the figure, the face with the default parameter values was recorded for 2000 ms preceding and 2000 ms following the time shown for a total visual stimulus of 4866 ms. A dark screen was presented for the auditory alone trials.

Following the synthesis a Betacam tape was dubbed to 3/4" U-Matic for editing. Only the final 4766 ms of each video sequence was used for each trial. A tone marker was dubbed onto the audio channel of the tape at the start of each syllable to allow the playing of the 400 ms auditory speech stimulus just following the consonant release of the visual stimulus. The marker tone on the video tape was sensed by a schmidt trigger on a PDP-11/34A computer which presented the auditory stimuli from digitized representations on the computer's disk. Figure 4 shows the temporal relationship between the auditory and visual parts of the stimulus. As can be seen in the figure, the parameter transitions specifying the consonantal release occurred at about the same time for both modalities.

4.1.3. Design and procedure

In this experiment, synthetic auditory and visual speech were manipulated in an expanded factorial design previously illustrated in Fig. 1. The onsets of the second and third formants were varied to give an auditory continuum between the syllables /ba/ and /da/. In analogous fashion, we systematically varied parameters of the facial model to give a continuum between visual /ba/ and /da/. Five levels of

audible speech varying between /ba/ and /da/ were crossed with five levels of visible speech varying between the same alternatives. In addition, the audible and visible speech also were presented alone for a total of $25 + 5 + 5 = 35$ independent stimulus conditions. Six random sequences were determined by sampling the 35 conditions without replacement giving six different blocks of 35 trials. These trials were recorded on videotape for use in the experiments.

Subjects were instructed to listen and to watch the speaker, and to identify the syllable as /ba/ or /da/. The native English speakers and the native Spanish speakers were tested at the University of California, Santa Cruz. The native Japanese speakers were tested at ATR in Japan. The same procedure was used for the three different groups with a few minor differences listed below. All subjects were tested on the same videotape. They all received 10 practice trials and the number of test trials were 840 ($35 \times 6 \times 4$). Thus there were 24 observations at each of the 35 unique experimental conditions. Subjects were given a short break after every 210 trials. The display monitor subtended a visual angle of 39 degrees. For the English speakers, the experimental tape was played to the subjects over individual NEC model C12-202A 12-inch color monitors. For the Japanese subjects, the monitor was a National (Panasonic) 14" diagonal screen TH-15B1 and the video tapes were played on a National (Panasonic) AG-6500 VHS recorder. For all three groups, the loudness level of the auditory stimuli was 79 dB (A). The measurement was done with the sound level meter (B&K 2231, with the Microphone Type 4133; Time Weighting, "Fast"; Frequency Weighting, "A"; Display Parameter, "SPL"). The background noise level was 47.5 dB (A).

For the English and Spanish speakers, up to four subjects could be tested simultaneously in individual sound-attenuated rooms. These rooms were each illuminated by two 60 watt incandescent bulbs in a frosted glass ceiling fixture. These subjects made their responses by pressing a key labeled "ba" or "da" on the terminal keyboard. The reaction times (RTs) of these key presses were measured. The Japanese subjects wrote their responses in kana on 3×5 " note cards. The experimenter entered these responses into the terminal after the experiment was over. In all cases, the experimenter was a native speaker of the subject's native language and all instructions and interactions were in the native language.

4.2. Results

Subjects' forced-choice response identifications were recorded for each stimulus. The mean observed proportion of identifications was computed for each subject for the unimodal and bimodal conditions. Separate analyses of variance were carried out on the auditory, visual and bimodal conditions. Both the auditory and the visual sources of information had a strong impact on the identification judgements. As illustrated in Fig. 8, the proportion of responses changed systematically across the visual continuum, both for the unimodal, $F(4, 236) = 244.10$, $p < 0.001$, and the bimodal, $F(4, 236) = 91.03$, $p < 0.001$ conditions. Note that the four degrees of freedom for the numerator comes from the five levels of the stimulus while the 236 degrees of freedom for the denominator comes from the 62 subjects minus the three groups, times the numerator degrees of freedom. Similarly, the pattern of responses changed in an orderly fashion across the auditory continuum, for both the unimodal, $F(4, 236) = 1038.48$, $p < 0.001$, and bimodal, $F(4, 236) = 671.23$, $p < 0.001$,




Figure 8. Probability of a /da/ response as a function of the visual and auditory levels of the speech stimulus for the visual alone (top left plot), auditory alone (bottom left plot) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the English, Japanese and Spanish native speakers.

conditions. Finally, the auditory and visual effects were *not* additive in the bimodal condition, as demonstrated by the significant auditory–visual interaction on response probability, $F(16, 944) = 45.37$, $p < 0.001$.

For the visual alone condition illustrated in the left-most plot of Fig. 8(a), there was no significant difference across the three languages groups, $F(2, 59) = 0.335$, $p = 0.72$. However, the bottom left plot shows a significant difference between language groups in the unimodal auditory condition, both as a simple effect, $F(2, 59) = 10.437$, $p < 0.001$, and as an interaction of auditory level and groups, $F(8, 236) = 8.216$, $p < 0.001$. Not surprisingly, the synthetic auditory speech did not match the natural language categories equivalently across the three languages. The second and third levels along the /ba/–/da/ continuum were identified as /da/ more frequently by the Spanish speakers. The Japanese speakers identified the third level as /da/ less frequently than either the Spanish and English speakers.

The five plots on the top right-hand side of Fig. 8(a) show $P(/da/)$ as a function of the five levels along the visual continuum. The five plots correspond to the five levels of the auditory continuum. Language group is the curve parameter within each plot. There was a significant difference for the auditory source of information in the bimodal condition, $F(2, 59) = 3.954$, $p < 0.05$. The differences are primarily due to the Spanish speakers identifying the three most /ba/-like auditory stimuli as more /da/-like and Japanese subjects identifying these same stimuli as more /ba/-like than the English speakers, as reflected in a significant group by auditory level interaction, $F(8, 236) = 5.120$, $p < 0.001$.

The three groups did not show any significant difference with respect to identification of the syllables along the visual continuum in either the visual [$F(8, 236) = 1.050, p = 0.399$] or bimodal [$F(8, 236) = 1.818, p = 0.074$] conditions.

4.2.1. Relative influence of visible and audible speech

One question of interest is the relative contribution of visible and audible speech in the bimodal condition. An index of the magnitude of the effect of one modality can be calculated by taking the difference in response probabilities to the two endpoint stimuli from that modality. This difference was computed for each subject for each level for both audible and visible sources of information. As an example, given some auditory level, a 0.9 probability of /da/ given the visual /da/ endpoint stimulus and an overall 0.2 probability of /da/ given the visual /ba/ endpoint stimulus would give a visual effect of 0.7. Analyses of variance were carried out on these visual and auditory effect scores. The left plot of Fig. 9 shows the visual effect as a function of the auditory level for the subjects in the bimodal condition for the English, Japanese and Spanish speakers. As can be seen, the visual effect was higher for the more ambiguous central auditory levels than for the end points, $F(4, 236) = 77.476, p < 0.001$. The overall magnitude of the visual effect (0.322, 0.219 and 0.353 for the English, Japanese and Spanish speakers, respectively) did not differ across the three groups, $F(2, 59) = 2.412, p = 0.097$, although there was an interaction between group and auditory level, $F(8, 236) = 4.029, p < 0.001$. We believe that this interaction occurred because the subjects in the three groups perceived the auditory levels somewhat differently. The Japanese speakers, for example had a wider range of marginal $P(/da/)$ judgments. It should be stressed that we do not expect the size of the effect of a given modality to be equivalent across the three languages. The synthesized continuum between /ba/ and /da/ will not match the prototypical values




Figure 9. Visual effect as a function of the auditory level (left plot) and as a function of the probability of a /da/ response given the auditory level (right plot) over subjects in the bimodal condition for the English, Japanese and Spanish speakers.

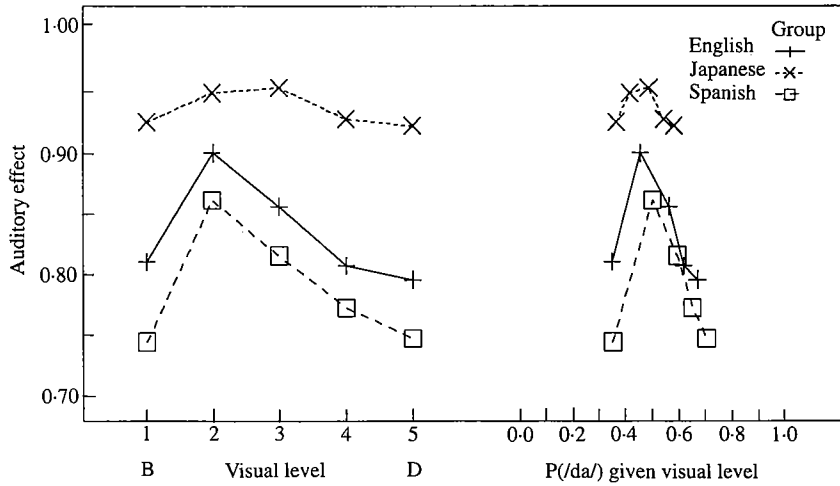


Figure 10. Auditory effect as a function of the visual level (left plot) and as a function of the probability of a /da/ response given the visual level (right plot) over subjects in the bimodal condition for the English, Japanese and Spanish speakers.

of these alternatives equally for the different languages (or even for different subjects within a language). When the visual effect is replotted in the right plot in Fig. 9 as a function of the marginal P(/da/) in the bimodal trials, the curves for the three language groups come more closely into agreement.

Figure 10 shows the analogous analysis for auditory effects. As can be seen in the left plot, there was a larger auditory effect for the more ambiguous visual levels, $F(4, 236) = 7.914$, $p < 0.001$. Similar to the visual effects, the magnitude of the auditory effect (0.834, 0.936 and 0.788 for the English, Japanese and Spanish speakers, respectively) was fairly constant across the three groups, $F(2, 59) = 2.988$, $p < 0.057$, nor was there any interaction of group by visual level. As with the visual effect, the auditory effect is also replotted as a function of the marginal P(/da/) in

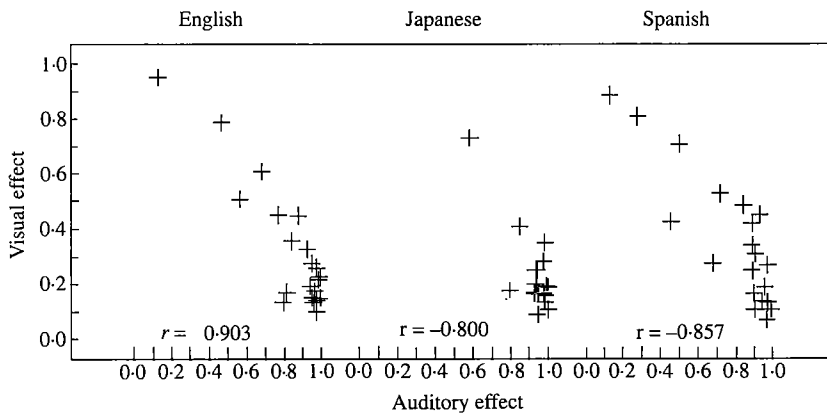


Figure 11. Mean visual effect as a function of the mean auditory effect for each of the subjects in the bimodal condition for the English, Japanese and Spanish speakers. The r value gives the correlation between the two effects.

the bimodal trials. The small group differences show that the auditory endpoints were not equally discriminable across the three languages. The endpoints were more prototypically /ba/ and /da/ for the Japanese than for the English speakers. These endpoints were least prototypical for the Spanish speakers. This is consistent with the significant group by auditory interactions for both the unimodal and bimodal conditions reported above.

Another way to consider the relative sizes of the visual and auditory effects is shown in Fig. 11. This figure plots the size of mean auditory effect as a function of the size of the mean visual effect for the subjects in each of the three groups. Figure 11 illustrates that audible speech had a larger influence than visible speech across the three languages. As can also be seen in the figure, the size of the effects varied significantly across subjects within each group. In addition, there was a strong negative correlation (given in each plot of Fig. 11) between the two effects. To the extent one modality has a large effect, the other had a small effect.

4.2.2. Reaction times

No reaction times (RTs) were recorded for the Japanese subjects. Figure 12 gives the RTs for the English and Spanish speakers as a function of the auditory level, the auditory and visual levels given bimodal speech, and the visual level. As can be seen in the figure, RTs were systematically related to the speech events for both groups

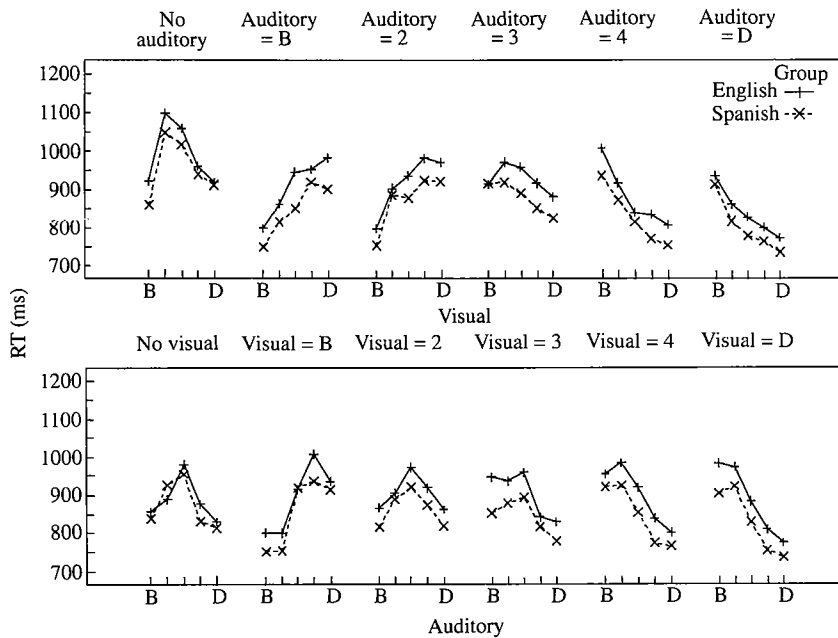


Figure 12. Reaction times as a function of the visual and auditory levels of the speech stimulus for the visual alone (top left plot), auditory alone (bottom left plot) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the English and Spanish native speakers.

of speakers. For the unimodal speech shown in the left hand plots of the two panels, RTs were an inverse U-function of the stimulus level [auditory $F(4, 156) = 39.269$, $p < 0.001$, visual $F(4, 156) = 36.171$, $p < 0.001$]. The other plots in Fig. 12 give the bimodal conditions. RTs varied significantly with changes in both the auditory and visual levels and their combination [auditory $F(4, 156) = 18.620$, $p < 0.001$, visual $F(4, 156) = 3.009$, $p = 0.020$, auditory by visual $F(16, 624) = 42.140$, $p < 0.001$]. Although some differences between the two language groups are apparent in the figure, the only significant difference involving groups was the Group by Auditory level in the auditory alone condition, $F(4, 156) = 2.63$, $p = 0.036$.

The RTs were longer to the extent the two sources of information in combination produced a relatively ambiguous speech event. Figure 13 gives the RTs for the English and Spanish speakers as a function of the overall probability of a /da/ identification to each of the 35 stimulus conditions in the expanded factorial design. This probability is taken as a measure of the ambiguity of the speech event. Probabilities near 0.5 reflect highly ambiguous speech events whereas values near 0 or near 1 reflect the least ambiguous speech events. As can be seen in the figure, RTs were shorter for the less ambiguous stimuli for both groups.

Further analyses of the RTs for identification provide an additional test between the FLMP and CMP. One explanation for the RT results is that subjects are slower to respond to the extent that the auditory and visual information give conflicting information about the speech event. A second explanation is that identification time is positively related to the overall ambiguity of the speech event given both sources of information. These two explanations are not equivalent because a /ba/ sound paired with a /da/ visual articulation could be as ambiguous as an ambiguous sound presented unimodally even though the first case would have more conflicting information.




Figure 13. Reaction times for the English and Spanish speakers as a function of the overall probability of a /da/ identification to each of the 35 stimulus conditions.

These two contrasting explanations may be seen as representative of the CMP and FLMP, respectively. If separate categorizations of the auditory and visual sources are made, then auditory and visual conflict should be the important factor in determining RT. Conflicting auditory and visual sources are likely to lead to different categorizations and thus longer identification RTs because of the additional time needed to resolve the conflict in favor of one source or the other. On the other hand, when the same categorization is made to each of the two sources, the identification RTs should be relatively fast even though one or both sources might be ambiguous. The CMP has no apparent mechanism to account for increases in RT with increases in the ambiguity of the speech event.

By contrast, if continuous information is available, as assumed by the FLMP, RTs should increase with increases in ambiguity of the speech event. Ambiguity should be critical no matter whether the ambiguity is inherent in one source or comes from conflicting information in the two sources. To the extent that the relative goodness of match to a speech category is ambiguous, it should take longer to decide on one of the discrete response alternatives. The FLMP predicts an increase in RT with increases in conflict of the visual and auditory sources only because conflicting sources of information will necessarily create an ambiguous speech event.

To quantify the amount of conflicting information, the parameters corresponding to the categorization probabilities [Equation (10)] for the auditory and visual sources of information given by the CMP were used. The measure of conflict was taken to be the proportion of trials on which auditory and visual information gave different categorizations:

$$C_{ij} = a_i(1 - v_j) + (1 - a_i)v_j, \quad (12)$$

where C_{ij} is a measure of the conflict given by the i th level of the auditory continuum and the j th level of visual articulation. The a_i and v_j values are the parameters representing the proportion of /da/ decisions for each modality from the fit of the CMP. It should be noted that, although conflict is either present or absent on a particular trial, C_{ij} can take on continuous values because it corresponds to the proportion of trials on which conflict occurs.

The amount of ambiguity was defined in terms of the relative goodness of match predicted by the FLMP. With two response alternatives, the most ambiguous speech event would be one in which the predicted $P(/da/)$ would be equal to 0.5. The least ambiguous speech events would be $P(/da/)$ equal to 0 or 1. More generally, ambiguity A_{ij} can be defined in terms of the degree to which the likelihood of a /da/ identification approximates 0.5:

$$A_{ij} = 0.5 - |P(/da/ | A_i \text{ and } V_j) - 0.5|, \quad (13)$$

where $P(/da/ | A_i \text{ and } V_j)$ is the FLMP's predicted probability of a /da/ identification given the i th auditory level and the j th visual level.

To evaluate the conflict and ambiguity explanations, correlation analyses of ambiguity and conflict were carried out separately on the 21 English speakers and 20 Spanish speakers. Each subject contributed 35 C_{ij} , 35 A_{ij} and 35 RT values. The mean correlation for the ambiguity over the 41 subjects was 0.597 *vs.* 0.201 for the conflict. This difference was significant by ANOVA, $F(1, 39) = 91.97$, $p < 0.001$. There was no effect of language group or interaction with this factor. Conflict




Figure 14. Reaction times for each of the 35 stimulus conditions for each of the English (+) and Spanish (x) speakers as a function of ambiguity. Ambiguity is given by the parameter values yielded by Equation (13). The line gives the fit of the linear regression RT as a function of ambiguity.

accounted for very little of the variance (4%) relative to that accounted for by ambiguity (36%).

The points in Fig. 14 give the RT for identification as a function of the ambiguity values [Equation (13)] for the 41 subjects. The English data are marked by (+) and the Spanish by (x). The large variability seen in Fig. 14 is not a result of differences in conflict between the auditory and visual sources. The variability is due to the fact that the different points at a given level of ambiguity in Fig. 14 are from different subjects and different stimulus conditions. It is not unreasonable that the RTs would differ greatly for different subjects and for different stimuli. We know that some subjects are faster than others even though they have the same value of ambiguity, and that RTs to one bimodal stimulus could differ from the RTs to another one even though they have the same value of ambiguity.

4.2.3. Tests of the models

The FLMP, ADM and CMP were fitted to the individual results of each of the 62 subjects. The quantitative predictions of the model are determined by using the program STEPIT (Chandler, 1969). A model is represented to the program in terms of a set of prediction equations and a set of unknown parameters. By iteratively

TABLE III. Average best fitting parameters for the FLMP model for the three language groups. The values index the degree of support for the alternative /da/. The support for /ba/ is 1 minus each value

Group	Modality	Level				
		B	2	3	4	D
English	Visual	0.09301	0.29927	0.59437	0.75356	0.82375
	Auditory	0.03038	0.07960	0.57874	0.96479	0.97723
Japanese	Visual	0.15033	0.30879	0.62895	0.77293	0.82184
	Auditory	0.02108	0.04323	0.33292	0.95303	0.99383
Spanish	Visual	0.10171	0.37746	0.59367	0.70617	0.81773
	Auditory	0.04668	0.20232	0.66521	0.95120	0.98118

adjusting the parameters of the model, the program minimizes the squared deviations between the observed and predicted points. The outcome of the program STEPIT is a set of parameter values which, when put into the model, come closest to predicting the observed results. Thus, STEPIT maximizes the accuracy of the description of a given model. We report the goodness-of-fit of a model by the root mean square deviation (RMSD)—the square root of the average squared deviation between the predicted and observed values.

The continuous lines in Figs 15, 16 and 17 give the average predictions of the FLMP for the English, Japanese and Spanish groups respectively. The FLMP model provides a good description of the identifications of both the unimodal bimodal syllables (average RMSDs of 0.0530, 0.0500 and 0.0595 for the English, Japanese and Spanish speakers, respectively). The predictions of the FLMP come very close to the observations for all three language groups.

Table III gives the average best fitting parameters of the FLMP. These parameter values index the degree of support for each response alternative by each level of the audible and visible stimuli. As can be seen in Table III, the parameter values change in a systematic fashion across the five levels of the audible and visible synthetic speech. For both modalities, the support for the alternative /da/ increases systematically from the B to D level along the continuum. The larger spread for the auditory than for the visual parameters indicates a larger influence for auditory than visual speech. Furthermore, the spread among values is roughly the same across the three language groups.

It should be noted that the predicted points for the FLMP shown in Figs 15–17 cannot be recovered from the parameter values shown in Table III. The figure and table give values averaged across the model fits of the individual subjects in each group. Given that the FLMP is non-linear, the average predictions cannot be computed from the average parameter values.

To fit the ADM to the results, each unique level of the auditory stimulus requires two unique parameters a_b and a_d for each of the five levels along the auditory continuum. Five free parameters are necessary for the five levels along the visual continuum. Finally, an auditory bias parameter w_b is necessary for a total of 21 parameters. The short dashed lines in Figs 15–17 give the average predictions of the ADM for the English, Japanese and Spanish groups, respectively. The RMSD was 0.0674 for the English, 0.0535 for the Japanese and 0.0677 for the Spanish group. An analysis of variance on the RMSD values showed that the FLMP, with just 10 free




Figure 15. Observed (points) and predicted (lines) probability of /da/ response as a function of the visual and auditory levels of the speech stimulus for the visual alone (left top plot), auditory alone (left bottom plot) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the FLMP, ADM and CMP models. Results for the English speakers.

parameters *vs.* the 21 parameters of the ADM, gave a significantly better description of the results, $F(1, 59) = 5.198$, $p = 0.025$.

To fit the CMP to the results, each unique level of the auditory stimulus requires a unique parameter a_{B_i} , and analogously for v_{B_j} . The modeling of /da/ responses thus requires five auditory parameters plus five visual parameters. The p value would be fixed across all conditions for a total of 11 parameters. Thus, we have a fair comparison to the FLMP which requires 10 parameters.

The CMP was fit to the individual results in the same manner as in the fit of the FLMP. The long dashed lines in Figs 15–17 give the average predictions of the CMP for the English, Japanese and Spanish groups, respectively. As can be seen in the figures, the CMP gave a poor description of the observed results. The RMSD was 0.0917 for the English, 0.0886 for the Japanese, and 0.1019 for the Spanish group. An analysis of variance on the RMSD values showed that the FLMP gave a significantly better description of the results than did the CMP, $F(1, 59) = 98.156$, $p < 0.001$.

The CMP is mathematically identical to weighted adding or a weighted averaging model (Massaro, 1987). Thus, a test of the CMP also allows a test of whether the inputs are added or combined in a non-additive manner. The good fit of the FLMP relative to the CMP is evidence against additive integration. The integration of the multiple sources appears to follow a multiplicative combination with the result that the less ambiguous source has the larger impact on performance. Given that the FLMP is mathematically equivalent to Bayes's theorem—an optimal algorithm for

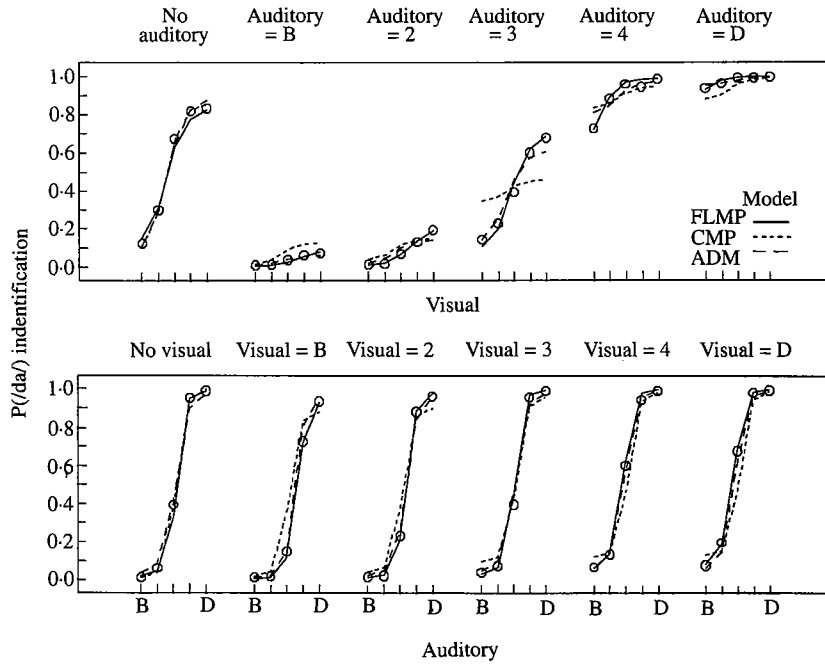


Figure 16. Observed (points) and predicted (lines) probability of a /da/ response as in Fig. 15. Results for the Japanese speakers.

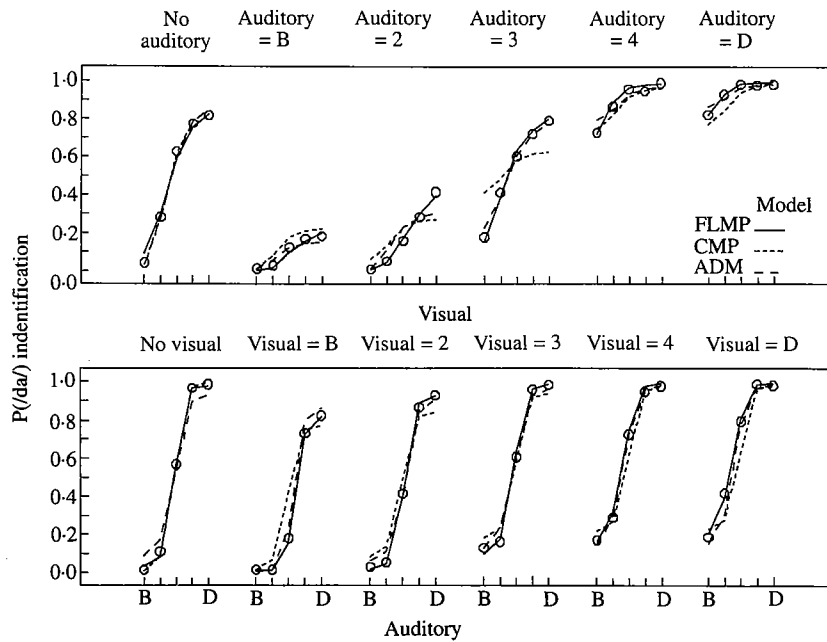


Figure 17. Observed (points) and predicted (lines) probability of a /da/ response in Fig. 15. Results for the Spanish speakers.

integrating multiple sources of information—the good fit of the model to the present results is evident for optimal speech recognition by humans (Massaro, 1987, 1989a; Massaro & Friedman, 1990).

4.3. Discussion

The phoneme repertoires and the phonotactic constraints on the sequential occurrence of segments are different in English, Spanish and Japanese. Thus, we would expect the responses given normal speech, degraded speech, or conflicting bimodal speech to differ across the three languages. Different configurations of responses make it very difficult to answer the question of whether speakers of the different languages are equally influenced by visible speech in bimodal speech perception. This question can be answered, however, by creating synthetic speech that is reasonably representative of two alternative speech categories and limiting all subjects to just those two response alternatives. The two-alternative experiment was successful in comparing bimodal speech perception across three different languages. The results indicated that the processes engaged by bimodal speech are fundamentally equivalent, even though the languages differ substantially. The model tests revealed that the FLMP gave a significantly better description of performance than either the ADM or the CMP. The results also revealed relatively similar visual influences across the three languages. There was no evidence that visible speech played a larger role in one language than another.

5. Experiment 2: open-ended alternatives

Given this outcome, a task with open-ended response alternatives is still of interest for several reasons. First, the nature of the responses in different languages is of interest. Research with English-speaking subjects has shown that subjects respond with the alternative that gives the best match to both the auditory and visual information. Will the same hold true for other languages? Second, will the FLMP continue to give a better description than the other models when speakers are permitted an open-ended set of alternatives. Finally, it is of interest whether visible speech will still have an important influence when open-ended alternatives are permitted. For these reasons, we carried out a second experiment with English and Japanese speakers tested with an open-ended set of response alternatives.

5.1. Method

The English speakers were 16 students. These subjects received \$6.00 an hour. Three students were eliminated because they responded on less than 90% of the test trials. The English-speaking subjects were instructed to listen and to watch the speaker, and to identify the syllable as /ba/, /da/, /bda/, /dba/, /ð/, /va/, /ga/ or "other". These response alternatives were determined from pilot studies in which the responses were not constrained. Each of the 35 possible stimuli were presented a total of 12 times during two sessions and the subject identified each stimulus during a 3 s response interval. Prior to the experimental stimuli, subjects were given 15 practice trials to familiarize them with the task. Subjects were given a short break, approximately 5 min, after completing the tape of 210 trials. Unknown to the subjects, the tape was rewound and played again repeating the 210 trials.

Ten Japanese subjects from the same subject pool used in Experiment 1 were tested. The Japanese subjects wrote their responses in kana. All procedural details were the same as described for the two-choice task.

5.2. Results

The points in Fig. 18 give the observed proportion of responses for the visual alone (left plots), auditory alone (second plots) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the English speakers. These subjects gave a variety of responses. In addition to /ba/ and /da/, /bda/, /ɔ̃a/, and /va/




Figure 18. Observed (points) and predicted (lines) proportion of /ba/, /da/ /bda/, /ɔ̃a/, /va/, /dba/ and /ga/ identifications for the visual alone (left plots), auditory alone (second plots) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the English speakers. The lines give the predictions for the FLMP, ADM and CMP.

were frequent responses. As expected, visible speech had a strong influence on the perceptual judgments of the English speakers. Visible articulations on the /ba/ end of the continuum increased the number of /ba/ judgments. The number of /bda/ judgments increased when a visible /ba/ was paired with an auditory syllable from the /da/ end of the continuum. Visible /da/ articulations increased the likelihood of /da/, /ða/ and /va/ responses. The visual information influenced the likelihood of a /va/ judgment primarily at the /ba/ end of the auditory continuum. The visual /ba/ endpoint stimulus decreased the number of /va/ responses, whereas the other four visual levels increased the number of /va/ judgements at the /ba/ end of the auditory continuum. These judgments reflect the contribution of both auditory and visual speech.

The points in Fig. 19 give the observed results for the Japanese speakers. Subjects mainly responded /ba/ and /da/, but also gave frequent /wa/ and /za/ judgments. The response /ga/ was infrequently given. "Other" trials occurred 1.2% of the time. The category "other" refers to those trials in which subjects did not respond with a specific segment or responded with a segment that was given very infrequently as a response. Japanese subjects were also influenced by visible speech, even when open-ended alternatives were permitted. The top left plot of the visual alone condition shows that the likelihood of a /ba/ judgment decreased as the visible

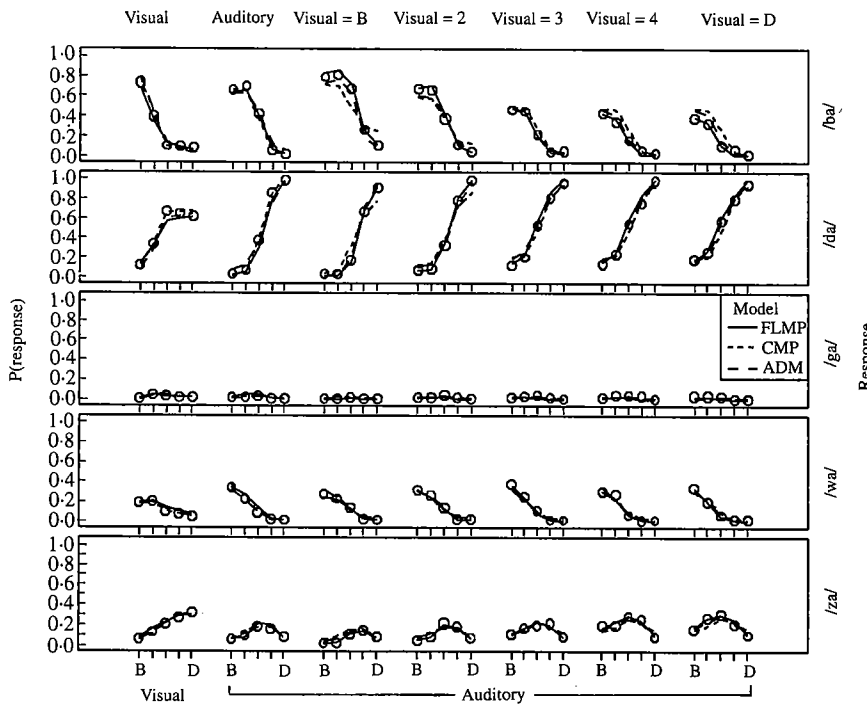


Figure 19. Observed (points) and predicted (lines) proportion of /ba/, /da/ /ga/, /wa/ and /za/ identifications for the visual alone (left plots), auditory alone (second plots) and bimodal (remaining plots) conditions as a function of the five levels of the synthetic auditory and visual speech varying between /ba/ (B) and /da/ (D) for the Japanese speakers. The lines give the predictions for the FLMP, ADM and CMP.

stimulus went from the /ba/ end to the /da/ end of the continuum. This effect of the visible speech also occurred in the bimodal condition, as shown in the top right five plots of Fig. 19. Similarly, the second panel from the top shows that the likelihood of a /da/ judgment increased as the visible stimulus went from the /ba/ end to the /da/ end of the continuum. There were more /za/ responses for the visible speech at the /da/ end of the continuum. The number of /wa/ responses increased at the /ba/ end of the visible continuum.

Comparing the results from Japanese and English speakers supports our analysis of the differences in the phonemic repertoires, phonetic realizations of the syllables, and phonotactic constraints in the two languages. Different responses are given by the two groups of speakers. English speakers tended to respond /da/ or /bda/ when a visual /ba/ was paired with an auditory /da/. Japanese speakers, on the other hand, responded /da/, /ba/ or /za/. When the audible and visible speech comes from the /ba/ side of the continuum, English speakers sometimes responded /va/ where Japanese speaker responded /wa/. Although different responses are given, speakers of both languages are influenced by visible speech.

5.3. Model tests

The FLMP, ADM and CMP were tested against the results. The predictions of the FLMP are given by Equations (3), (4) and (5) in Section 2.1 above. The fit of the FLMP requires 10 free parameters for each response alternative N , unless the response probabilities are constrained to add to 1, in which case only $10(N-1)$ free parameters are necessary. Thus, 70 free parameters are necessary for the English speakers with eight alternatives whereas 50 free parameters are necessary for the Japanese speakers with six alternatives. The number of free parameters for the fit of the ADM also depends of the number of alternatives, N . For the ADM, $(N-1)a + (N-1)v + (N-1)$ free parameters are necessary. For the Japanese speakers, with five specific alternatives plus the other category, $(6-1)5 + (6-1)5 + (6-1) = 55$ free parameters are necessary. For the English speakers, with seven alternatives plus the other category, $(8-1)5 + (8-1)5 + (8-1) = 77$ free parameters are necessary. For the CMP, $N(a+v) + 1$ parameters are necessary, $6(5+5) + 1 = 61$ for the Japanese speakers and $8(5+5) + 1 = 81$ for the English speakers.

The lines in Fig. 18 give the predictions of the FLMP, ADM and CMP. For the English speakers, the seven judgments shown in Fig. 18 accounted for 95.6% of the responses, as a weighted average of the three presentation conditions. For the Japanese speakers, the five judgments shown in Fig. 19 accounted for 98.8% of the responses, computed as for the English speakers.

The RMSD values for the fit of the FLMP, CMP and ADM to the English speakers were 0.0524, 0.1113 and 0.0981, respectively. The RMSD values for the fit of the FLMP, CMP and ADM to the Japanese speakers were 0.0439, 0.0636 and 0.0492, respectively. Analyses of variance were carried out on the RMSD values with language group and model as factors. Comparing the FLMP and ADM, the FLMP gave a significantly better fit than the ADM, $F(1, 21) = 40.464$, $p < 0.001$. In addition, the interaction between language group and model, $F(1, 21) = 25.665$, $p < 0.001$, indicated that the FLMP advantage was greater for the English than for

TABLE IV. Average best fitting parameters for the FLMP model for the English speakers with open-ended responses

Modality	Level	Response						
		/ba/	/da/	/bda/	/dba/	/ða/	/va/	/ga/
Visual	B	0.4607	0.0284	0.1715	0.0392	0.0193	0.1841	0.0179
	2	0.0765	0.0462	0.1019	0.0669	0.0449	0.5216	0.0432
	3	0.0183	0.0720	0.0580	0.0634	0.0804	0.5185	0.0741
	4	0.0235	0.1324	0.0560	0.0417	0.1428	0.3834	0.0763
	D	0.0260	0.1601	0.0490	0.0686	0.1944	0.2666	0.0914
Auditory	B	0.4654	0.0125	0.0414	0.0063	0.1249	0.2984	0.0028
	2	0.2795	0.0278	0.0305	0.0223	0.3375	0.2229	0.0410
	3	0.1274	0.0993	0.0507	0.0541	0.4675	0.1090	0.0455
	4	0.0119	0.4265	0.1033	0.0573	0.3229	0.0183	0.0345
	D	0.0014	0.7121	0.1236	0.0461	0.0688	0.0042	0.0217

the Japanese speakers. Separate ANOVAs indicated that difference between the ADM and FLMP was only marginally significant for the Japanese, $f(1, 9) = 4.972$, $p = 0.051$, *vs.* a more robust difference for the English speakers, $F(1, 12) = 36.428$, $p < 0.001$.

Comparing the somewhat larger differences of the FLMP and CMP, the FLMP gave a significantly better fit than the CMP, $F(1, 21) = 63.447$, $p < 0.001$. Again, the interaction between language group and model, $F(1, 21) = 13.665$, $p = 0.002$, indicated that the FLMP advantage was greater for the English than for the Japanese speakers.

The details of the judgements are best captured in the predictions of the FLMP. Tables IV and V give the average best fitting parameters of the FLMP for the English and Japanese speakers, respectively. These parameter values index the degree of support for each response alternative by each level of the audible and visible stimuli. As can be seen in the tables, the parameter values change in a systematic fashion across the five levels of the audible and visible synthetic speech.

TABLE V. Average best fitting parameters for the FLMP model for Japanese speakers with open-ended responses

Modality	Level	Response				
		/ba/	/da/	/ga/	/wa/	/za/
Visual	B	0.5425	0.1431	0.0000	0.2375	0.0769
	2	0.2797	0.2942	0.0336	0.2068	0.1705
	3	0.1163	0.5034	0.0076	0.1345	0.2382
	4	0.0853	0.4991	0.0032	0.1057	0.3067
	D	0.0691	0.5078	0.0124	0.0890	0.3185
Auditory	B	0.5508	0.0345	0.0125	0.3448	0.0573
	2	0.5640	0.0612	0.0058	0.2744	0.0945
	3	0.3413	0.2573	0.0261	0.1483	0.2185
	4	0.0703	0.6491	0.0000	0.0166	0.2639
	D	0.0268	0.8496	0.0000	0.0094	0.1141

These parameter values account for the relative contributions of the audible and visible speech to the judgements shown in Figs 18 and 19.

6. General discussion

One goal of the present study was to broaden the domain of inquiry in bimodal speech perception by cross-linguistic comparisons. We assessed the contribution of language and/or culture on speech perception in face-to-face communication. The prediction of the RTs, like the identification judgements, show a large advantage for the predictions of the FLMP over the ADM and the CMP. Both the identification responses and their RTs show that speech perception by speakers of different languages can be accounted for by the same processes. Japanese, Spanish and English speakers appear to process bimodal speech in fundamentally the same manner. All subjects were influenced by both auditory and visual sources of information. When the different phonemic repertoires of the two languages, the phonetic realizations of the syllables, and the phonotactic constraints are taken into account, the magnitude of the influence of visible speech appears to be roughly the same across the three languages.

Three models of how auditory and visible sources of information are processed were tested against the results. A reasonable hypothesis is that the contribution of visible speech is dependent on poor-quality audible speech. This hypothesis was formulated in terms of a model in which the perceiver either identifies the auditory information, or else bases the decision on the visual information. The results were also used to test categorical speech perception in which the auditory and visual sources are categorized before they are combined. These two models were contrasted with a model in which multiple sources of continuous information are evaluated and integrated in speech perception. The outcome of the model tests provided unambiguous support for the latter description of all speakers of the three languages. Although there are significant differences in the languages and culture, it appears that the underlying mechanisms for speech perception are similar for the three languages.

The research reported in this paper and the writing of the paper were supported, in part, by grants from the Public Health Service (PHS R01 NS 20314), the National Science Foundation (BNS 8812728), the graduate division of the University of California, Santa Cruz, and support from ATR. The authors would like to thank Mary Beckman and two anonymous reviewers for their comments on the paper.

References

- Breeuwer, M. & Plomp, R. (1984) Speechreading supplemented with frequency-selective sound-pressure information, *Journal of the Acoustical Society of America*, **76**, 686–691.
- Cohen, M. M. & Massaro, D. W. (1990) Synthesis of visible speech, *Behavioral Research Methods and Instrumentation*, **22**, 260–263.
- Chandler, J. P. (1969) Subroutine STEPIT—Finds local minima of a smooth function of several parameters, *Behavioral Science*, **14**, 81–82.
- Gouraud, H. (1971) Continuous shading of curved surfaces, *IEEE Transactions on Computers*, **C-20**, 623–628.
- Klatt, D. H. (1980) Software for a cascade/parallel formant synthesizer, *Journal of the Acoustical Society of America*, **67**, 971–995.
- Maddieson, I. (1984) *Patterns of sounds*. Cambridge: Cambridge University Press.

- Massaro, D. W. (1987) *Speech perception by ear and eye: A paradigm for psychological inquiry*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Massaro, D. W. (1989a) Multiple Book Review of *Speech Perception by Ear and Eye: A Paradigm for Psychological Inquiry, Behavioural and Brain Sciences*, **12**, 741–794.
- Massaro, D. W. (1989b) Testing between the TRACE model and the Fuzzy Logical Model of Perception, *Cognitive Psychology*, **21**, 398–421.
- Massaro, D. W. (1990) A Fuzzy logical Model of Speech Perception, *Proceedings of the XXIV International Congress of Psychology*. In *Human Information Processing: Measures, Mechanisms, and Models* (D. Vickers & P. L. Smith, editors). Amsterdam: North Holland.
- Massaro, D. W. (1992) Broadening the domain of the fuzzy logical model of perception. In *Cognition Conceptual and Methodological Issues* (H. L. Pick, Jr., P. Van den Broek, & D. C. Knill, editors). Washington, DC: American Psychological Association.
- Massaro, D. W. & Cohen, M. M. (1990) Perception of synthesized audible and visible speech, *Psychological Science*, **1**, 55–63.
- Massaro, D. W. & Friedman, D. (1990) Models of integration given multiple sources of information, *Psychological Review*, **97**, 225–252.
- McGurk, H. & MacDonald, J. (1976) Hearing lips and seeing voices, *Nature*, **264**, 746–748.
- Mills, A. E. & Theim, R. (1980) Auditory–visual fusions and illusions in speech perception, *Linguistische Berichte*, **68/80**, 85–108.
- Parke, F. I. (1974) A parametric model for human faces. *Technical Report UTEC-CSc-75-047*. Salt Lake City: University of Utah.
- Parke, F. I. (1975) A model for human faces that allows speech synchronized animation, *Computers and Graphics Journal*, **1**, 1–4.
- Parke, F. I. (1982) Parameterized models for facial animation, *IEEE Computer Graphics*, **2**, 61–68.
- Pearce, A., Wyvill, B., Wyvill, G. & Hill, D. (1986) Speech and expression: A computer solution to face animation. *Paper presented at Graphics Interface '86*.
- Platt, J. R. (1964) Strong inference, *Science*, **146**, 347–353.
- Sekiyama, K. & Tohkura, Y. (1991) McGurk effect in non-English listeners: Few visual effects for Japanese subjects hearing Japanese syllables of high auditory intelligibility, *Journal of the Acoustical Society of America*, **90**, 1797–1805.
- Sekiyama, K. & Tohkura, Y. (1993) Inter-language differences in the influence of visual cues in speech perception, *Journal of Phonetics*, **21**, 427–444.
- Summerfield, A. Q. (1979) Use of visual information in phonetic perception, *Phonetica*, **36**, 314–331.
- Summerfield, A. Q. (1983) Audio-visual speech perception. In *Hearing science and hearing disorders*. (M. E. Lutman & M. P. Haggard editors). London: Academic Press.
- Thompson, L. A. & Massaro, D. W. (1989) Before you see it, you see its parts: Evidence for feature encoding and integration in preschool children and adults, *Cognitive Psychology*, **21**, 334–362.
- Vroomen, J. H. M. (1992) Hearing voices and seeing lips: Investigations in the psychology of lipreading. Doctoral Thesis, Katholieke Universiteit Brabant.