1 **Modelling multimodal integration**
2 **– The case of the McGurk-Effect –**

Fabian Tomaschek¹, Daniel Duran²

3 ¹ Department of General Linguistics, University of Tübingen, Germany

4 ² Deutsches Seminar, Albert-Ludwigs-Universität Freiburg, Germany.

5 **Corresponding author:**

6 Fabian Tomaschek

7 Seminar für Sprachwissenschaft

8 Eberhard Karls University Tübingen

9 Wilhelmstrasse 19

10 Tübingen

11 e-mail: fabian.tomaschek@uni-tuebingen.de

12 Version: September 27, 2019

13 **Abstract**

14 The McGurk effect is a well known perceptual phenomenon in which lis-
15 teners perceive [da], when presented with visual instances of [ga] syllables
16 combined with the audio of [ba] syllables. However, the underlying cognitive
17 mechanisms are not yet fully understood. In this study, we investigated the
18 McGurk effect from the perspective of two learning theories – the Exemplar
19 theory and Discriminative Learning. We hypothesized that the McGurk ef-
20 fect arises from distributional differences of the [ba, da, ga] syllables in a
21 lexicon of a given language. We tested this hypothesis using computational
22 implementations of these theories, simulating learning on the basis of lexica
23 in which we varied the distributions of these syllables systematically. These
24 simulations support our hypothesis.

25 For both learning theories, we found that the probability of observing
26 the McGurk effect in our simulations was greater, when lexica contained a
27 larger percentage of [da] and [ba] instances. Crucially, the probability of the

28 McGurk effect was inversely proportional to the percentage of [ga], what ever
29 the percentage of [ba] and [da].

30 To validate our results, we inspected the distributional properties of [ba,
31 da, ga] in different languages for which the McGurk effect was or was not
32 previously attested. Our results mirror the findings for these languages. Our
33 study shows that the McGurk effect – an instance of multi-modal perceptual
34 integration – arises from distributional properties in different languages and
35 thus depends on language learning.

1 Introduction

It is commonly known what the the McGurk effect is (McGurk & MacDonald, 1976). Hearing [ba] and seeing ⟨ga⟩, test subjects perceived the *percepteme* [da]¹. While it is meanwhile very clear that the human brain integrates visual and acoustic information, because “[why] would the brain not use such available information”, as MacDonald (2018) puts it, it is still not fully understood, how the fused perceptual outcome comes to being. In this present study, we claim that distributional properties of [ba, da, ga] in the language are the reason for the fused outcome. We therefore investigate how distributional properties of speech affect the McGurk effect. Phonetic categories, such as [ba, da, ga] are not signalled by discrete acoustic cues. Rather, acoustic cues show a large variability, for example due to speaking rate (Gay, 1978; Hirata, 2004), word length (Altmann, 1980), coarticulation (Öhman, 1966; Magen, 1997), practice effects (Tomaschek et al., 2013, 2014, 2018a,c), as well as idiosyncratic speaker variation (Tomaschek & Leeman, 2018; Weirich & Fuchs, 2006). There is a large number of studies showing that listeners are highly sensitive to the distributional characteristics of phonetic cues indicating membership to a phonemic category (Clayards et al., 2008; Yoshida et al., 2010; Nixon et al., 2014, 2015; Nixon & Best, 2018; Nixon, 2018). Ignoring the number of phonemic categories that can be discriminated by an acoustic continuum, two sources of variance have been shown to affect perceptual outcomes in listeners: the amount of overlap between categories along and the variance within each category. The stronger the overlap of cues between

¹The classic example of the McGurk effect is described as follows: Hearing [ba] and seeing ⟨ga⟩, test subjects perceive a fused perceptual outcome [da]. We will call the recognized linguistic items or stimulus responses, be they consistent with the perceptual cues or inconsistent, resulting in fused perceptual outcome, *perceptemes*. The term *percepteme* is established in order to distinguish the cognitive response in a categorization task from its phonetic auditory and visual input. Such phonetic input items, i.e. specific instances of realized speech items, are denoted in usual phonetic transcription with square brackets as [da]. Phonemes, i.e. abstract linguistic categories, are denoted in usual phonemic/phonological transcription as /da/. Visual representations of speech items are denoted with angle brackets as ⟨da⟩. Note that we do not use the term *viseme* for the visual representations of speech items. Fisher (1963) introduced this term as an abbreviation for “visual phonemes”, denoting “mutually exclusive classes of sounds visually perceived” – i.e. groups of responses in a confusion matrix (Fisher, 1968, cf.). This term, however, is often used today with a somewhat different meaning, referring to the smallest visual unit of speech in visual or multimodal speech synthesis or recognition. *Perceptemes* will be indicated by white square brackets, e.g. [da]. And, of course this works with [bi, di, gi] and to a certain degree with voiceless plosives as well as glides. See furthermore Rosenblum (2008) for a short review of multimodal speech perception and Sumbly & Pollack (1954) for an earlier finding of the effect.

59 phonemic categories, and the larger the variance within a phonemic category,
60 the less concise are the listeners’ decisions about to which phonemic category
61 stimuli belong.

62 We approached the problem of distributional properties by means of com-
63 putational simulations. We investigated the McGurk effect with two compu-
64 tationally formalized cognitive models that both account for distributional
65 effects in speech production and speech perception. To our knowledge, only
66 Massaro and colleagues used computer simulations to explain the McGurk
67 effect (Massaro & Cohen, 1995). They explained the fused perception of
68 [da] by calculating the probability of confusing auditory and acoustic cues
69 as [da] given the acoustics of the stimulus. To assess the probability of con-
70 fusion, perceived cues are “matched against prototype descriptions [...] and
71 an identification decision is made on the basis of the relative goodness of
72 match of the stimulus information with the relevant prototype description”
73 (Massaro & Cohen, 1995).

74 In the last decade, new models were presented that explain linguistic and
75 cognitive processes during speech production and perception on the basis of
76 the speaker’s experience of language. We use the computational implementa-
77 tions of two of those models in the present study to investigate the origins of
78 the McGurk effect. The first is Exemplar Theory (Johnson, 1997; Goldinger,
79 1998; Pierrehumbert, 2001) which assumes (on an abstract functional level)
80 that individual instances of perceived events are stored in memory as fully
81 specified exemplars. Categories, e.g. equivalent or similar speech sounds,
82 form “clouds” of exemplars which are clustered close together in a percep-
83 tual feature space. Within these exemplar clouds, variability and frequency of
84 occurrence are preserved implicitly. Categorization of new inputs is based on
85 a comparison computing the similarity between the new perceptual stimulus
86 with all stored exemplars. Frequency of occurrence, reflected by the number
87 of exemplars, is an important factor in these similarity computations: the
88 more exemplars in memory with the same associated labels, the higher this
89 category label is weighted in categorization (Duran, 2013, 2015)².

90 The second is Naive Discriminative Learning (Baayen et al., 2011), a
91 computational algorithm based on classical conditioning that simulates error-
92 driven learning (Rescorla & Wagner, 1972; Ramscar et al., 2010, 2013) ac-
93 cording to which phonemic categories are discriminated by means of physical
94 cues such as formant frequencies and, in our case, visual information of the
95 lips. The strength with which cues discriminate outcomes depends on the
96 frequency of co-occurrence of the outcomes and the cues, as well as the dis-
97 tributional properties of cues within and across categories. Frequency of

²See appendix A for a discussion of “similarity” in exemplar-theoretic models.

98 occurrence is implicitly represented by means of stronger association weights
99 between cues and outcomes. However, this strength is further modulated by
100 the distribution of cues between outcomes.

101 We hypothesized that the McGurk effect is a result of the distributional
102 properties of the acoustic and visual cues signalling [ba, da, ga] (see Section
103 2). The more frequent the [da] category, the more often we should be able
104 to attest the McGurk effect in our models. More precise details about the
105 results are not possible at this point because of the high dimensionality of
106 the data and the non-linear behavior of the models under investigation. A
107 verbal description of the expected results would be circular, as the mathe-
108 matical calculations of the present study can be regarded as predictions for
109 participants' behavior in actual experiments. Furthermore, the current study
110 serves as a test bed to compare these two theories and their predictions for
111 linguistic behavior.

112 In the remainder of the paper, we first present the results of a survey
113 on the distributional properties of [ba, da, ga] cues and their frequency of
114 occurrence in different languages. Subsequently, we present the two models.
115 After presenting the results, how consistent the models are in predicting [ba,
116 da, ga] using the consistent cues, we use them to predict the McGurk effect
117 given different distributions of phonetic and visual cues in the phonemic
118 category.

119 2 Analysis of [ba, da, ga]

120 2.1 Distributional properties of physical cues

121 In order to investigate the distributional properties of the acoustic and visual
122 cues of [ba, da, ga], we recorded one female speaker of German (24 years)
123 reading the words that begin with the demisyllable [ba, da, ga] aloud from
124 a sheet of paper. Words were presented in randomized order. The recording
125 was performed in a sound treated booth at the Department of Linguistics,
126 University of Tübingen. As a proxy for visual cues we recorded the lip move-
127 ments of the speaker using electromagnetic articulography (NDI wave 3D ar-
128 ticulograph, sampling frequency of 400 Hz). We will call the cues extracted
129 from articulation *oral cues*. Simultaneously, the audio signal was recorded
130 (Sampling rate: 22.05 kHz, 16bit) and synchronized with the articulatory
131 recordings. We recorded four sensor positions: lower and upper lip, glued
132 inside of the mouth to the center of the lips; left and right corner of the lips,
133 both glued inside of the mouth. We focused on a single speaker in order to
134 reduce noise from idiosyncratic variation. Word boundaries were manually

Figure 1: *Distributions of F1, F2, F3 values (a-c) in the first 50 ms of [ba, da, ga] words, and horizontal and vertical lip distances (d-e) in the 50 ms preceding [ba, da, ga] words. [ba], [da], [ga].*

135 annotated using Praat (Boersma & Weenink, 2015). Annotations were used
 136 to epoch the electromagnetic articulography.

137 In total, we recorded 304 disyllabic German words beginning with either
 138 [ba] (162), [da] (67) or [ga] (75). Figures 1 (a-c) illustrate the distributions
 139 of the mean frequencies of the first three formants F1, F2 and F3 which
 140 we calculated for each word within an interval of 50 ms that began at
 141 the vowel onset. F1 values are lowest for [ga] demisyllables, and highest
 142 for [ba] demisyllables. Although [da] demisyllables would yield a mean F1
 143 value between the other two categories, their acoustic cues strongly overlap
 144 between both [ba] and [ga]. F2 values are minimal for [ba]. Higher F2 values
 145 for [da] and [ga] strongly overlap. By contrast, F3 values are minimal for
 146 [ga] and strongly overlap with [ba]. [da] demisyllables yield the highest F3
 147 values.

148 Figures 1 (d-e) illustrate the distributions of the mean Euclidean dis-
 149 tance between the upper lip sensor and the lower lip sensor (vertical lip
 150 distance), and the left and right lip sensor (horizontal lip distance) which
 151 we calculated for each word in an interval ranging from 50 ms before, until
 152 the acoustic onset of the [ba, da, ga] word. [ba] demisyllables yielded on
 153 average the smallest horizontal lip distance, [ga] demisyllables the largest.
 154 There is a strong overlap in the horizontal distance between [ba], [da] and

Figure 2: *The dotplot illustrates, how [ba, da, ga] are distributed in different languages discussed in Section 2, with 'X' representing languages for which a weak or no McGurk effect has been attested in the literature. The colored dots represent one data set for a simulation run using NDL or Exemplar-based calculations.*

155 [ga] demisyllables, especially between [da] and [ga]. Vertical lip distances
 156 for ipa [ba] demisyllables are smallest and show almost no overlap with [da]
 157 and [ga] demisyllables. Although [da] demisyllables have smaller vertical
 158 lip distances than [ga] demisyllables, there is nevertheless a large amount
 159 of overlap. In summary, the demisyllables [ba, da, ga] are discriminated by
 160 means of overlapping, multi-dimensional and, if vision is taken into account,
 161 di-modal cues. In both modalities there is strong overlap between the three
 162 demisyllables.

163 2.2 Frequency of occurrence in different languages

164 The McGurk effect has been investigated with respect to change in voicing
 165 ([ba], [da], [ga] vs. [pa], [ta], [ka]), change in manner of articulation ([ma],
 166 [na], [fa], [sa], [ra], [wa]) and embedding in the phonetic structure ([aba],
 167 [ada], [aga], [aɖa], [ava], [gy:g], [ge:g]). It has been found that the percentage
 168 of fused responses is higher when stimuli are voiced than when they are
 169 voiceless (McGurk & MacDonald, 1976; Sekiyama & Tohkura, 1991) (but see

170 Sekiyama et al. (1995) for a reversed finding); percentage of fused responses is
 171 highest for plosive stimuli, intermediate for fricatives and lowest for sonorants
 172 (Sekiyama et al., 1995; Massaro & Cohen, 1995) (but see Paré et al. (2003)
 173 for a reversed finding). Knowing that the arisal of the McGurk effect strongly
 174 depends on the temporal synchronization of the articulatory and the auditory
 175 signal (Hertrich et al., 2009), these findings make perfect sense as different
 176 manners of articulation, as well as voicing change the temporal patterns of
 177 the speech signal (for complementary vowel length see Kleber, 2017).

178 Overall, the percentage of fused responses varies strongly between $\sim 95\%$
 179 and $\sim 5\%$ (Paré et al., 2003; Magnotti & Beauchamp, 2015; Sekiyama et al.,
 180 1995; Sekiyama, 1997; Traunmüller, 2009; Schwartz et al., 1998; Schwartz,
 181 2010). It has also been shown that bilinguals give more fused responses
 182 than monolinguals (Marian et al., 2018). However, not only is there large
 183 variation across speakers within one language group, but also between lan-
 184 guages (Sekiyama & Tohkura, 1991; Sekiyama et al., 1995; Aloufy et al.,
 185 1996; Fuster-Duran, 1996; Sekiyama, 1997; Chen & Hazan, 2007; Majewski,
 186 2008; Bovo et al., 2009; Traunmüller, 2009; Magnotti & Beauchamp, 2015;
 187 Mildner & Dobrić, 2015; Zhang et al., 2018). Fused responses have been
 188 attested in German, English, Swedish, Italian and Turkish. By contrast, it
 189 seems as though speakers of Cantonese, Chinese, and Japanese are less likely
 190 to perceive [da] in the McGurk condition.

191 We hypothesize that the reason for this might be that words with [ba],
 192 [da] and [ga] in onset positions might be differently distributed in these lan-
 193 guages. In the following, we refer to [ba], [da] and [ga] holistically as *demi-*
 194 *syllables*, because they are usually attested in words with a coda. We counted
 195 the number of words beginning with the demi-syllables in question for the
 196 following languages within large text corpora: Cantonese (Luke & Wong,
 197 2015), Chinese (Sun, 2018), English (Balota et al., 2011), German (Arnold
 198 & Tomaschek, 2016), Italian (Lison & Tiedemann, 2016), Japanese (Bccwj-
 199 Consortium, 2016), Polish (Lison & Tiedemann, 2016), Swedish (Lison &
 200 Tiedemann, 2016) and Turkish (Lison & Tiedemann, 2016).

201 Indeed, we found that the distribution of the demi-syllables differs strongly
 202 between the languages, as can be seen in Figure 2. It turns out that English,
 203 Swedish, Turkish, Polish, Italian and German form a cluster of languages
 204 with relatively similar [ga] frequencies (Figure 2, b&c), but vary with respect
 205 the trade off between [ba] and [da] (Figure 2, a&b). Cantonese, Chinese and
 206 Japanese, by contrast, have a very low frequency of occurrence of [ba] demi-
 207 syllables (Figure 2, a&c). They vary with respect the distribution of [ga] and
 208 [da] demi-syllables (Figure 2, b).

209 3 Modelling multi-modal perception

210 3.1 Modeling with Naïve Discriminative Learning

211 Naïve Discriminative Learning (NDL, Baayen et al., 2011) is the computa-
 212 tional formalization of the Rescorla-Wagner learning algorithm (Rescorla &
 213 Wagner, 1972; Rescorla, 1988), an error-driven learning algorithm based on
 214 Pavlovian conditioning which is closely related to the perceptron (Rosenblatt,
 215 1962) and adaptive learning in electrical engineering (Widrow & Hoff, 1960).
 216 General cognitive mechanisms such as the blocking effect (Kamin, 1967) or
 217 the feature-label ordering effect (Ramscar et al., 2010) have been modeled
 218 with discriminative learning. Beyond that, NDL has repeatedly been shown
 219 to provide accurate predictions for human behavior in language tasks such
 220 as response times in lexical decision (Baayen et al., 2011), accuracy of the
 221 decision (Arnold et al., 2017), acquisition of morphological categories such as
 222 case and number and their phonological markup (Arnon & Ramscar, 2012;
 223 Ramscar et al., 2013). The NDL is a two layer network which learns to
 224 discriminate a set of output units –*outcomes*– on the basis of a set of in-
 225 put units –*cues*. The trained network can be used to provide predictions on
 226 how strongly acoustic and visual cues activate perceptual outcomes such as
 227 demisyllables.

228 Equation (??) formalizes discriminative learning. It traces one cue-to-
 229 outcome combination across all learning events in the learning set. In order
 230 to calculate the weights for the entire network, the calculation has to be
 231 repeated for all cue-to-outcome combinations.

$$232 \quad 233 \quad w_i^{t+1} = w_i^t + \Delta w_i^t$$

$$234 \quad \Delta w_i^t = \begin{cases} \text{a) } 0 & \text{if ABSENT}(C_i, t) \\ \text{b) } \alpha_i \beta \left(\lambda - \sum_{\text{PRESENT}(C_k, t)} w_j \right) & \text{if PRESENT}(C_i, t) \ \& \ \text{PRESENT}(O, t) \\ \text{c) } \alpha_i \beta \left(0 - \sum_{\text{PRESENT}(C_k, t)} w_j \right) & \text{if PRESENT}(C_i, t) \ \& \ \text{ABSENT}(O, t) \end{cases}$$

235 with w_i representing connection strength between a cue and an outcome,
 236 t iterating across all the learning events specified in a learning set. In every
 237 t event, it is checked whether the cue C and the associated outcome O are

present and their weight w_i^t is adjusted by means of Δw_i^t . The result is used in the next event, w_i^{t+1} . Δw_i^t is calculated in the following way:

- **Condition a)** If neither the cue in question nor its associated outcome are present during an event, Δw_i^t is assigned to zero, i.e. no adjustment happens.
- **Condition b)** If the cue in question and its associated outcome are present during an event, Δw_i^t is calculated by subtracting the sum across the weights of all cues present in the learning event from λ , the maximum learnability of the association between cues and outcomes. λ is set by default to 1 in all our models. The result is weighted by the product of α , representing the salience of the present cue, and β , representing the salience of the situation in which the outcome can be found. We set the product of α and β to 0.001, a value we have found in pilot studies to optimally represent the learning rate of humans.
- **Condition c)** If the cue in question is present, but its associated outcome is not present during an event, Δw_i^t is calculated by subtracting the sum across the weights of all cues present in the learning event from 0. Again, the result is multiplied by α and β .

The error Δw_i^t is the difference between the intended activation (1 when cues and outcomes co-occur together, and 0 when they do not) and the cues' true activation $\sum_{\text{PRESENT}(C_k, t)} w_j$. This results in the general pattern that whenever an event matches a prediction, associations between cues and outcomes are strengthened. When an event does not match a prediction, associations are weakened. The amount of strengthening / weakening is proportional to an outcome's frequency of occurrence. Take a conversation as an example. A speaker utters a word with specific formant frequencies (i.e. acoustic cues). A listener understands $\llbracket \text{ba} \rrbracket$ (i.e. predicts that the speaker meant $\llbracket \text{ba} \rrbracket$) and gets feedback from the speaker that this is correct. Then the association between those specific formant frequencies and $\llbracket \text{ba} \rrbracket$ are strengthened. When the speaker signals that the listener misunderstood, the listener experiences an error and the association between formant frequencies and $\llbracket \text{ba} \rrbracket$ is weakened.

Note that the intended activation is calculated by summing across all cues present in one event. This describes *cue competition*, i.e. the cues compete for the association strength to their outcome. To illustrate how this works in principle, consider the following example. When cues (e.g. the aforementioned formant frequencies) co-occur with different outcomes, association strength between cues and the outcome of interest (the aforementioned $\llbracket \text{ba} \rrbracket$)

are weakend. This process can result in *unlearning*, i.e. cues do not predict certain outcomes anymore. For example, a listener learns that specific formant frequencies are associated with [ba]. But at some point in their lives these specific formant frequencies become associated with [da] and [ba] ceases to occur. With every occurrence of [da] the association between the specific formant frequencies and [ba] becomes weaker. In a framework that is based purely on counts, the relation between these specific formant frequencies and [ba] would have stayed the same as the counts would not have changed³. In NDL, however, the relation changes.

In its current form, NDL requires cues and outcomes to be distinct units, which means that it is not capable of associating gradually changing continua with an outcome, as happens in phonetic learning. Therefore, continuous cues need to be transformed into discrete representations before training. Cue creation is described in the following section.

3.2 Data for the NDL simulation

In this section, we describe how we obtained the data for training NDL (and subsequently the exemplar based model). In order to manipulate frequency of occurrence and the breadth of the variance within each phonemic category, we generated normally distributed values for F1, F2, F3, horizontal and vertical lip distance. Means and standard deviations were equivalent to the raw data described above. Data sets for training were created by manipulating the number of samples in each phonemic category to range between 1000 and 9000 in increments of 1000. In each data set, the *relative frequency* of each phonemic category was calculated. In a second sweep across the frequency ranges, the size of the *standard deviation* of each cue was manipulated to be twice as large as the original standard deviation.

In pilot studies, we tested how discretization of the continua affected the classification task and found no differences depending on whether 15, 20, 25, or 30 bins along each continuum were used. This is why we binned the continuum into 20 steps. In addition to the acoustic and visual cues, an environmental cue as well as a noise cue were added to the training set for NDL. NDL training was performed with the NDL implementation **pyndl** (Sering et al., 2017). ‘McGurk’ stimuli were created by combining each of the acoustic cues of each [ba] instance with the visual cues of each [ga] instance. A new set of normally distributed values was created for each simulation.

In NDL, classification of perceptemes for both, the correct acoustic-visual combination as well as for the McGurk stimuli, were obtained by calculating

³Only in case no memory decay (forgetting) occurs.

Figure 3: *Dotplots (with lines) illustrating the average prediction accuracy in the NDL simulation depending on different standard deviations around the mean (top panel), different numbers of steps to bin the gradual data (second panel) and the proportion of the number of [ba, da, ga] stimuli in the training data set (bottom panel). Line types represents data in the acoustic modality (A), the visual modality (V) and both joined together (AV). Note that A, AV and V overlap to a large degree.*

the activation of each cue set, i.e. the sum of weights between a set of cues and all outcomes in the network, and choosing the phonemic category with the highest activation. Also, we classified acoustic and visual cues separately to assess their predictive power independently.

3.3 NDL-based classification of [ba, da, ga]

Before we inspect the predictions for McGurk stimuli, we first inspect how well the two models predict the perceptemes [ba, da, ga] for consistent stimuli. We first turn our attention to the results of NDL.

A Spearman’s rank-correlation indicates that increasing the *standard deviation* of the distribution around the mean increases the overlap between the [ba] and [ga] category ($\rho = 0.9$) and between the [ba] and [da] category ($\rho = 0.53$). Given that overlap was calculated on the basis of F1 values and the horizontal distance, it is not surprising that the effect was only minimal on the overlap between the [da] and [ga] category ($\rho = 0.11$).

Consequently, we can observe that the proportions of correctly classified perceptemes is lower, when the *standard deviation* within each phonemic category is greater (Figure 3,a). This effect can be observed for the [ba] category when both modalities are taken into account. When the modalities

are split, proportions dropped to around 0.2. This effect is surprising, because $\llbracket \text{ba} \rrbracket$ has one acoustic and one visual cue that stands apart from the other two phonemic categories: the F2 cues and vertical distance cues of $\llbracket \text{ba} \rrbracket$ are separated from the cues for the two other categories (cf. Figure 1, b&e).

Turning to the $\llbracket \text{da} \rrbracket$ and $\llbracket \text{ba} \rrbracket$ category (cf. Figure 1, c&d), we observe that proportion of correctly classified items across *standard deviation* is similar in the $\llbracket \text{da} \rrbracket$ and $\llbracket \text{ga} \rrbracket$ condition independently of whether the modalities were split or not.

The proportion of correctly identified items across the *relative frequency* of each percepteme ranged between 0.4 and 1 for the $\llbracket \text{ba} \rrbracket$ category, between 0.4 and 1 for the $\llbracket \text{da} \rrbracket$ category and 0 and 1 for the $\llbracket \text{ga} \rrbracket$ category (Figure 3, i-k), with higher proportions for greater *relative frequencies*. When the two modalities were split, $\llbracket \text{ba} \rrbracket$ categorizations dropped.

Overall, we can conclude that NDL is well capable to classify the three perceptemes $\llbracket \text{ba} \rrbracket$, $\llbracket \text{da} \rrbracket$, $\llbracket \text{ga} \rrbracket$ on the basis of the cues from both modalities. When the two modalities are split, NDL is capable to classify $\llbracket \text{da} \rrbracket$ and $\llbracket \text{ga} \rrbracket$ with a high accuracy but fails to classify $\llbracket \text{ba} \rrbracket$. Why this is the case is beyond the scope of the present paper.

3.4 Modeling with Exemplar theory

In this section we present a computational model based on exemplar-theoretic principles. Exemplar theory is based on the idea that speech items are represented in the mental lexicon as individual, detailed or even fully specified *memory traces*, *episodes* or *exemplars*⁴. This basic idea has already been formulated early, for example by Paul (1880) or Semon (1909).

However, exemplar theory as it is applied today in phonetics and phonology, has its origins in much later work in cognitive psychology. Psychological studies by Rosch and her colleagues during the 1970s have shown that natural categories have no well-defined boundaries. Instead, individual instances are judged with varying degrees as being prototypical (good) exemplars of a category (Medin & Schaffer, 1978; Mervis & Rosch, 1981; Rosch, 1998). A central question in that line of research was how individual stimuli are categorized (e.g. individual perceived objects, events, etc.). More specifically, the problem that psychologists faced was to develop a formal model that could explain how categories can be learned based on individual concrete

⁴A note on terminology: The terms *memory trace*, *episode* or *exemplar* are used interchangeably in the literature. In the remainder of this paper we use the term *exemplar* for an individual stored speech item and refer to the set of stored exemplars including their associations to linguistic labels as the (*mental*) *lexicon*.

instances of experience and how they are represented cognitively and stored in memory (cf. Hintzman, 1984, 1986; Nosofsky, 1988).

An early psychological exemplar-theoretic simulation model is “MINERVA 2” which was developed by Hintzman (1984, 1986). In MINERVA 2, the long-term memory is modelled as “a vast collection of episodic memory traces, each of which is a record of an event of experience”. Hintzman (1984) emphasizes that the long-term memory contains traces of all experiences, even if they are very similar to previous experiences. Stimuli (i.e. novel experiences or events, called *probes* by Goldinger, 1998) are represented within that framework as n -dimensional vectors $x \in \{-1, 0, +1\}^n$ which represent n features. Values of -1 or 1 encode absence or presence of a feature, and a value of 0 was used to encode “irrelevant” features in stimuli or features which have been forgotten or which have never been stored in long-term memory (Hintzman, 1984). The model proposed by Hintzman (1986) assumes that there is a large set of “primitive properties” which “are not acquired by experience”. Crucially, every new experience leaves behind a new memory trace or *exemplar*. Frequency of occurrence associated with different types (e.g. categories) is reflected implicitly by the number of exemplars in memory associated with that given type. We do not include the MINERVA 2 model in the following discussion, because we consider its assumption about a feature space defined by (probably innate) “primitive properties” encoded as 1 , 0 or -1 , not unlike *distinctive features* in phonology, as cognitively implausible. Although we do not directly implement and test MINERVA 2 (e.g. in section A), many basic concepts developed by Hintzman (1984, 1986) are inherent to later exemplar models.

However, later exemplar models, e.g. by Lacerda (1995); Johnson (1997); Wade et al. (2010); Duran et al. (2013); Duran & Lewandowski (accepted), employ a different feature representation: as in MINERVA 2, features correspond to individual dimensions of an n -dimensional representation space. Their values, however, are real numbers. Exemplar-theoretic models in phonetics and phonology share the following basic set of assumptions:

- Perceived speech items are stored in memory. These stored items are referred to as *memory traces*, *episodes* or *exemplars*.
- Categorisation in the perception of a novel speech item is based on *similarity* to the collection of stored exemplars.
- Novel experiences (can) create new exemplars in memory, thus constantly changing the categorisation system.

Figure 4: *Exemplar clouds. Left: Illustration of the McGurk effect in an exemplar-theoretic model. Points represent exemplars within a multi-modal feature space. Red dashed lines indicate the distances between an inconsistent stimulus (XX) and the centers of the three clusters corresponding to the three syllables [ba], [da], [ga]. Right: Data set (gray points and density contours) and grid of stimuli (XX) for the preliminary simulations (see section A for details).*

3.5 Formal specification and implementation of Exemplar theory

We implemented a basic exemplar model based on the above mentioned principles. In order to minimize assumptions about memory representation of speech items and their organization, we assume that exemplars are represented as individual points within a multi-dimensional feature space. This space may comprise auditory, visual and other kinds of information. Formally, exemplars are represented as point-vectors in an n -dimensional real space $x \in \mathbb{R}^n$, where each dimension is associated with an individual feature. We refer to the collection of all stored exemplars as the *mental lexicon*, or short, the *memory* (formally a set $\mathcal{M}$ of vectors). The model itself is agnostic to the dimensionality of the exemplar space and to the specific associations between its dimensions and phonetic features. We set aside the question of exactly *which* mental features are employed for exemplar representations in memory. We also do not address the alleged head-filling-up problem⁵ (Johnson, 1997) in this present study. Instead, we simply assume that *all* individual instances of perceived (or produced) items are stored as

⁵The “head-filling-up problem” is the intuitive assumption that we cannot store everything we perceive in memory.

419 exemplars in memory.

420 In addition to the multi-dimensional feature representation, we assume
 421 that exemplars are labelled with category labels. Goldinger (1996), for
 422 example, notes that exemplars are not just simple “perceptual analogues
 423 that are totally defined by stimulus properties” but “complex perceptual-
 424 cognitive objects, jointly specified by perceptual forms and linguistic func-
 425 tions”. Formally, categories are represented by sets of exemplar vectors
 426 $\mathcal{C}_A = \{x : “x \text{ is labelled } A”\}$. For the current implementation, we assume
 427 that each exemplar is associated with no more than one category⁶.

428 Similarity is defined as a function which maps a pair of exemplars onto
 429 a non-negative real value⁷. Often, similarity computations incorporate the
 430 distance between the two exemplar vectors in the underlying feature space.
 431 In the case of n -dimensional exemplars in real space, $x \in \mathbb{R}^n$, the most
 432 common distance metric is the Euclidean distance between two vectors x_p
 433 and x_q (equation 1).

$$d(x_p, x_q) = \sqrt{\sum_{i=1}^n (x_p[i] - x_q[i])^2} \quad (1)$$

434 We use a representation of exemplars as individual data points in our im-
 435 plementation of the exemplar model without any reference to their sequential
 436 context within the speech stream or other levels of linguistic representation
 437 (compare, for example, Wade et al., 2010; Walsh et al., 2010). Mathemati-
 438 cally, this representation is discrete in so far as each exemplar is represented
 439 by one specific vector, i.e. an element within the set of all exemplars. Note
 440 that often, a mathematically more abstract representation of exemplar col-
 441 lections is employed, like probability distributions or connectionist models /
 442 artificial neural networks (Johnson, 1997; Pierrehumbert, 2001, 2016). We
 443 favoured a straight-forward representation for the present study, which in-
 444 troduces less assumptions about mental organisation and abstraction. A
 445 caveat with this approach, however, is that albeit the required computa-
 446 tions are rather simple, they need to be carried out in very large numbers.

⁶In other words, we assume that the categories $\mathcal{C} \subseteq \mathcal{M}$ form a partition of the set of all stored exemplars, i.e. an exemplar (x) can be an element of no more than one category ($x \in \mathcal{C}$). We also assume for the sake of the presented model for the McGurk effect that there are at least three categories in $\mathcal{M}$. Thus, the case $\mathcal{C} = \mathcal{M}$ does not apply here. More realistically, however, categories could be represented by fuzzy sets, with varying degrees of membership of its elements. The question, how this would change the exemplar model is left open for future work.

⁷Depending on the specific model, the similarity value can be interpreted as *activation* (in cognitive terms) of an exemplar in memory. We do not strictly distinguish these two notions with respect to our exemplar model and use term similarity.

447 The formal exemplar-theoretic model requires the pair-wise computation of
 448 the Euclidean distance between the stimulus and each exemplar in memory.
 449 Our goal is to keep our implementation of the model at the *computational*
 450 *level*, i.e. at an abstract level describing *what* cognitive functions are exe-
 451 cuted instead of more specific details about *how* these mechanisms may be
 452 implemented, e.g. at a neural level (cf. Marr (2000)). In our simulations,
 453 the straight-forward implementation of a set of exemplars is to store each
 454 exemplar as an individual object rather than employing some kind of proba-
 455 bility distributions Lacerda (1995) or neural network representations Johnson
 456 (1997).

457 We decided to use the generalized weighted neighborhood classification
 458 as a classification method. The definition and our considerations for this
 459 decision can be found in the appendix (A), in which we demonstrate how it
 460 outperforms other classification algorithms.

461 3.6 Data

462 The data we use for the exemplar-based simulations is created in analogy to
 463 the data we use with the NDL-based simulations, described in section 3.2.
 464 Due to the different nature of exemplar-based classification, we do not com-
 465 pute discretized cues. All individual data points with their real-valued artic-
 466 ulatory and acoustic features are used as *exemplars* within the simulations⁸.
 467 The empirical standard deviations were multiplied by factors 1 and 2. With
 468 these parameters, we generate a total of 4,394 different data sets of vari-
 469 ous sizes, with 1,723 unique combinations of relative frequencies of the three
 470 percepteme categories [ba, da, ga] (cf. Figure 2).

471 Since there is no batch learning in exemplar-based classification, we im-
 472 plemented a 100-fold cross-validation method which partitions the data into
 473 a large set of memory exemplars and a small set of classification stimuli (i.e.,
 474 in machine-learning terms, *train* and *test* sets, respectively).

475 3.7 Exemplar-based classification of [ba, da, ga]

476 The following description of the exemplar model of the McGurk effect paral-
 477 lels the description of the NDL-based classification (section 3.3) as closely as
 478 possible, in order to allow for a direct comparison of the two models. Again,
 479 we first investigate the model performance on consistent [ba, da, ga] stimuli
 480 using numerical simulations. The results are visualized in Figure 5 (page 16).

⁸We use the **pdist** package to compute pairwise distances between exemplars in the data sets (Wong, 2013).

Figure 5: *Dotplots (with lines) illustrating the average prediction accuracy in the Exemplar simulation depending on different standard deviations around the mean (top panel), and the proportion of the number of [ba, da, ga] stimuli in the training data set (bottom panel). Line types represents data in the acoustic modality (A), the visual modality (V) and both joined together (AV).*

481 We can observe that the proportions of correctly classified items in the
 482 entire data set decrease, the larger the *standard deviation* within each cate-
 483 gory (Figure 5, top panel). The only exception to this pattern is the results
 484 for [da] with monomodal acoustic data (A), which has slightly higher mean
 485 accuracy with a standard deviation of 2, though this difference is not sig-
 486 nificant. The exemplar-based classification seems to benefit most from the
 487 visual articulatory data.

488 Given this general observation, it is not surprising that the mean classi-
 489 fication accuracy on monomodal visual (V) and bimodal joined data (AV)
 490 is highest for [ba] over most of the range of the *relative frequency* of each
 491 phonemic category.

492 4 Classification of McGurk stimuli

493 In the previous section, we have shown that NDL and Exemplar-theory per-
 494 form very well in classifying multi-modal cues as [ba, da, ga]. In the following
 495 section, we turn our attention to the models' prediction for the McGurk stim-
 496 uli, i.e. stimuli that were created by combining the acoustic features of [ba]
 497 with the visual features of [ga] (called McGurk test from now on). We first
 498 discuss the results for NDL followed by the results for Exemplar theory.

499 To investigate the prediction accuracy, we used a Generalized Additive

Model (GAM, package **mgcv**, Version 1.8-23, Wood, 2006). GAMs allow to investigate non-linear functional relations between dependent and multiple independent variables using smooths on the basis of thin plate regression splines and tensor product smooths on the basis of cubic regression splines (see also Tomaschek et al. (2018c); Wieling et al. (2016) for their application in articulatory). We used the ‘betar’ family for proportions data on a scale between 0 and 1 (Wood et al., 2016).

In total, we fitted six GAM models to investigate how the proportions of [da] in the McGurk tests can be predicted by the *relative frequencies* of each phonemic category, i.e. three models for each pair-wise combination of [ba, da, ga] ([ba]-[da], [ga]-[da] and [ga]-[ba]) for each theoretical background. Each model contained a tensor product smooth fitting a two-way interaction between the *relative frequencies* for [ba] and [da], [ga] and [da], and [ga] and [ba], in addition to a fixed effects for *standard deviation*⁹. All of the models had significant non-linear tensor product smooths ($p < 0.0001$). The model summaries can be inspected in the Supplementary Materials downloadable from <https://osf.io/v76b2>. Across all NDL models, increasing the standard deviation significantly increased the probability of [da] ($\beta = 0.9$, $sd = 0.01$, $t = 79.2$, $p < 0.001$). The increase in standard deviation around a category’s mean probably results in a greater overlap between the categories, i.e. a greater uncertainty which cues support which percepteme. Simultaneously, more cues are learned in an inconsistent way, i.e. cues for [ba] are also learned to indicate [da]. Consequently, a stronger McGurk effect should be expected. This finding could explain, why bilinguals show a larger probability to perceive [da] in the McGurk test than monolinguals (Marian et al., 2018). They are simply faced with a larger distribution of cues to the respective categories.

Figure 6 illustrates the predictions for the interactions between the *relative frequencies* of the phonemic categories in the McGurk tests. The results on the basis of the NDL simulations are represented in the top row, those for Exemplar-theory in the bottom row. Each column represents a model with a tensor product smooth for one pair-wise combination of [ba, da, ga] *relative frequencies*. The x-axis and the y-axis of each plot represent the relative frequency of one of the phonemic categories in the entire training set. Contour lines and colors represent the interaction’s partial effect of proportions of [da] in the McGurk test, with blue colors representing lower accuracy and yellow colors representing higher accuracy. The black circles represent the

⁹It is not possible to fit a model which contains the *relative frequencies* of all three categories, because the data is fully collinear and results become uninterpretable (Tomaschek et al., 2018b). The reason for this is that the relative frequency of one category is the result of 1 minus the summed relative frequency of the other two categories.

Figure 6: *Predictions for the McGurk probability depending on the proportion of [ba], [da], [ga] in the lexicon, simulated with Exemplar-based algorithms. a-c: results based on NDL. d-f: results based on Exemplar-theoretic algorithms.*

proportions of [ba, da, ga] in the discussed languages (section 2).

4.1 Results for NDL

We observe greater proportions of [da] in the McGurk test as the *relative frequencies* of [ba] and [da] increase (y-axis and x-axis in Figure 6 (a), indicated by more green and yellow colors from left to right). This result supports our hypothesis that the McGurk effect arises when the frequency of occurrence of [da] is relatively high.

The two predictors interact insofar as when the *relative frequency of [da]* decreases, the proportion of [da] in the McGurk test remains high when the *relative frequency of [ba]* increases. This effect creates an area of high proportions of [da] at the diagonal of the plot.

Implicitly, this diagonal encodes a low *relative frequency* of [ga]. Thus, it seems that when the *relative frequency* of [ga] is low, the proportion of [da] in the McGurk test is high and vice versa. This is supported in Figure 6 (b), where the effect on proportions is mirrored in the vertical plane. The propor-

tion of [da] instances is larger when the *relative frequency* of both categories, [da] and [ga] decreases. They interact insofar as very high proportions of [da] arise only when the *relative frequency* of [ga] is low.

Turning our attention to the interaction between the *relative frequency* of [ba] and the *relative frequency* of [ga] (Figure 6, c), we can observe that the proportion of [da] in the McGurk effect increases, when both frequencies decrease. They interact insofar as high proportions emerge for low *relative frequency* of [ga].

The circles in the plots represent the languages we have discussed above. Apart from Polish, those languages for which the McGurk effect has been attested can be found in the area of high proportions of [da] instances in the McGurk test. Chinese and Japanese stand apart from that group and are located in an area for which our simulations predict smaller proportions of [da] instances in the McGurk effect.

4.2 Results for the Exemplar-theoretic model

The plots in Figure 6 (d-f) illustrate the results for the Exemplar-theoretic model. The direction of the effects across all plots is consistent with the findings in the NDL simulations. The smaller the *relative frequency* of [ga], the larger the proportion of [da] in the McGurk test. This effect, like above, is further modulated by the relative frequencies of [ba] and [da]. The crucial difference between NDL and Exemplar-theory is that the probability of obtaining [da] in the McGurk test is drastically larger. Whereas in the NDL simulation the languages, for which the McGurk effect was attested, were located in a probability area ranging between 0.15 to 0.3, the Exemplar-based simulations allocate them in an area of 0.5 to 0.9. Thus, the Exemplar-based simulations make more concise predictions about the effect than NDL.

5 Discussion

In the present paper, we investigated the origins of the McGurk effect using two computational models of cognitive processing, Naïve Discriminative Learning (NDL Baayen et al., 2011), a computational algorithm based on classical conditioning that simulates error-driven learning (Rescorla & Wagner, 1972; Ramscar et al., 2010, 2013) and Exemplar Theory, which assumes that individual instances of perceived events are stored in memory as fully specified exemplars (Johnson, 1997; Goldinger, 1998; Pierrehumbert, 2001). On the basis of these models, we hypothesized that the McGurk effect arises due to the distributional properties of the acoustic and visual cues of the

labial, coronal, and velar consonants in each language. Concretely, we hypothesized that if a language has a large relative frequency of [da], a stronger McGurk effect should be observed in that language. Using articulographic and acoustic recordings of real German words beginning with [ba, da, ga], we found support for our hypothesis. Moreover, our simulations revealed that it is actually the interplay between the relative frequencies of [ba, da, ga]. The crucial parameter for whether or not a McGurk effect arises is the relative frequency of [ga]: the larger it is, the less of a McGurk effect should arise. A comparison of the simulated predictions of the McGurk effect based on the relative frequencies of [ba, da, ga] with the distribution of [ba, da, ga] in English, Swedish, Turkish, Italian, and German, for which a McGurk effect was observed, and in Chinese, Cantonese and Japanese, for which none was observed, supported the current results.

The question, however, arises, why Polish can be found among the “McGurk languages” in our simulations although no effect has been attested empirically. The only one study, which the present authors could find, investigating the McGurk effect in Polish speakers reports that only 4.5% of the responses were consistent with the fused perceptual outcome in the entire testing set (Majewski, 2008). The testing set also contained stimuli that joined voiced and voiceless consonants as well as stimuli that joined visual [ba] with acoustic [ga]. Interestingly, the percentage of “true” McGurk items (i.e. articulatory [ga] + acoustic [ba]) in the testing set was also roughly 4.7%. It would be a big coincidence that this percentage is very similar to the number of fused perceptual outcomes. Given this consideration and our results we claim that speakers of Polish actually do experience the McGurk effect.

Furthermore, our results predict that native speakers of Cantonese, Chinese and Japanese do not perceive [da] in the McGurk test because of the distributional properties of [ga] in the respective languages. This finding is at odds with extant explanations for these languages. Sekiyama (1997) explains the result for these languages by means of the ‘face-avoidance hypothesis’, according to which speakers of these languages avoid to watch at the faces of their interlocutors, thus need to rely primarily on the acoustic cues to process speech as they have not learned to use the visual cues. While the present results do not contradict the face-avoidance hypothesis, they add another dimension to these findings, namely the distributional properties of sounds in a language. Furthermore, given rise of TV ownership in China from 0 to 300 per 1000 persons (Wang et al., 2002) since the 1980 and a high ownership in Japan, it is possible that native speakers of these languages still get a high ‘dose’ of visual cues in spite of a cultural doctrine not to look in other people’s faces, even though Senju et al. (2013) reports that Japanese native speakers focus more on the eyes of the interlocutor while English native speakers focus

629 more on their mouths (see also (Blais et al., 2008)).

630 Another question arises why NDL predicts lower probability of [da] in the
631 McGurk effect than Exemplar theory. One potential outcome is that NDL,
632 according to error-driven learning, assumes cue competition. In other words,
633 during learning, the connection weight between (acoustic and visual) cues
634 and perceptual outcomes is adjusted not only on the basis of co-occurrences,
635 but also on the basis of non-co-occurrences as well as on the basis of the
636 amount of cues. This introduces a larger uncertainty about the relation
637 between cues and outcomes. By contrast, Exemplar Theory does not know a
638 procedural training with cue competition. ‘Learning’ in Exemplar Theory is
639 the accumulation of instances and its acoustic and visual cues, independently
640 of to which other categories their cues are related. Thus, less uncertainty is
641 introduced to the system. For the current simulation, Exemplar Theory
642 outperforms the NDL simulation. However, the probability to perceive [da]
643 in real speech ranges between $\sim 95\%$ and $\sim 5\%$, depending on the speaker
644 (see Nath & Beauchamp (2011) for a neural explanation for inter-perceiver
645 variability) even in languages for which the McGurk effect is attested (Paré
646 et al., 2003; Magnotti & Beauchamp, 2015; Traunmüller, 2009; Schwartz
647 et al., 1998; Schwartz, 2010).

648 It should also be mentioned that the current simulations used McGurk
649 stimuli which were constructed by combining the visual cues from every [ga]
650 instance with the acoustic cues from every [ba] instance based on real words.
651 This means that potential combinations might have arisen which simply will
652 not result in [da]. This is not the standard procedure in traditional McGurk
653 test. Rather, studies investigating the McGurk effect use idealized articu-
654 lations of demisyllables such as [ba, da, ga]. The question thus arises, how
655 participants would perform if they were faced with true speech. Whereas the
656 simulation on the basis of Exemplar Theory predicts that such experiments
657 should still acquire high probabilities of [da] in the McGurk test, NDL is less
658 optimistic.

659 References

- 660 Aloufy, S., M. Lapidot & M. Myslobodsky. 1996. Differences in susceptibility
661 to the “Blending Illusion” among native Hebrew and English speakers.
662 Brain and Language 53(1). 51–57. doi:10.1006/brln.1996.0036.
- 663 Altmann, G. 1980. Prolegomena to menzerath’s law. Glottometrika 2. 1–10.
- 664 Arnold, D. & F. Tomaschek. 2016. The karl eberhards corpus of spon-
665 taneously spoken southern german in dialogues - audio and articulatory

- 666 recordings. In Christoph Draxler & Felicitas Kleber (eds.), Tagungsband
667 der 12. tagung phonetik und phonologie im deutschsprachigen raum. p
668 und p12. 12. tagung phonetik und phonologie im deutschsprachigen raum,
669 12. - 14. oktober 2016, 9–11. Ludwig-Maximilians-Universität München.
670 <https://epub.ub.uni-muenchen.de/29405/>.
- 671 Arnold, D., F. Tomaschek, F. Lopez, K. Sering, M. Ramscar & R. H. Baayen.
672 2017. Words from spontaneous conversational speech can be recognized
673 with human-like accuracy by an error-driven learning algorithm that dis-
674 criminate between meanings straight from smart acoustic features, by-
675 passing the phoneme as recognition unit. PLOS ONE .
- 676 Arnon, I. & M. Ramscar. 2012. Granularity and the acquisition of grammat-
677 ical gender: How order-of-acquisition affects what gets learned. Cognition
678 122(3). 292–305.
- 679 Baayen, R. H., P. Milin, D. F. Durdevic, P. Hendrix & M. Marelli. 2011. An
680 amorphous model for morphological processing in visual comprehension
681 based on naive discriminative learning. Psychological review 118(3). 438–
682 481.
- 683 Balota, D.A., M.J. Yap, M.J. Cortese, K.A. Hutchison, B. Kessler, B. Loftis,
684 J.H. Neely, D.L. Nelson, G.B. Simpson & R. Treiman. 2011. The english
685 lexicon project. Behavior Research Methods 39. 445–459.
- 686 Bccwj-Consortium. 2016. Modern japanese written balance corpus (bccwj).
687 <http://pj.ninjal.ac.jp/corpus.center/bccwj/>.
- 688 Blais, Caroline, Rachael E Jack, Christoph Scheepers, Daniel Fiset & Roberto
689 Caldara. 2008. Culture shapes how we look at faces. PloS one 3(8). e3022.
- 690 Boersma, P. & P. Weenink. 2015. Praat: doing phonetics by computer [com-
691 puter program], version 5.3.41, retrieved from <http://www.praat.org/> .
- 692 Bovo, Roberto, Andrea Ciorba, Silvano Prosser & Alessandro Martini. 2009.
693 The mcgurk phenomenon in italian listeners. Acta Otorhinolaryngologica
694 Italica 29(4). 203.
- 695 Chen, Yuchun & Valerie Hazan. 2007. Language effects on
696 the degree of visual influence in audiovisual speech per-
697 ception. In Proceedings of the 16th icphs, Saarbrücken.
698 <http://www.icphs2007.de/conference/Papers/1271/index.html>. ID
699 1271.

- 700 Clayards, Meghan, Michael K Tanenhaus, Richard N Aslin & Robert A Ja-
701 cobs. 2008. Perception of speech reflects optimal use of probabilistic speech
702 cues. Cognition 108(3). 804–809.
- 703 Duran, Daniel. 2013. Computer simulation experiments in phonetics and
704 phonology: simulation technology in linguistic research on human speech:
705 Universität Stuttgart Doctoral dissertation. doi:10.18419/opus-3202.
- 706 Duran, Daniel. 2015. Perceptual magnets in different neighborhoods. In
707 A. Leemann, M.-J. Kolly, S. Schmid & V. Dellwo (eds.), Phonetics and
708 Phonology: Studies from German speaking Europe, 225–237. Frankfurt
709 am Main / Bern: Peter Lang.
- 710 Duran, Daniel, Jagoda Bruni & Grzegorz Dogil. 2013. Modeling multi-
711 modal factors in speech production with the Context Sequence Model.
712 In Elektronische Sprachsignalverarbeitung 2013, 86–92. TUDpress.
- 713 Duran, Daniel & Natalie Lewandowski. accepted. Cognitive factors in speech
714 production and perception: a socio-cognitive model of phonetic conver-
715 gence. In CALS 2018 proceedings, .
- 716 Fisher, Cletus G. 1968. Confusions among visually perceived conso-
717 nants. Journal of Speech and Hearing Research 11(4). 796–804. doi:
718 10.1044/jshr.1104.796.
- 719 Fisher, Cletus Graydon. 1963. Confusions within six types
720 of phonemes in an oral-visual system of communication.
721 Columbus: The Ohio State University Doctoral dissertation.
722 http://rave.ohiolink.edu/etdc/view?acc_num=osu1486553441673462.
- 723 Fox, John & Sanford Weisberg. 2019. An R companion to applied regression.
724 Thousand Oaks, California: Sage Publications, Inc third edition edn.
- 725 Fuster-Duran, Angela. 1996. Perception of conflicting audio-visual speech:
726 an examination across spanish and german. In David G. Stork & Mar-
727 cus E. Hennecke (eds.), Speechreading by humans and machines: Models,
728 systems, and applications, 135–143. Berlin, Heidelberg: Springer Berlin
729 Heidelberg.
- 730 Gay, Thomas. 1978. Effect of speaking rate on vowel formant movements.
731 The Journal of the Acoustical Society of America 63(1). 223–230.
- 732 Goldinger, Stephen D. 1996. Words and voices: Episodic traces in spo-
733 ken word identification and recognition memory. Journal of Experimental
734 Psychology: Learning, Memory, and Cognition 22(5). 1166–1183.

- Goldinger, Stephen D. 1998. Echoes of echoes? An episodic theory of lexical access. Psychological Review 105(2). 251–279.
- Hertrich, I., K. Mathiak, W. Lutzenberger & H. Ackermann. 2009. Time course of early audiovisual interactions during speech and nonspeech central auditory processing: A magnetoencephalography study. Journal of cognitive neuroscience 21 2. 259–74.
- Hintzman, Douglas L. 1984. MINERVA 2: A simulation model of human memory. Behavior Research Methods, Instruments, & Computers 16(2). 96–101. doi:10.3758/BF03202365.
- Hintzman, Douglas L. 1986. “Schema abstraction” in a multiple-trace memory model. Psychological Review 93(4). 411–428. doi:10.1037/0033-295X.93.4.411.
- Hirata, Yukari. 2004. Effects of speaking rate on the vowel length distinction in Japanese. Journal of Phonetics 32(4). 565 – 589. doi: <http://dx.doi.org/10.1016/j.wocn.2004.02.004>.
- Johnson, Keith. 1997. Speech perception without speaker normalization: An exemplar model. In Keith Johnson & John Mullenix (eds.), Talker variability in speech processing, 145–165. Academic Press.
- Kamin, Leon J. 1967. Predictability, surprise, attention, and conditioning .
- Kleber, Felicitas. 2017. Complementary length in vowel–consonant sequences: Acoustic and perceptual evidence for a sound change in progress in Bavarian German. Journal of the International Phonetic Association 1–22. doi: 10.1017/S0025100317000238.
- Lacerda, Francisco. 1995. The perceptual-magnet effect: An emergent consequence of exemplar-based phonetic memory. In K. Elenius & P. Branderyd (eds.), Proceedings of the 13th international congress of phonetic sciences, vol. 2, 140–147. Stockholm. http://www.ling.su.se/staff/frasse/Lacerda_ICPhS95a.pdf.
- Lison, Pierre & Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles. In Proceedings of the 10th international conference on language resources and evaluation (LREC 2016), <http://stp.lingfil.uu.se/~joerg/paper/opensubs2016.pdf>. Corpus data based on <http://www.opensubtitles.org/>.

- 763 Luke, Kang-Kwong & May LY Wong. 2015. The hong kong cantonese corpus:
764 design and uses. Journal of Chinese Linguistics 25(2015). 309–330.
- 765 MacDonald, John. 2018. Hearing lips and seeing voices: the origins and develop-
766 ment of the ‘mcgurk effect’ and reflections on audio–visual speech perception
767 over the last 40 years. Multisensory Research 31(1-2). 7–18.
- 768 Magen, H. S. 1997. The extent of vowel-to-vowel coarticulation in english.
769 Journal of Phonetics 25. 187–205.
- 770 Magnotti, J. & M. Beauchamp. 2015. The noisy encoding of disparity model
771 of the mcgurk effect. Psychonomic Bulletin and Review 22(3). 701–709.
- 772 Majewski, Wojciech. 2008. McGurk effect in polish listeners. Archives of
773 Acoustics 33(4). 447–454.
- 774 Manning, Christopher D., Prabhakar Raghavan & Hinrich Schütze. 2008.
775 Introduction to Information Retrieval. New York: Cambridge Univer-
776 sity Press online edition edn. [http://nlp.stanford.edu/IR-book/information-](http://nlp.stanford.edu/IR-book/information-retrieval-book.html)
777 [retrieval-book.html](http://nlp.stanford.edu/IR-book/information-retrieval-book.html).
- 778 Marian, Viorica, Sayuri Hayakawa, Tuan Lam & Scott Schroeder. 2018. Lan-
779 guage experience changes audiovisual perception. Brain Sciences 8(5). 85.
780 doi:10.3390/brainsci8050085.
- 781 Marr, David. 2000. Vision: a computational investigation into the human
782 representation and processing of visual information. New York: Freeman
783 14th edn. OCLC: 248012301.
- 784 Massaro, D. & M. Cohen. 1995. Cross-linguistic comparison in the integration
785 of visual and auditory speech. Memory and Cognition 23(1). 113–131.
- 786 McGurk, Harry & John MacDonald. 1976. Hearing lips and seeing voices.
787 Nature 264(5588). 746–748. doi:10.1038/264746a0.
- 788 Medin, Douglas L. & Marguerite M. Schaffer. 1978. Context theory of clas-
789 sification learning. Psychological Review 85(3). 207–238. doi:10.1037/0033-
790 295X.85.3.207.
- 791 Mervis, Carolyn B. & Eleanor Rosch. 1981. Categorization of Nat-
792 ural Objects. Annual Review of Psychology 32(1). 89–115. doi:
793 10.1146/annurev.ps.32.020181.000513.

- Mildner, Vesna & Arnalda Dobrić. 2015. Reconsidering the McGurk Effect. In Proceedings of the 18th International Congress of Phonetic Sciences, Glasgow, UK: University of Glasgow.
- Nath, A. R. & M.S. Beauchamp. 2011. A neural basis for interindividual differences in the mcgurk effect, a multisensory speech illusion. NeuroImage 59(1). 781–787.
- Nixon, Jessie S. 2018. Effective acoustic cue learning is not just statistical, it is discriminative. Proc. Interspeech 2018 1447–1451.
- Nixon, Jessie S & Catherine T Best. 2018. Acoustic cue variability affects eye movement behaviour during non-native speech perception. In Proc. 9th international conference on speech prosody 2018, 493–497.
- Nixon, Jessie S, Jacolien van Rij, Peggy Mok, Harald Baayen & Yiya Chen. 2015. Eye movements reflect acoustic cue informativity and statistical noise. Experimental Linguistics 50.
- Nixon, Jessie Sophia et al. 2014. Sound of mind: electrophysiological and behavioural evidence for the role of context, variation and informativity in human speech processing. Leiden University Centre.
- Nosofsky, Robert M. 1988. Exemplar-based accounts of relations between classification, recognition, and typicality. Journal of Experimental Psychology: Learning, Memory, and Cognition 14(4). 700–708. doi:10.1037/0278-7393.14.4.700.
- Öhman, S.E.G. 1966. Coarticulation in vcv utterances: Spectrographic measurements. Journal of the Acoustical Society of America 39(151). 151–168.
- Paré, Martin, Rebecca C. Richler, Martin ten Hove & K. G. Munhall. 2003. Gaze behavior in audiovisual speech perception: The influence of ocular fixations on the McGurk effect. Perception & Psychophysics 65(4). 553–567. doi:10.3758/BF03194582. <http://link.springer.com/10.3758/BF03194582>.
- Paul, Hermann. 1880. Principien der Sprachgeschichte. Halle: Max Niemeyer Verlag.
- Pierrehumbert, J. B. 2001. Exemplar dynamics: Word frequency, lenition and contrast. In J Bybee & P. Hopper (eds.), Frequency and the emergence of linguistic structure, 137–157. Amsterdam, Netherlands: John Benjamins Publishing Company.

- Pierrehumbert, Janet B. 2016. Phonological Representation: Beyond Abstract Versus Episodic. Annual Review of Linguistics 2(1). 33–52. doi: 10.1146/annurev-linguistics-030514-125050.
- Ramscar, M., M. Dye & S. McCauley. 2013. Error and expectation in language learning: The curious absence of ‘mouses’ in adult speech. Language 89(4). 760–793.
- Ramscar, M., D. Yarlett, M. Dye, K. Denny & K. Thorpe. 2010. The effects of feature-label-order and their implications for symbolic learning. Cognitive Science 34(6). 909–957.
- Rescorla, R. 1988. Pavlovian conditioning - it’s not what you think it is. American Psychologist 43(3). 151–160.
- Rescorla, R. & A. Wagner. 1972. A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W.F. Prokasy (eds.), Classical conditioning ii: Current research and theory, 64–69. Appleton Century Crofts, New York.
- Rosch, Eleanor. 1998. Principles of Categorization. In George Mather, Frans Verstraten & Stuart Anstis (eds.), The Motion Aftereffect, 251–270. The MIT Press.
- Rosenblatt, Frank. 1962. Principles of neurodynamics. Spartan Book.
- Rosenblum, Lawrence D. 2008. Speech perception as a multimodal phenomenon. Current Directions in Psychological Science 17(6). 405–409.
- Schwartz, Jean-Luc. 2010. A reanalysis of mcgurk data suggests that audio-visual fusion in speech perception is subject-dependent. The Journal of the Acoustical Society of America 127(3). 1584–1594.
- Schwartz, Jean-Luc, Jordi Robert-Ribes & Pierre Escudier. 1998. Ten years after summerfield: A taxonomy of models for audio-visual fusion in speech perception. Hearing by eye II: Advances in the psychology of speechreading and auditory-visual speech 85–108.
- Sekiyama, K., L. D. Braid, K. Nishino, M. Hayashi & M. M. Tuyo. 1995. The mcgurk effect in japanese and american perceivers. In Proceedings of the XIIIth icphs, vol. 3, 214–217.
- Sekiyama, Kaoru. 1997. Cultural and linguistic factors in audio-visual speech processing: The mcgurk effect in chinese subjects.

- 860 *Perception & Psychophysics* 59(1). 73–80. doi:10.3758/BF03206849.
861 <http://dx.doi.org/10.3758/BF03206849>.
- 862 Sekiyama, Kaoru & Yoh’ichi Tohkura. 1991. McGurk effect in non-English
863 listeners: Few visual effects for Japanese subjects hearing Japanese syllables
864 of high auditory intelligibility. *The Journal of the Acoustical Society of*
865 *America* 90(4). 1797–1805. doi:10.1121/1.401660.
- 866 Semon, Richard. 1909. *Die Mnemischen Empfindungen in ihren Beziehungen*
867 *zu den Originalempfindungen*. Leipzig: Wilhelm Engelmann.
- 868 Senju, Atsushi, Angelina Vermetti, Yukiko Kikuchi, Hironori Akechi, Toshikazu
869 Hasegawa & Mark H Johnson. 2013. Cultural background modulates how we
870 look at other persons’ gaze. *International journal of behavioral development*
871 37(2). 131–136.
- 872 Sering, Konstantin, Marc Weitz, David-Elias Künstle & Lennart Schnei-
873 der. 2017. Pyndl: Naive discriminative learning in python. doi:
874 10.5281/zenodo.597964.
- 875 Shi, Lei, Thomas L. Griffiths, Naomi H. Feldman & Adam N. Sanborn.
876 2010. Exemplar models as a mechanism for performing Bayesian inference.
877 *Psychonomic Bulletin & Review* 17(4). 443–464. doi:10.3758/PBR.17.4.443.
- 878 Sumby, William H & Irwin Pollack. 1954. Visual contribution to speech in-
879 telligibility in noise. *The journal of the acoustical society of america* 26(2).
880 212–215.
- 881 Sun, Ching Chu. 2018. *Lexical processing in simplified chinese: an investigation*
882 *using a new large-scale lexical database*: Universität Tübingen dissertation.
- 883 Tomaschek, F., D. Arnold, Franziska Broeker & R. H. R. Baayen. 2018a. Lexical
884 frequency co-determines the speed-curvature relation in articulation. *Journal*
885 *of Phonetics* 68. 103–116.
- 886 Tomaschek, F., P. Hendrix & R. H. Baayen. 2018b. Strategies for managing
887 collinearity in multivariate linguistic data. *Journal of Phonetics* 71. 249–267.
- 888 Tomaschek, F. & A. Leeman. 2018. The size of the tongue movement area affects
889 the temporal coordination of consonants and vowels – a proof of concept
890 on investigating speech rhythm. *The Journal of the Acoustical Society of*
891 *America* 144(5). EL410–EL416.

- 892Tomaschek, F., B. V. Tucker, R. H. Baayen & M. Fasiolo. 2018c. Practice makes
893 perfect: The consequences of lexical proficiency for articulation. Linguistic
894 Vanguard 4(s2). 1–13.
- 895Tomaschek, F., B. V. Tucker, M. Wieling & R. H. Baayen. 2014. Vowel articu-
896 lation affected by word frequency. In Proceedings of the 10th issp, 425–428.
897 Cologne.
- 898Tomaschek, F., M. Wieling, D. Arnold & R. H. Baayen. 2013. Word frequency,
899 vowel length and vowel quality in speech production: An ema study of the
900 importance of experience. In Proceedings of the interspeech, Lyon.
- 901Traunmüller, Hartmut. 2009. Factors affecting visual influence on heard vowel
902 roundedness: Web experiments with swedes and turks. In Peter Branderud
903 & Hartmut Traunmüller (eds.), Proceedings FONETIK 2009: The XXIIth
904 Swedish phonetics conference, 166–171. Department of Linguistics, Stock-
905 holm University.
- 906Venables, W. N. & B. D. Ripley. 2002. Modern applied statistics with s. New
907 York: Springer 4th edn. <http://www.stats.ox.ac.uk/pub/MASS4>. ISBN 0-
908 387-95457-0.
- 909Wade, Travis, Grzegorz Dogil, Hinrich Schütze, Michael Walsh & Bernd
910 Möbius. 2010. Syllable frequency effects in a context-sensitive segment pro-
911 duction model. Journal of Phonetics 38(2). 227–239.
- 912Walsh, Michael, Bernd Möbius, Travis Wade & Hinrich Schütze. 2010. Multi-
913 level Exemplar Theory. Cognitive Science 34(4). 537–582. doi:10.1111/j.1551-
914 6709.2010.01099.x.
- 915Wang, Youfa, Carlos Monteiro & Barry M Popkin. 2002. Trends of obesity and
916 underweight in older children and adolescents in the united states, brazil,
917 china, and russia. The American journal of clinical nutrition 75(6). 971–977.
- 918Weirich, M. & S. Fuchs. 2006. Palatal morphology can influence speaker-specific
919 realizations of phonemic contrasts. Journal of Speech, Language, and Hearing
920 Research 56. 1894–1908.
- 921Wickham, Hadley. 2016. ggplot2: Elegant graphics for data analysis. Springer-
922 Verlag New York. <http://ggplot2.org>.
- 923Widrow, B. & M. E. Hoff. 1960. Adaptive switching circuits. 1960 WESCON
924 Convention Record Part IV 96–104.

- 925 Wieling, M., F. Tomaschek, D. Arnold, M. Tiede, Franziska Bröker, Samuel
 926 Thiele, Simon N. Wood & R. H. Baayen. 2016. Investigating dialectal differ-
 927 ences using articulography. Journal of Phonetics .
- 928 Wong, Jeffrey. 2013. pdist: Partitioned distance function. [https://CRAN.R-](https://CRAN.R-project.org/package=pdist)
 929 [project.org/package=pdist](https://CRAN.R-project.org/package=pdist). R package version 1.2.
- 930 Wood, S., N. Pya & B. Säfken. 2016. Smoothing parameter and model selection
 931 for general smooth models. Journal of the American Statistical Association
 932 111(516). 1548–1563. doi:10.1080/01621459.2016.1180986.
- 933 Wood, S. N. 2006. Generalized additive models: an introduction with r. Boca
 934 Raton, Florida, U. S. A: Chapman and Hall/CRC.
- 935 Yoshida, Katherine A, Ferran Pons, Jessica Maye & Janet F Werker. 2010.
 936 Distributional phonetic learning at 10 months of age. Infancy 15(4). 420–
 937 433.
- 938 Zhang, Juan, Yaxuan Meng, Catherine McBride, Xitao Fan & Zhen Yuan. 2018.
 939 Combining behavioral and ERP methodologies to investigate the differences
 940 between McGurk effects demonstrated by Cantonese and Mandarin speakers.
 941 Frontiers in Human Neuroscience 12.

942 **A Appendix: Considerations on similarity and** 943 **classification in Exemplar-based simulations**

944 Exemplar theory builds on the idea that input stimuli are compared with
 945 the set of stored exemplars and then classified according to their similarity
 946 to known items. The similarity function, however, can be implemented in
 947 different ways. The behavior of an exemplar model thus largely depends
 948 on the particular formalization of the classification function. In this section
 949 we give a brief overview and comparison of different exemplar-based classi-
 950 fications that have been proposed in seminal papers on exemplar-theoretic
 951 speech perception and illustrate their properties in numerical simulations.

952 There are three basic classification methods in exemplar-theoretic mod-
 953 els: (1) The first one is ε -neighborhood classification (also called radius-based
 954 classification) as proposed, for example, by Lacerda (1995) or Pierrehumbert
 955 (2001) (see also: Walsh et al., 2010). (2) The second basic classification
 956 method is distance-based classification as proposed, for example, by John-
 957 son (1997) (see also Duran (2015) for a variant implementation or Shi et al.

(2010) for a Bayesian model). (3) The third method is k -nearest neighbors classification (kNN). While the similarity function is monotonically decreasing over the feature space in distance-based classification, it is non-monotonic in neighborhood classification and kNN.

962 A.1 The “Lacerda classification”

963 The exemplar classification proposed by Lacerda (1995) is an example of
 964 an ε -neighborhood classification. In this approach, a stimulus item is com-
 965 pared with all exemplars within a specified radius ε in the multi-dimensional
 966 feature space. The comparison is either computed by integrating over the
 967 distributions of the different categories or, in the case of individually rep-
 968 resented exemplars, counting the numbers of exemplars. The stimulus item
 969 is then assigned to the category which contributes the highest number of
 970 exemplars within the neighborhood. If there are no exemplars within the
 971 neighborhood, the stimulus is not assigned to any category. Since we em-
 972 ploy a set-based n -dimensional exemplar representation, we adapt Lacerda’s
 973 functions as follows:

974 Given an exemplar memory $\mathcal{M}$ with categories $\mathcal{C} \subseteq \mathcal{M}$, $x \in \mathcal{M}$ denotes
 975 an exemplar x stored in memory, and $x \in \mathcal{C}_A$ denotes an exemplar x belonging
 976 to a specific category A . Formally, the memory $\mathcal{M}$ and the categories $\mathcal{C} \subseteq \mathcal{M}$
 977 are sets and exemplars are vectors as described in section 3.4.

978 The similarity of a stimulus x_0 to a category A according to Lacerda
 979 (1995) is given by equation 2.

$$s_L(x_0, \mathcal{C}) = \frac{|\{x \in \mathcal{C} | d(x_0, x) \leq \varepsilon\}|}{|\{x \in \mathcal{M} | d(x_0, x) \leq \varepsilon\}|} \quad (2)$$

980 The classification of a stimulus x_0 is then given as the category which
 981 maximizes the similarity (equation 3).

$$\text{category}_L(x_0) = \underset{\mathcal{C}}{\operatorname{argmax}} s_L(x_0, \mathcal{C}), \quad \text{for all } \mathcal{C} \subseteq \mathcal{M} \quad (3)$$

982 In the following simulations we set $\varepsilon = 1$ (which is equal to the standard
 983 deviation we set for the normally distributed exemplars in each category).
 984 There is no general rule for the value of ε . It depends on the data distribution
 985 and is often chosen based on experience.

986 A.2 The “Pierrehumbert classification”

987 The exemplar classification proposed by Pierrehumbert (2001) is an example
 988 of an ε -neighborhood classification which takes into account time. In Pierre-

humbert’s model, exemplars are weighted with an exponential decay according to their age. Since we do not consider time as a factor in our model of the McGurk effect, we set this temporal weight to 1 for all exemplars. This effectively reduced Pierrehumbert’s model to an unnormalized ε -neighborhood classification as given in equations 4 and 5.

$$s_P(x_0, \mathcal{C}) = |\{x \in \mathcal{C} | d(x_0, x) \leq \varepsilon\}| \quad (4)$$

$$\text{category}_P(x_0) = \underset{\mathcal{C}}{\operatorname{argmax}} s_P(x_0, \mathcal{C}), \quad \text{for all } \mathcal{C} \subseteq \mathcal{M} \quad (5)$$

A.3 The “Johnson classification”

The exemplar classification proposed by Johnson (1997) is an example of a general, unrestricted distance-based classification which follows Nosofsky (1988). In this model, the distances between stimuli and exemplars are weighted by an *attention weight* vector w which can assign different weights to the feature dimensions. The weighted euclidean distance between x_0 and an exemplar x is given by equation 6, where $w[i]$ denotes the *attention weight* on feature dimension i . The similarity between a stimulus x_0 and an exemplar $x \in \mathcal{M}$ is given by equation 7, where c is a *sensitivity constant* which can be set to reduce the impact of distant exemplars “to almost nothing” (Johnson, 1997)¹⁰.

Based on the instance-wise similarity, the cognitive similarity of an exemplar $x \in \mathcal{C}_A$ to a stimulus x_0 is given by their similarity weighted by a base activation $\bar{a}$ of that exemplar and some additional added exemplar-specific Gaussian noise N (equation 8). Finally, classification of a stimulus x_0 is given by the category for which the summed similarity of all its member exemplars is maximal as defined in equation 9.

$$d_w(x_p, x_q) = \sqrt{\sum_{i=1}^n w[i](x_p[i] - x_q[i])^2} \quad (6)$$

$$s_J(x_0, x) = e^{-cd_w(x_0, x)}, \quad \text{for all } x \in \mathcal{M} \quad (7)$$

¹⁰Johnson (1997) notes that setting the sensitivity constant accordingly, “the similarity function provides a sort of K nearest-neighbors classification, in which only nearby neighbors are considered.” While this is true in a first approximation, we present proper kNN classification in section A.5. The difference is, while the exponentially decreasing similarity with larger distances in Johnson’s model quickly approaches zero, the contribution of all exemplars beyond the k nearest ones is exactly zero in kNN.

$$a(x_0, x) = \bar{a}(x)s_J(x_0, x) + N(x) \quad (8)$$

$$\text{category}_J(x_0) = \underset{\mathcal{C}}{\operatorname{argmax}} \sum_{x \in \mathcal{C}} a(x_0, x), \quad \text{for all } \mathcal{C} \subseteq \mathcal{M} \quad (9)$$

1011 In the following simulations, we set the attention weight w to a vector of
 1012 ones, sensitivity $c = 1$, the base similarity $\bar{a} = 1$ and the classification noise
 1013 to $N = 0$ for all exemplars (thus, the similarity level is equal to the similarity
 1014 s_J).

1015 **A.4 The “Duran (global-similarity) classification”**

1016 The global-similarity classification has been proposed by Duran (2015). It is
 1017 inspired by the neighborhood classification according to Lacerda (1995) but
 1018 considers all exemplars in a modification to the distance-based classification
 1019 as proposed by Johnson (1997). Based on the euclidean distance (equation 1),
 1020 similarity between a stimulus x_0 and a category $\mathcal{C} \subseteq \mathcal{M}$ is given by equation
 1021 10.

$$s_D(x_0, \mathcal{C}) = \frac{1}{|\mathcal{C}|} \sum_{x \in \mathcal{C}} e^{-d(x_0, x)} \quad (10)$$

1022 The classification of a stimulus x_0 is then given as the category which
 1023 maximizes the overall similarity (equation 11).

$$\text{category}_D(x_0) = \underset{\mathcal{C}}{\operatorname{argmax}} s_D(x_0, \mathcal{C}), \quad \text{for all } \mathcal{C} \subseteq \mathcal{M} \quad (11)$$

1024 **A.5 The kNN classification**

1025 The k -nearest neighbors classification algorithm (kNN) assigns each stimulus
 1026 to the majority category of its k closest neighbors within the exemplar feature
 1027 space. The simplest case is $k = 1$ which is not robust as its classification
 1028 can be based on a single outlier. There is no general rule for the value of
 1029 k . It is usually desirable that $k > 1$ (often $k \gg 1$) and that k is odd such
 1030 that decision ties are less likely (Manning et al., 2008). In general, small k
 1031 introduce classification noise while large k introduce smooth (i.e. blurred,
 1032 less precise) decision boundaries.

1033 For the following simulations we use the kNN algorithm provided by the
 1034 R library **class** (version 7.3.15, Venables & Ripley, 2002).

1035 **A.6 The generalized weighted neighborhood classifica-** 1036 **tion**

1037 Finally, we implement a generalized weighted neighborhood classification.
1038 It combines ε -neighborhood classification with distance based classification.
1039 For all exemplars within an ε -neighborhood around a stimulus x_0 , similarity
1040 is weighted decreasing exponentially according to the euclidean distance, as
1041 shown in equation 12.

$$s_G(x_0, \mathcal{C}) = \sum_{x \in \mathcal{C} \wedge d(x_0, x) \leq \varepsilon} e^{-d(x_0, x)} \quad (12)$$

1042 Again, the classification of a stimulus x_0 is given as the category which
1043 maximizes the similarity (equation 13).

$$\text{category}_G(x_0) = \underset{\mathcal{C}}{\operatorname{argmax}} s_G(x_0, \mathcal{C}), \quad \text{for all } \mathcal{C} \subseteq \mathcal{M} \quad (13)$$

1044 **A.7 Simulations and results**

1045 We compare this set of classification methods in a number of numerical sim-
1046 ulations. All simulations are implemented in R. The source code is avail-
1047 able online at <https://github.com/simphon/...> **URL einfügen, Repository ein-**
1048 **richten!**

1049 For the sake of illustration, we use two dimensional data with three nor-
1050 mally distributed categories according to the specifications shown in table 1.
1051 Following convention, we denote the first feature dimension with x and the
1052 second dimension with y . The three categories are denoted as “ba”, “da”
1053 and “ga”¹¹.

1054 Within this two-dimensional space, we define a set of stimuli along a reg-
1055 ular grid from -5 to 10 (in both dimensions) and compute the classification
1056 for each stimulus according to the various classification methods presented
1057 above.

1058 A data set with the specified numbers of individual exemplars drawn from
1059 a normal distribution is generated according to the specifications given in
1060 table 1. The numbers of exemplars in each category are tested in three
1061 variations with each category in turn being 10 times larger than the remaining
1062 two. The neighborhood and the kNN classifications are affected by random
1063 exemplars sampled from a normal distribution, especially on the fringes of
1064 each category. In order to minimize this effect, we repeat each simulation

¹¹Note that this is an arbitrary denomination. We chose these terms based on the McGurk effect for the sake of this illustration.

category	μ_x	μ_y	σ_x	σ_y	num. exemplars
ba	0	0	1	1	{10000, 1000, 1000}
da	3	3	1	1	{1000, 10000, 1000}
ga	6	6	1	1	{1000, 1000, 10000}

Table 1: *Data specifications for the preliminary simulations.*

with the same parameters for 100 iterations and average the classification results over the set of stimuli.

As a result, we can draw a map with the predicted category at each stimulus location. Since most computational exemplar models are discussed on a one-dimensional feature space, we also present the similarity functions along one dimension. For this purpose, we take the similarity values along the diagonal in the grid of stimuli (i. e. the stimuli from $(-5, -5)$ to $(10, 10)$ with $x = y$). This diagonal intersects with all category centers in our data sets. The resulting plots allow a comparison with other models in the literature (and especially the originals of the methods implemented here).

We tested $k = \{3, 5, 7, 9, 11, 13\}$ and $k = 157$ for kNN classification (the latter value corresponding to the mean number of exemplars in the ε -neighborhood¹²). We show the results for the largest k value. The number of stimuli predicted to belong to the largest category decreased with smaller k (affecting mainly stimuli at the fringes). The similarity curves along the diagonal are essentially the same.

Predictions for the case of a large “ba” category are shown in Figure 7¹³. Results of the other two cases with large “da” and large “ga” are symmetric. Predictions for all classification methods show a larger area of stimuli predicted to belong to the larger category. The figures show how the neighborhood methods restrict their predictions to the area close to the actual data points. The other methods predict a category even for the most distant stimuli. Also, it can be observed that the peak similarity does not depend on the relative frequencies of the categories for the Lacerda, Duran and kNN classifications. With the other methods, the peaks are proportional to the relative number of exemplars within each category. This can also be observed clearly in figure 8, which shows the similarity curves along the diagonal through the category centers. The peak similarities are equal to 1 for all three categories

¹²The mean value of $k \approx 157$ was set empirically based on the given data set. Due to the random distribution of exemplars, the mean number of exemplars in the ε -neighborhood differed from run to run.

¹³Plots are produced with the **ggplot2** library in R (Wickham, 2016). Data ellipses are computed with the “stat_ellipse” function, which is based on **car::ellipse** (Fox & Weisberg, 2019)

1093 with Lacerda and kNN classification. The remaining methods have a similar-
1094 ity approximately proportional to the relative category frequencies (except
1095 for the Duran classification which has a summed total similarity of 1 over all
1096 three categories).

1097 **A.8 Conclusions on exemplar-based similarity and clas-** 1098 **sification**

1099 Based on general considerations for *cognitively plausible* exemplar-based clas-
1100 sification of speech items we have employed the generalized weighted neigh-
1101 borhood classification (section A.6) in our exemplar model of the McGurk
1102 effect. One particularly desirable property is that it does not predict any cat-
1103 egory for very distant stimuli. As Walsh et al. (2010) note, ε -neighborhood
1104 classification is superior to kNN and general distance-based classification
1105 since the latter always predict a category because even very distant (“un-
1106 grammatical”) stimuli always have nearest neighbors.¹⁴ Another desirable
1107 property of generalized weighted neighborhood classification is that its sim-
1108 ilarity depends on the number of exemplars in a category reflecting higher
1109 confidence in the classification of a stimulus as belonging to a larger cate-
1110 gory. It shows non-linear similarity curves (potentially corresponding to a
1111 perceptual magnet effect Lacerda, 1995; Duran, 2015) with decision bound-
1112 aries shifted away from the category with more exemplars. Though plausible
1113 in general, we do not take into account a priori similarity (Johnson, 1997)
1114 or a “resting activation level” (Pierrehumbert, 2001) which would allow for
1115 incorporation of priming effects or top-down expectations into the model of
1116 recognition.

1117 Considerations regarding computational cost within the implemented model
1118 are not taken into account, given the fact that neural computations in the
1119 human brain are massively parallelized, distributed and associative instead
1120 of sequential and numerical as in our digital computers.

¹⁴Note that recognition/categorization is not a prerequisite for the storage of an exemplar. In order for categories to emerge from exemplar distributions (e.g. during acquisition), all exemplar occurrences have to be stored in the first place. Thus, exemplars do not necessarily need to be associated with a category label.

Figure 7: *Category predictions.* Gray dots indicate the grid of stimuli. Colors indicate the predicted category for the stimulus at the corresponding location. Large circles show data ellipses (note that only the number of exemplars differed, while the standard deviations were equal for all categories). The dotted lines indicate the convex hull around each category (combined over all iterations). The red circle marks the approximate location of stimuli with inconsistent cues as in the case of the McClelland effect (cf. the illustration in figure 4).

Figure 8: *Similarity curves for stimuli along the x-y diagonal through the category centers. Colored lines correspond to the similarity strength for each category. Dotted vertical lines indicate the category centers.*