

Brownian dynamics for the vowel sounds of human language

J. Burridge*

School of Mathematics and Physics, University of Portsmouth

B. Vaux

Faculty of Modern and Medieval Languages and Linguistics, University of Cambridge

(Dated: November 25, 2019)

We present a model for the evolution of vowel sounds in human languages, in which words behave as Brownian particles diffusing in acoustic space, interacting via the vowel sounds they contain. Interaction forces, derived from a simple model of the language learning process, are attractive at short range and repulsive at long range. This generates sets of acoustic clusters, each representing a distinct sound, which form patterns with similar statistical properties to real vowel systems. Our formulation may be generalised to account for spontaneous self actuating shifts in system structure which are observed in real languages, and to combine in one model two previously distinct theories of vowel system structure: dispersion theory, which assumes that vowel systems maximize contrasts between sounds, and quantal theory, according to which non linear relationships between articulatory and acoustic parameters are the source of patterns in sound inventories. By formulating the dynamics of vowel sounds using inter-particle forces, we also provide a simple unified description of the linguistic notion of *push* and *pull* dynamics in vowel systems.

I. INTRODUCTION

Each human language has its own inventory of sounds. Within these inventories a distinction is made between consonants, which require some restriction of airflow for their production, and vowels; resonant sounds, for which airflow is relatively unrestricted [1]. Vowels are the primary carriers of linguistic information in connected human speech [2, 3]. The sound of a vowel is determined largely by the position and configuration of the tongue, although other parts of the vocal apparatus can be involved [4]. Linguists have traditionally represented vowels as points in a two dimensional *articulatory* domain (the *vowel quadrilateral*) with coordinates given by tongue *height* and *backness* [5]. To produce the vowel sound in ⟨cat⟩, represented phonetically as [æ], the tongue is in a low, forward position. By contrast the vowel in ⟨foot⟩, pronounced [u], is articulated with the tongue dorsum in a relatively high and backed position. Vowels may also be reliably identified from the first two peaks, or *formants* (F_1, F_2), of their frequency spectrum [6] (Figure 1). Models of vowel production [7] and experiments [8, 9] suggest that F_1 strongly correlates to tongue height, and F_2 to backness, although recent work suggests that F_2 may depend on both [9]. Mathematically, there appears to be a bijective map from articulatory to acoustic (F_1, F_2) space which is approximately affine [10], with some exceptional regions [11]. For this reason, we will view vowel sounds as existing in a closed two dimensional domain: *vowel space*.

Vowel systems exhibit recurring patterns and regularity: The majority of languages (64.6%) have between 5 and 7 different qualities of vowel with /i/ /a/, and /u/ occurring in over 80% of languages [12]. Moreover, certain

arrangements within vowel space are particularly common [12–14]. Vowel systems, like most elements of languages, evolve over time and may therefore be viewed as dynamical systems coupled to the social dynamics of human society, and also to geography and social networks [15–17]. Cross-linguistic similarities suggest that their internal dynamics may play a particularly powerful role, and numerous models have been proposed [13, 14, 18–25]. In the early work of Liljencrants and Lindblom [13], vowels were modelled as electrical charges, based on the principle of *maximal contrast* [26], yielding a single idealised vowel system for each cardinality (number of vowels). *Focalization* theory, which adds an attractive interaction to vowel dynamics, hypothesizes the convergence of formants [18, 27], increasing acoustic salience and “perceptual value”. Other models involve iterative construction driven by contrast maximization [20], or agent based imitation games between speech synthesizers [14, 21]. One inspiration for our work is *exemplar dynamics* [22–25, 28, 29], where agents store in memory a large number of exemplars for each sound, organised into sound categories. When uttering a sound, an exemplar is reproduced (with noise) from memory and accepted as valid by listeners based on how easily it can be identified and, in some cases, how typical it is of its category [24]. When two categories get close, they may overlap, making identification of utterances in the intersecting regions more difficult, leading to their rejection as exemplars. This shifts the mean sounds of the two categories apart, as if subject to a repulsive force. Exemplar models have recently been analysed [23, 30], with the aim of determining conditions for the merger of sound categories, and the behaviour of the boundaries between them in vowel space. Repulsive mechanisms are common to all the models above and generate distributions of sounds which are dispersed in vowel space. They are known as *dispersion theories* [12, 31].

* james.burridge@port.ac.uk

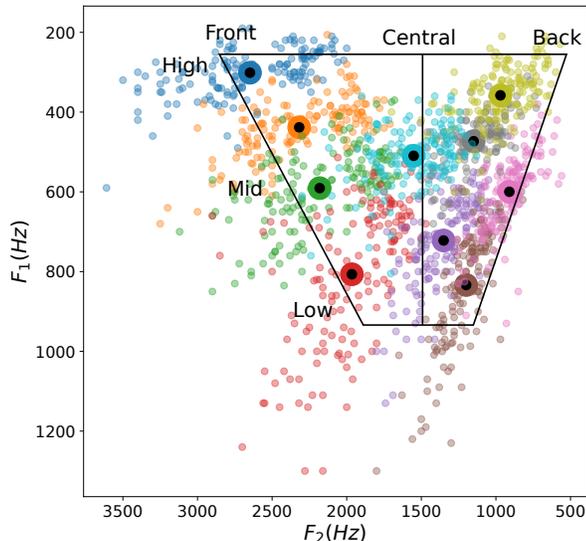


FIG. 1. Formant data for the vowel sounds of 76 speakers of American English from the original study of Peterson and Barney (1952) [6], where speakers were recorded reading the words *heed*, *hid*, *head*, *had*, *hod*, *hawed*, *hood*, *who'd*, *hud* and *heard*. Axes have been reversed so that the positions of vowels correspond to those in the IPA vowel chart [5]. The approximate shape of this chart has been superimposed on the formant data, to illustrate the correspondence between mean formant positions (large points), and the arrangement of the traditional chart. Front vowels: Blue *i*, Orange *ɪ*, Green *ɛ*, Red *æ*. Back vowels Light green *u*, Grey *ʊ*, Pink *ɔ*, Purple *ʌ*, Brown *ɑ*. Central vowel: Turquoise *ɜ*.

While there is clearly value in defining models with greater linguistic realism (using realistic synthesizer equations [21], or metrics of *perceptual distance* [19]), this approach also makes the determination of their general behaviour difficult. It is hard to evaluate whether the complexity they add in order to better match reality is scientifically justified, or whether it is a form of interpolation. Exemplar theory is a step away from this approach, and may be defined as a simple iterative computational model. However, rigorous analysis is challenging [30], and because the atomic constituents of the theory are tokens in the memories of speakers, simulations of large numbers of different sounds and words are also computationally expensive. The theory has been used to explore interactions between words and phonemes [23, 24], but it has not so far been used to model the evolution of realistic vowel systems.

The model that we present builds on the ideas described in the above models, but we aim for an analytical definition which is simpler to simulate and analyse. Sound change in language may be thought of as the diffusion of word pronunciations in acoustic space, making Brownian (Langevin) dynamics [32, 33] the natural

mathematical description of their motion. We think of vowel systems as a “soup” of words or, more technically, *phonological frames*. These frames interact via the vowel sounds that they contain, and the interaction forces may be seen either as a phenomenological model [34], based on established qualitative models of sound change [35], or as a simplified version of exemplar dynamics. Interactions between frames are mediated by a cloud of utterances – in our case formant clouds – but unlike exemplar theory, we do not explicitly model this collection of sounds, only the mean sound for each frame. This dramatically simplifies the model, yielding one stochastic differential equation per frame, facilitating analytical calculations, and allowing the simulation of a large number of words. Our formulation is analogous to a physical model (frames form a charged colloid [36–39]) allowing us to clarify traditional qualitative descriptions of change in terms of *pushing* and *pulling*. Its simplicity also allows a number of extensions to describe a range of different phenomena in one unified model. These include self actuating sound change [40–42], allophonic sound variations, the effect of non-linearities in the relationship between articulatory and acoustic parameters [11, 43] and word-frequency effects which have recently been observed in a purely computational exemplar model [24].

II. A BRIEF INTRODUCTION TO PHONOLOGY

For the benefit of non-linguists, we now review the relevant elements of *phonology*: the branch of linguistics that deals with systems of sounds.

A. Phonemes and allophones

Individual speech sounds are usefully viewed as existing within phonological frames [24]. For instance, in English the frame /m_p/, if it forms a single word, will accept one of two sound categories, creating either ⟨map⟩ or ⟨mop⟩. Similar variants appear in other words, with subtle variations depending on the frame. Because of these variations, a distinction is made between the contrastive sound categories: *phonemes* and their context dependent versions, termed *allophones*. For example ⟨pea⟩, ⟨spin⟩ and ⟨sip⟩ all contain what English speakers might call a p-sound, but these sounds are all slightly different. The transcriptions of these words into phonetic symbols, which represent specific sounds, are [p^hi:], [spm] and [sɪp̚]. The three p-sounds here are, respectively *aspirated* (followed by a burst of breath), *unaspirated* and *unreleased* (meaning that there is no audible end to the temporary occlusion of airflow needed to make the sound). These three variants are allophones of the English phoneme /p/

$$/p/ = \{[p^h], [p], [p̚]\}. \quad (1)$$

Phonemes differ between languages: Thai speakers, for example, consider [p^h] and [p] to be manifestations of distinct phonemes, /p^h/ and /p/ respectively. Allophones may be grouped into phonemes by examining the sounds that surround them (their phonological environment). Consider the frame /kæ./ where /æ/ is the vowel phoneme in the word ⟨cat⟩. Inserting any of the allophones of /p/ would produce an utterance recognised by English speakers as the word ⟨cap⟩. However inserting the sound [b] produces an utterance with a different meaning. We say that [b] is in *contrastive distribution* with [p]; if we switch the sounds in the frame, we change the meaning. Allophones of the same phoneme are generally not in contrastive distribution; they are (normally) in *complementary distribution*. This means that they are not found in the same immediate phonological environment. For example [p^ɹ] can only occur in syllable-final consonant clusters and [p^h] only occurs either word-initially or at the beginning of a stressed syllable. Therefore they never contrast, and given a phoneme we can predict which allophone will appear simply by knowing what other sounds surround it. The definition of a phoneme and its allophones applies identically to consonants and vowels. For example, an English vowel phoneme is /æ/, which occurs in ⟨bad⟩ and ⟨bat⟩. Here the allophone in ⟨bad⟩ is longer, transcribed [æː]. Another characteristic which distinguishes vowel allophones is *nasalization* as in ⟨ban⟩[bæ̃n].

Since the allophones of a phoneme typically exhibit quite subtle variations, the term *phoneme* is often used as if it referred to a single sound in the language. Doing so is no less consistent than using *allophone* in the same way. In reality the units of sound uttered by speakers vary widely between and within individuals. Both phonemes and allophones are categories, with one being a subcategory of the other. We will exploit this idea when we come to explore allophonic variations in section VII.

B. Vowel systems

As explained above, the first two formants, or equivalently the position of the tongue, are the primary determiners of vowel sounds, but other articulatory variations can be involved as well. After height and backness, the next most common is lip rounding [12]. There is, however, a strong cross-linguistic correlation between rounding and backness (94.0% of front vowels are unrounded, and 93.5% of back vowels are rounded), making the parameter redundant in most cases. Beyond lip rounding, many languages have separate series of vowels, each distinguished by some additional characteristic such as length (long vs short) or nasalisation (nasalised vs oral). Often the sounds in both series occupy the same positions in vowel space. For example in Mazatec [44] there are four oral vowels /i,ɛ,a,o/ and four corresponding nasalised versions /ĩ,ɛ̃,ã,õ/. In this case we say that Mazatec has four vowel *qualities*. This matching be-

tween series is the norm so we can ignore additional characteristics and still provide description of vowel systems which captures the essential properties of their structure for most languages. This will be our approach, until we consider allophonic variation in section VII.

III. THE MODEL

We view words as particles in acoustic space, with positions determined by the vowel sounds they contain. A cluster of words then implies the existence of a frequently occurring sound in a language. We propose short-range attractive and long-range repulsive forces between words based on some ideas from linguistics about the interactions between phonemes [13, 17, 45]. Before setting out the details we describe how such forces can lead to the formation of a vowel system. Consider a large number words distributed uniformly at random throughout a chamber representing vowel space as in Figure 2. At first, particles near to each other will be drawn together, forming loose clusters whose typical size will be determined by the radius at which the attractive force becomes repulsive, as in Figure 3. As time progresses, provided the volatility of the random component of their motion is not too large, these clusters will become tighter and more separated due to long range repulsive effects. Depending upon the shape of the chamber, and the range of the short and long range components of the interaction force, we will obtain a number of different arrangements of particle clusters within the system, each of which represents a different inventory of sounds.

We now set out the details. For words with only one vowel, their position is unambiguous, but words with two or more vowels occupy multiple positions. We view such words as generating one or more phonological frames, or environments, for the vowels they contain. For example

$$\langle \text{hoodwink} \rangle \rightarrow /h_dwɪŋk/ + /hʊdw_ɪŋk/. \quad (2)$$

Each such frame now has a unique position in the acoustic domain, determined by the sound which is placed in its gap. For simplicity we assume that two frames from the same word interact in the same way as frames from different words.

We now consider a large group of speakers who are sufficiently socially connected to pronounce the words of their language in a roughly similar way. Acoustic experiments show that vowel sounds are subject to variation. Peterson and Barney (1952) [6] collected formant data for a group of 76 American speakers pronouncing a series of ten words which differed only in their vowel, allowing F_1/F_2 values to be collected for ten vowels. Because these utterances vary between speakers, the formant values form clouds in acoustic space. In Figure 1 the large coloured dots show the centroids of each vowel phoneme cloud. We have superimposed the standard IPA vowel quadrilateral [5] on the scatter plot, and we note that to a good approximation the centroids sit in the chart

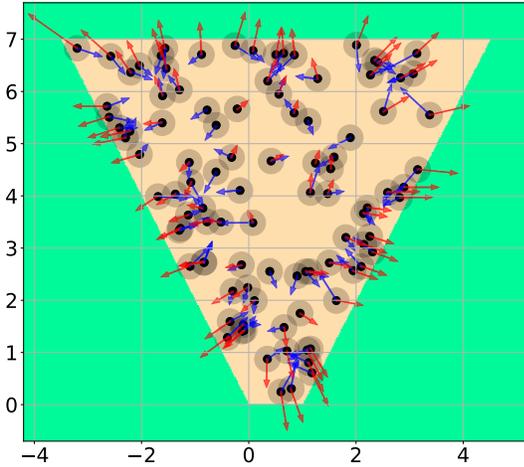


FIG. 2. Early stage configuration of 100 particles (words) in a chamber representing acoustic (F_1, F_2) vowel space. Range of attractive force $\alpha = 1$, cloud radius $\sigma = 1$, diffusion coefficient $D = 0.1$. Blue arrows show attractive forces, red show repulsive.

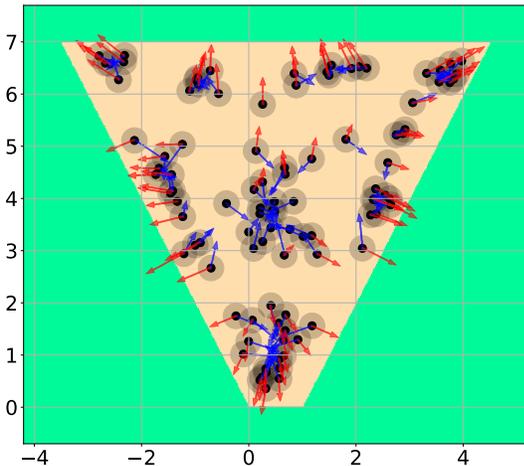


FIG. 3. Later stage configuration ($t=0.5$) of 100 particles (words) in a chamber representing acoustic (F_1, F_2) vowel space. Range of attractive force $\alpha = 1$, cloud radius $\sigma = 1$, diffusion coefficient $D = 0.1$.

positions assigned by linguists over half a century earlier [46].

We let $\mathbf{x}_i(t) \in \mathbb{R}^2$ be the population average vowel sound uttered for frame i , and we refer to $\mathbf{x}_i(t)$ as the position of this frame. Because utterances of the vowel sound in each word form a cloud in acoustic space [6], the next utterance to be heard from frame i will be a

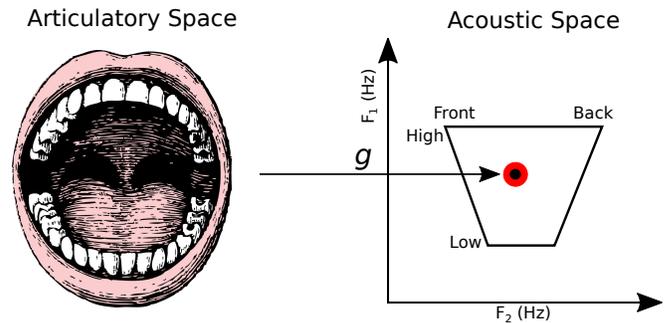


FIG. 4. The articulatory to acoustic map, \mathbf{g} . In order to utter a target sound in acoustic space, a speaker must arrange her articulatory apparatus (tongue, lips, laryngeal structures etc.) into the correct position, \mathbf{y} . The uttered sound will be $\mathbf{x} = \mathbf{g}(\mathbf{y})$.

random variable $\mathbf{X}_i(t)$. The speaker who produces this utterance must arrange her vocal apparatus into an appropriate configuration to create the desired sound. We denote by \mathbf{g} the function which maps points in articulatory space (mouth cavity shape etc.) to points in acoustic space as illustrated in Figure 4. To generate the desired sound, the speaker must have an intuitive knowledge of the inverse mapping \mathbf{g}^{-1} so that given the target sound \mathbf{x} she is able to form her mouth parts into the appropriate configuration $\mathbf{y} = \mathbf{g}^{-1}(\mathbf{x})$. We assume, for now, that on average the articulatory states \mathbf{Y}_i used by the population to generate sounds \mathbf{X}_i produce unbiased results in the sense that the average values of outputs over all speakers and utterances are equal to the frame position

$$\mathbb{E}[\mathbf{g}(\mathbf{Y}_i(t))] = \mathbb{E}[\mathbf{X}_i(t)] = \mathbf{x}_i(t). \quad (3)$$

We also assume that articulatory states used to generate \mathbf{x}_i are normally distributed, having mean $\mathbf{g}^{-1}(\mathbf{x}_i)$. In this case, if the map \mathbf{g} is affine

$$\mathbf{x} = \mathbf{g}(\mathbf{y}) = A\mathbf{y} + \mathbf{c} \quad (4)$$

where $A \in \mathbb{R}^{2 \times 2}$, then the utterances \mathbf{X}_i will also be normally distributed, having density function

$$\psi_i(\mathbf{x}) = \frac{1}{2\pi|\Sigma|} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}_i)^T \Sigma^{-1}(\mathbf{x} - \mathbf{x}_i)\right). \quad (5)$$

Here, the covariance matrix Σ determines the shape of the cloud of utterances [6, 47]. In the spherical case we set $\Sigma = \sigma\mathbf{I}$ where \mathbf{I} is the identity matrix, and σ is the *cloud radius*. Strictly speaking we could have taken (5) as the definition of the shape of the utterance cloud in acoustic space, without needing to consider the map \mathbf{g} . However, experiments show that there are regions of acoustic space where \mathbf{g} is not affine [11] and, as we will see in section VI, this affects both cloud shape and system dynamics.

We define the state of the vowel system of our language to be the set of all frame positions $\{\mathbf{x}_i(t)\}_{i=1}^n$, where n is the number of frames. Interactions between these frames are generated by the language learning process. When

a speaker learns how to pronounce the word which generated frame i , not only will utterances of the word itself provide templates for its vowel sound, but so will similar sounds in other words. If two frames contain vowel sounds which can be used as templates for one another, that is, speakers consider them to contain *the same sound*, we write

$$\mathbf{x}_i \stackrel{T}{=} \mathbf{x}_j, \quad (6)$$

where $\stackrel{T}{=}$ denotes *template equality*. We write S_i for the set of frames whose vowel sounds act as templates for frame i

$$S_i = \{j \text{ such that } \mathbf{x}_j \stackrel{T}{=} \mathbf{x}_i\} \quad (7)$$

where time dependence is implicit. In this paper we use proximity in acoustic space to define template equality

$$\mathbf{x}_i \stackrel{T}{=} \mathbf{x}_j \equiv |\mathbf{x}_i - \mathbf{x}_j| < \alpha \quad (8)$$

where $\alpha \geq 0$ is the *template range*. The set of templates for each frame will evolve over time, as frames change position. The overall density of utterances which may be used as templates for learning the vowel sound in frame i is then defined

$$\hat{\psi}_i(\mathbf{x}) := \sum_{j \in S_i} (1 + \omega \delta_{ij}) f_j \psi_j(\mathbf{x}) \quad (9)$$

where f_j is the relative frequency with which frame j is uttered, and $\omega \geq 0$ (the *self-focus*) is the extra weight placed on utterances of the frame as a template for its own vowel sound. We call $\hat{\psi}_i(\mathbf{x})$ the *template density* for frame i . When we have a set of frames for which every frame is a template for every other, then we call this set a phoneme. Given the template density for a frame i we can compute the *template mean* for that frame

$$\hat{\mathbf{x}}_i = \frac{\int_{\mathbb{R}^2} \hat{\psi}_i(\mathbf{x}) \mathbf{x} d\mathbf{x}}{\int_{\mathbb{R}^2} \hat{\psi}_i(\mathbf{x}) d\mathbf{x}} \quad (10)$$

$$= \frac{1}{\mathcal{N}_i} \int_{\mathbb{R}^2} \hat{\psi}_i(\mathbf{x}) \mathbf{x} d\mathbf{x} \quad (11)$$

$$= \frac{\sum_{j \in S_i} (1 + \omega \delta_{ij}) f_j \mathbf{x}_j}{\sum_{j \in S_i} (1 + \omega \delta_{ij}) f_j} \quad (12)$$

where \mathcal{N}_i is the normalising constant for the template density. The template mean, $\hat{\mathbf{x}}_i$, is the mean value, weighted for self-focus, of all the utterances from the frames which a language learner uses when learning the sound in frame i . The linguistic environment of each new learner will be different, and they will inevitably introduce their own idiosyncrasies driven by learning mistakes, the desire to emulate certain individuals, and variations in their own physiology. However, on average we have no reason to expect anything other than unbiased variations around the template mean. Ignoring the effect of frames

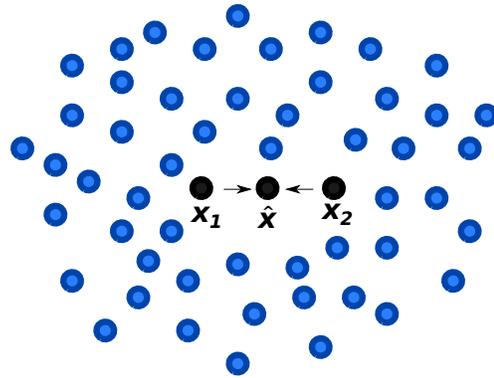


FIG. 5. The set of utterances for two different frames whose vowels are part of the same phoneme. The template mean is the average location of the locations of the frames, and since frames are attracted toward their template mean, they are attracted to other frames in the same phoneme.

which are not templates, we therefore expect new speakers coming of age to use sounds which on average match the template mean. As older speakers with more archaic forms of speech die, the speech sounds of the population as a whole will move in the direction of the template mean. This behaviour will induce an effective force on frame i which draws it toward the mean of its templates

$$\mathbf{f}_i^{\text{att}} := \hat{\mathbf{x}}_i - \mathbf{x}_i. \quad (13)$$

This is the simplest choice of inter-particle force consistent with the above considerations, which are summarised visually in Figure 5.

We now consider repulsive interactions, which are induced by disruptions to the learning process caused by frames which are nearby in acoustic space, but not sufficiently near to be templates. Given a frame i , the density of such *anti-template* frames is

$$\tilde{\psi}_i(\mathbf{x}) := \sum_{j \notin S_i} f_j \psi_j(\mathbf{x}). \quad (14)$$

We call $\tilde{\psi}_i(\mathbf{x})$ the *anti-template density*. Sounds from anti-templates will interfere with language learners' ability to recognise nearby template sounds, making the template sounds less likely to be copied [17, 23, 45]. To see how this might occur, consider a language that contains two acoustically similar vowel phonemes. Suppose that a language learner has noticed that these two sounds play two different roles in the language. We do not know how the developing mind achieves this but it is necessary in order to make sense of language. For example, there are *minimal pairs* of words such as ⟨pen⟩[p^hɛn] and ⟨pan⟩[p^hæn] which contrast only in a single vowel. Here it is essential to be able to distinguish /ɛ/ from /æ/. When learning how these two phonemes should sound, there will be many occasions where words are uttered using a sound which, from a purely acoustic perspective, is hard to categorise as one or the other. Evidence for this is shown

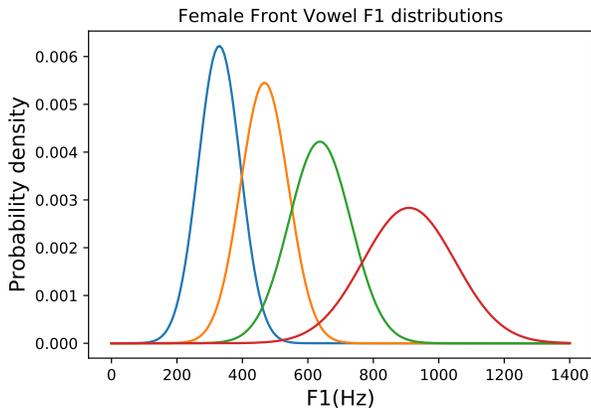


FIG. 6. Approximate distribution of F_1 values for female front vowels, derived from Peterson and Barney vowel data [6]. Curves are normal densities having means and standard deviations equal to those of the first formants of female front vowels. Colour coding matches that in Figure 1: Blue i , Orange ɪ , Green ε , Red æ .

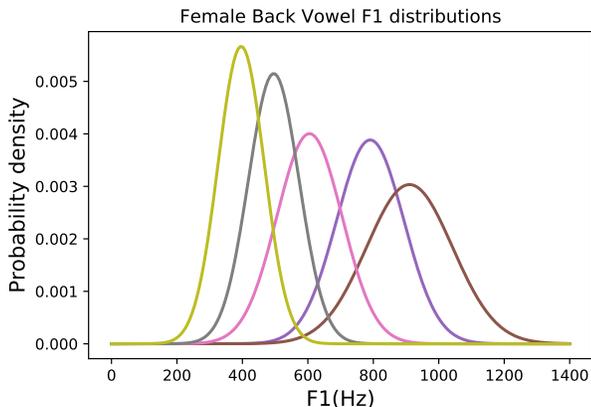


FIG. 7. Approximate distribution of F_1 values for female back vowels, derived from Peterson and Barney vowel data [6]. Curves are normal densities having means and standard deviations equal to those of the first formants of female back vowels. Colour coding matches Figure 1: Light green u , Grey $ɔ$, Pink $ɔ$, Purple ʌ , Brown ɑ .

in Figures 6 and 7 where, at least in American English, phonemes can be separated by as little as one standard deviation of their acoustic distribution. In this situation, an experienced speaker may be able to use a word’s context, and the vowel’s phonological environment, to efficiently identify what has been said. However a younger speaker must build up this ability over time. Until then, there will be cases where the meaning of a word is ambiguous. Even if the word meaning is understood, the utterance may not be perceived as a legitimate pronunciation, if its phonemes are too close to other recognisable sounds. In both these cases the uttered sound is, we assume, less likely to influence the language devel-

opment of the speaker. An argument similar to this is made by Labov [35] and is the basis of exemplar theory [23, 24, 45]. The effect of rejecting or otherwise reducing the importance of acoustically ambiguous phonemes during the learning process will be to push the sounds of the language away from regions of acoustic space which are overcrowded with phonemes. This will induce an effective repulsive force between frames that are not close enough to be considered as templates of one another.

To arrive at a plausible form for the repulsive effect on frame i from its anti-template frames, consider a single frame $j \notin S_i$. A simple measure of the extent to which j crowds i is the ratio of the density of frame j utterances at \mathbf{x}_i , to the total template and anti-template frame density in the same location, weighted for *functional load*, $\gamma \in [0, 1]$: the amount of work that individual vowel phonemes do in distinguishing words [48, 49]

$$\frac{\gamma f_j \psi_j(\mathbf{x}_i)}{\widehat{\psi}_i(\mathbf{x}_i) + \gamma \widetilde{\psi}_i(\mathbf{x}_i)}. \quad (15)$$

As $\gamma \rightarrow 0$ sounds become irrelevant to word identification and so the effect of overcrowding becomes negligible. The effect of overcrowding by frame j will be proportional both to this overcrowding ratio, and also to its distance away from i : if more distant templates are rejected, the effect on the mean of the templates which are accepted will be larger. The total repulsive force on frame i consistent with these assumptions is then

$$\mathbf{f}_i^{\text{rep}} := \frac{\gamma}{\widehat{\psi}_i(\mathbf{x}_i) + \gamma \widetilde{\psi}_i(\mathbf{x}_i)} \sum_{j \notin S_i} f_j \psi_j(\mathbf{x}_i)(\mathbf{x}_i - \mathbf{x}_j), \quad (16)$$

and the total force on frame i is $\mathbf{f}_i = \mathbf{f}_i^{\text{att}} + \mathbf{f}_i^{\text{rep}}$. To summarise: in the absence of phonemic overcrowding, the mean acoustic position of all the utterances of a frame made by new learners is equal to the mean position of all the templates for that frame. These include the frame itself, and other frames containing sufficiently similar sounds. Phonemic overcrowding, realised as the close proximity of anti-templates, causes the rejection of templates in overcrowded regions, creating a bias away from these regions. Attraction toward the template mean generates a short range attractive force between frames, and bias away from overcrowded regions generates a long range repulsive force. The template range α determines the radius of the attractive region around a frame. Outside this radius, the strength and range of repulsive inter-frame forces are controlled, respectively, by the functional load γ and by the cloud radius σ .

The interactions defined by (13) and (16) are not symmetric or additive. More common frames exert a greater influence on their surroundings, and the effect of one frame on another depends on its relative rather than absolute density. We capture stochasticity in speaker behaviour with a diffusion process with coefficient D , so each frame obeys

$$d\mathbf{x}_i(t) = \mathbf{f}_i dt + \sqrt{2D} d\mathbf{W}_i \quad (17)$$

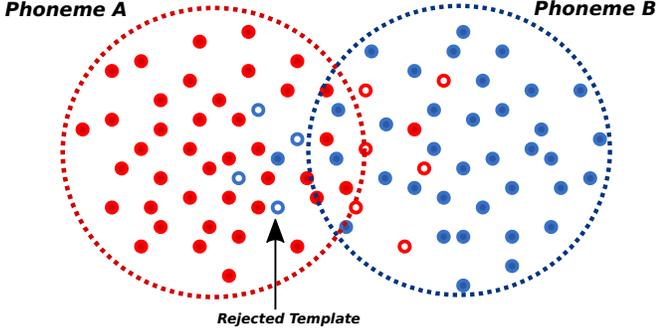


FIG. 8. The set of utterances for two different phonemes, each consisting of at least one frame. When two phonemes are near to each other in acoustic space, utterances of one may sound like the other increasing the likelihood that these utterances will not be accepted as valid templates by language learners. Such utterances are represented as open circles in the above diagram.

where \mathbf{W}_i is an n -dimensional Brownian motion [32]. This Itô stochastic differential equation (SDE) is equivalent to inertia free Langevin dynamics [33].

IV. DERIVATION AS AN EXEMPLAR MODEL

Our model (17) is phenomenological in the Landau sense [34]. We now derive it as an approximate exemplar model, partly in order to connect it with recent theory (exemplar dynamics [22–24]) but also because it is useful (for Sections VI and VIII and for future work) to have an explicit model of language learners accepting and rejecting templates. Exemplar dynamics in its purest form is an explicit computational model of a large population of sound units (exemplars) in the memories of speakers. Typically these sounds are characterised by a single acoustic variable [23, 24]. Exemplars for us are a subset of the utterances $\{\mathbf{X}_i\}_{i=1}^n$ (those which are not rejected by learners). Excepting the means $\{\mathbf{x}_i\}_{i=1}^n$, the distributions of the utterances \mathbf{X}_i are specified exogenous to the model. This simplification allows the dynamics of our language to be specified as a set of SDEs.

If an utterance $\mathbf{X}_j(t) \in S_i$ influences how a young speaker learns the vowel sound, $\mathbf{x}_i(t)$, in frame i then it is an *exemplar* of that sound. Suppose that a template of frame i is uttered as sound \mathbf{x} and $p_i(\mathbf{x})$ is the probability that this sound is accepted as an exemplar of the frame. The *exemplar mean* for frame i is then the expected value of its exemplars

$$\bar{\mathbf{x}}_i := \frac{\int \hat{\psi}_i(\mathbf{x}) p_i(\mathbf{x}) \mathbf{x} d\mathbf{x}}{\int \hat{\psi}_i(\mathbf{x}) p_i(\mathbf{x}) d\mathbf{x}}. \quad (18)$$

We may approximate the exemplar mean by expanding $p_i(\mathbf{x})$ to first order about the template mean. The template density (a superposition of Gaussian densities having overall mean $\hat{\mathbf{x}}_i$) is approximated as a Gaussian with

mean $\hat{\mathbf{x}}_i$ (so $\nabla \hat{\psi}_i(\hat{\mathbf{x}}_i) = 0$) and covariance $c_i \Sigma$. We have

$$\bar{\mathbf{x}}_i \approx \frac{\int \hat{\psi}_i(\mathbf{x}) \mathbf{x} (p_i(\hat{\mathbf{x}}_i) + (\mathbf{x} - \hat{\mathbf{x}}_i) \cdot \nabla p_i(\hat{\mathbf{x}}_i)) d\mathbf{x}}{\int \hat{\psi}_i(\mathbf{x}) (p_i(\hat{\mathbf{x}}_i) + (\mathbf{x} - \hat{\mathbf{x}}_i) \cdot \nabla p_i(\hat{\mathbf{x}}_i)) d\mathbf{x}} \quad (19)$$

$$= \frac{p_i(\hat{\mathbf{x}}_i) \mathcal{N}_i \hat{\mathbf{x}}_i + \left(\int \hat{\psi}_i(\mathbf{x}) \mathbf{x} \otimes (\mathbf{x} - \hat{\mathbf{x}}_i) d\mathbf{x} \right) \nabla p_i(\hat{\mathbf{x}}_i)}{p_i(\hat{\mathbf{x}}_i) \mathcal{N}_i + \left(\int \hat{\psi}_i(\mathbf{x}) (\mathbf{x} - \hat{\mathbf{x}}_i) d\mathbf{x} \right) \cdot \nabla p_i(\hat{\mathbf{x}}_i)} \quad (20)$$

$$= \hat{\mathbf{x}}_i + \frac{\left(\int \hat{\psi}_i(\mathbf{x}) \mathbf{x} \otimes (\mathbf{x} - \hat{\mathbf{x}}_i) d\mathbf{x} \right) \nabla p_i(\hat{\mathbf{x}}_i)}{\mathcal{N}_i p_i(\hat{\mathbf{x}}_i)} \quad (21)$$

$$= \hat{\mathbf{x}}_i + \frac{c_i \Sigma \nabla p_i(\hat{\mathbf{x}}_i)}{p_i(\hat{\mathbf{x}}_i)} \quad (22)$$

where \otimes is the outer product. This relation links the exemplar mean to the template mean via the acceptance probability. The number $c_i > 1$ is the ratio of the width of the template density to the width of the densities of individual frames. Assuming that frame clusters are tight compared to cloud size, then $c_i \approx 1$.

We now explicitly define the acceptance probability based on an overcrowding argument similar to that used in section III. Consider frame i , located at \mathbf{x}_i , and also another point \mathbf{x} in vowel space. If \mathbf{x} is not too far from \mathbf{x}_i , and a large fraction of the sounds at \mathbf{x} come from templates of i , then these templates are unlikely to be confused with anti-templates of i near \mathbf{x} . However, if location \mathbf{x} is overcrowded with nearby anti-templates then rejection becomes likely. The simplest choice consistent with these considerations is to let $p_i(\mathbf{x})$ be the relative density of i -templates at \mathbf{x} , corrected for functional load

$$p_i(\mathbf{x}) = \frac{\hat{\psi}_i(\mathbf{x}) \mathbf{1}_{|\mathbf{x} - \mathbf{x}_i| < R}}{\hat{\psi}_i(\mathbf{x}) + \gamma \tilde{\psi}_i(\mathbf{x})}, \quad (23)$$

where $R > 0$ is a cut-off radius beyond which sounds are rejected outright. In the absence of functional load ($\gamma = 0$), $p_i(\mathbf{x}) = \mathbf{1}_{|\mathbf{x} - \mathbf{x}_i| < R}$, so vowel sounds are rejected only when excessively distant from the current frame position, without reference to the crowding effects of other phonemes. Using (23) we have

$$\frac{\Sigma \nabla p_i(\hat{\mathbf{x}}_i)}{p_i(\hat{\mathbf{x}}_i)} = \frac{-\gamma \hat{\psi}_i(\hat{\mathbf{x}}_i) \Sigma \nabla \tilde{\psi}_i(\hat{\mathbf{x}}_i) \left(\hat{\psi}_i(\hat{\mathbf{x}}_i) + \gamma \tilde{\psi}_i(\hat{\mathbf{x}}_i) \right)}{\left(\hat{\psi}_i(\hat{\mathbf{x}}_i) + \gamma \tilde{\psi}_i(\hat{\mathbf{x}}_i) \right)^2 \hat{\psi}_i(\hat{\mathbf{x}}_i)} \quad (24)$$

$$= -\frac{\Sigma \nabla \tilde{\psi}_i(\hat{\mathbf{x}}_i)}{\gamma^{-1} \hat{\psi}_i(\hat{\mathbf{x}}_i) + \tilde{\psi}_i(\hat{\mathbf{x}}_i)} \quad (25)$$

$$= -\sum_{j \notin S_i} \frac{f_j \Sigma \nabla \psi_j(\hat{\mathbf{x}}_i)}{\gamma^{-1} \hat{\psi}_i(\hat{\mathbf{x}}_i) + \tilde{\psi}_i(\hat{\mathbf{x}}_i)} \quad (26)$$

$$= -\sum_{j \notin S_i} \frac{f_j \Sigma \Sigma^{-1}(\mathbf{x}_j - \hat{\mathbf{x}}_i) \psi_j(\hat{\mathbf{x}}_i)}{\gamma^{-1} \hat{\psi}_i(\hat{\mathbf{x}}_i) + \tilde{\psi}_i(\hat{\mathbf{x}}_i)} \quad (27)$$

$$= \sum_{j \notin S_i} \frac{f_j \psi_j(\hat{\mathbf{x}}_i) (\hat{\mathbf{x}}_i - \mathbf{x}_j)}{\gamma^{-1} \hat{\psi}_i(\hat{\mathbf{x}}_i) + \tilde{\psi}_i(\hat{\mathbf{x}}_i)} \quad (28)$$

yielding the following expression for the exemplar mean

$$\bar{\mathbf{x}}_i \approx \hat{\mathbf{x}}_i + \sum_{j \notin S_i} \frac{f_j \psi_j(\hat{\mathbf{x}}_i)(\hat{\mathbf{x}}_i - \mathbf{x}_j)}{\gamma^{-1} \tilde{\psi}_i(\hat{\mathbf{x}}_i) + \tilde{\psi}_i(\hat{\mathbf{x}}_i)}. \quad (29)$$

If acoustic differences between the mean vowel sounds of frames within one phoneme greatly exceed the difference between phonemes, or the cloud radius, then we can approximate the template means in the summand of (29) with \mathbf{x}_i . In this case we have

$$\bar{\mathbf{x}}_i - \mathbf{x}_i \approx \hat{\mathbf{x}}_i - \mathbf{x}_i + \sum_{j \notin S_i} \frac{f_j \psi_j(\mathbf{x}_i)(\mathbf{x}_i - \mathbf{x}_j)}{\gamma^{-1} \tilde{\psi}_i(\mathbf{x}_i) + \tilde{\psi}_i(\mathbf{x}_i)} \quad (30)$$

$$= \mathbf{f}_i^{\text{att}} + \mathbf{f}_i^{\text{rep}} \quad (31)$$

where $\mathbf{f}_i^{\text{att}}$ and $\mathbf{f}_i^{\text{rep}}$ are the phenomenological forces defined in section III. If we consider a single frame, and assume that on average young speakers match the mean of the sounds they accept as exemplars for this frame, then the community's speech will evolve towards that mean. The simplest force consistent with this assumption is $\mathbf{f}_i = \bar{\mathbf{x}}_i - \mathbf{x}_i$, and we will show below that this force may be derived from a simple agent based model. The above calculations showed that in the exemplar acceptance/rejection picture, this force decomposes into short range attractive and long range repulsive components given by (13) and (16).

To derive the SDE (17) we consider a community of N speakers whose evolution is driven by the replacement of older speakers by new speakers who learn from the community. We divide time into intervals of length $\delta t = N^{-1}$. At each interval a speaker is selected uniformly at random from the population, "retired", and replaced with a new speaker whose language state is a random variable with expectation equal to the exemplar mean. We can write this new state $\bar{\mathbf{x}}_i + \boldsymbol{\epsilon}_i$ where $\boldsymbol{\epsilon}_i = s\mathbf{Z}_i$, with $\mathbf{Z}_i \sim \mathcal{N}(0, \mathbf{I})$, is a random variable which captures the stochasticity injected by the new speaker, perhaps by their own free will, physiological differences, or selective copying of certain individuals. This produces speakers who live for a geometrically distributed number of time intervals, with mean N , giving an average lifespan of one time unit. To calculate the state of the community after a replacement, we first write the state of the speaker who is removed as

$$\mathbf{x}_i^{\text{dead}} = \mathbf{x}_i + \boldsymbol{\delta}_i \quad (32)$$

where $\boldsymbol{\delta}_i$ is a zero mean random variable with approximately the same statistical properties as $\boldsymbol{\epsilon}_i$. After a replacement the new frame position will be

$$\mathbf{x}_i(t + \delta t) = \mathbf{x}_i(t) - \frac{\mathbf{x}_i^{\text{dead}}(t)}{N} + \frac{\bar{\mathbf{x}}_i(t) + \boldsymbol{\epsilon}_i}{N} \quad (33)$$

$$= \left(\frac{N-1}{N} \right) \mathbf{x}_i(t) + \frac{1}{N} (\bar{\mathbf{x}}_i(t) + \boldsymbol{\epsilon}_i - \boldsymbol{\delta}_i). \quad (34)$$

Defining $\delta \mathbf{x}_i(t) = \mathbf{x}_i(t + \delta t) - \mathbf{x}_i(t)$, we have

$$\delta \mathbf{x}_i = (\bar{\mathbf{x}}_i - \mathbf{x}_i) \delta t + \sqrt{\frac{2s^2}{N}} \mathbf{Z}_i \sqrt{\delta t}. \quad (35)$$

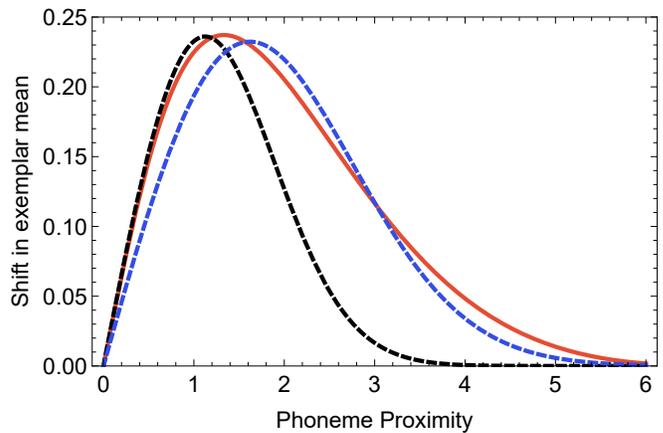


FIG. 9. Shift in the exemplar mean away from the frame position $\bar{\mathbf{x}}_i - \mathbf{x}_i$ for one phoneme (modelled by unit variance Gaussian) induced by another identically sized phoneme. Red: exact calculation using acceptance probability (23) when $\sigma = 1, \gamma = 0.5, R = 3$. Black: approximation with the same parameter values (30). Blue: approximation with larger clouds $\sigma = 1.5$ and lower functional load $\gamma = 0.3$.

In this simple model, stochasticity in the population as a whole is smaller in larger populations. However, an implicit assumption of the model is that new speakers are exposed to the whole community, because their linguistic state depends on the exemplar mean. N is therefore not realistically equal to the number of speakers of an entire language, because there are likely to be many smaller communities with dialects. It is possible to extend the above model to account for geographically or socially separated groups but this is beyond the scope of the current paper. Letting $\mathbf{W}_i \in \mathbb{R}^2$ be a standard two dimensional Brownian motion then we have

$$\sqrt{\delta t} \mathbf{Z}_i \stackrel{d}{=} \mathbf{W}_i(t + \delta t) - \mathbf{W}_i(t) := \delta \mathbf{W}_i \quad (36)$$

where $\stackrel{d}{=}$ denotes equality in distribution. We may therefore write our discrete SDE (35) as

$$\delta \mathbf{x}_i = (\bar{\mathbf{x}}_i - \mathbf{x}_i) \delta t + \sqrt{\frac{2s^2}{N}} \delta \mathbf{W}_i \quad (37)$$

$$= (\mathbf{f}_i^{\text{att}} + \mathbf{f}_i^{\text{rep}}) \delta t + \sqrt{2D} \delta \mathbf{W}_i \quad (38)$$

where

$$D = \frac{s^2}{N}. \quad (39)$$

From this we see that our phenomenological equation (17) is the continuous time equivalent of (38).

Figure 9 illustrates the difference between the exemplar model and its approximate form (the phenomenological model (17)). In Figure 9 we compare the shift in the exemplar mean calculated exactly using the acceptance probability (23), and the approximate shift calculated by expanding $p_i(\mathbf{x})$ to first order about $\hat{\mathbf{x}}_i$. We have

considered two phonemes each composed of a set of co-located frames, and set the template range to $\alpha = 0^+$, so we only see the repulsive component of the interaction. As expected the approximation converges to the exact result as the phoneme separation tends to zero. For larger separations the approximate shift falls away more quickly, creating a shorter range repulsive interaction. This difference in range appears because in (17) we are measuring interference effects on frame i using the relative density of anti-templates at the location of i , rather than in the outer reaches of its cloud, so anti-template frames must be closer to have an effect. By increasing the cloud radius relative to the exemplar model, and lowering the functional load, we can achieve a similar interaction force in both models. The two forms of the model are alternative but qualitatively similar ways of characterising how phoneme overcrowding affects vowel sound evolution, based on the same underlying ideas. We work with the approximate model because of its simpler form which is efficient to simulate and to analyse mathematically. However, we take account of the effects above when selecting the cloud radius.

V. COMPARISON TO REAL SYSTEMS

A. Defining vowel space

Although the traditional vowel chart is a wide-based trapezium, average formant values for vowel sounds suggest that the accessible region in (F_1, F_2) space is closer to triangular in shape, with /a/ forming the lowest apex. Figure 10 shows the mean formants of the *peripheral* (outermost) vowels of Northern Standard Dutch. In this particular language /o/ and /ɔ/ are very close, suggesting that the difference between them is captured by something other than their first two formants. In many other languages /o/ and /ɔ/ are not close, for example in American English $\Delta F_1 = 226$ Hz [50]. This highlights the fact that the meanings of the phonetic vowel symbols are not precisely defined in terms of any measurable quantity. Rather, they are a tool for describing the general structure of the phonemic systems of languages. The approximate correspondence between typical formant values and the traditional IPA vowel chart is all the more remarkable for this. A particularly notable discrepancy is the relative heights of /a/ and /ɑ/, which are identical in the standard IPA vowel chart but in reality appear to differ in their first formant.

Motivated by the above considerations, and to avoid introducing unjustified complexity into the model, we opt for a symmetric trapezoidal (approximately triangular) space, as is used to tabulate differences between vowel systems in large inventories [12, 52]. Figure 11 illustrates this and also shows the approximate position of the nine most common short vowels [53]. Their relative frequencies, together with example words, are shown in Table I. Letting the bottom left vertex of vowel space define the

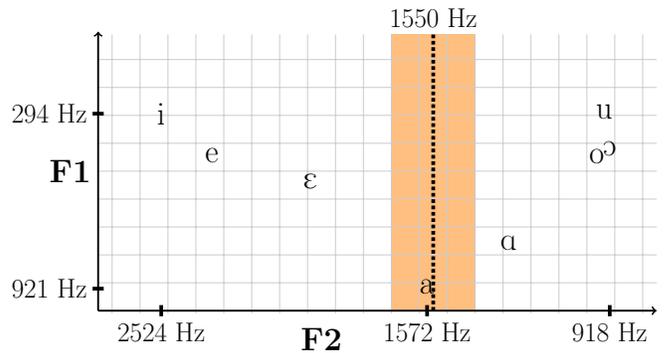


FIG. 10. Average formant values of the peripheral vowels /i,e,ε,a,ɑ,o,ɔ,u/ for female speakers of Northern Standard Dutch [51]. Marked formant frequencies give the positions of the extreme vowels /i,a,u/. Orange region shows the typical frequency range $F_2 \in [1400, 1700]$ for the second subglottal resonance in female speakers [4] (see section VI).

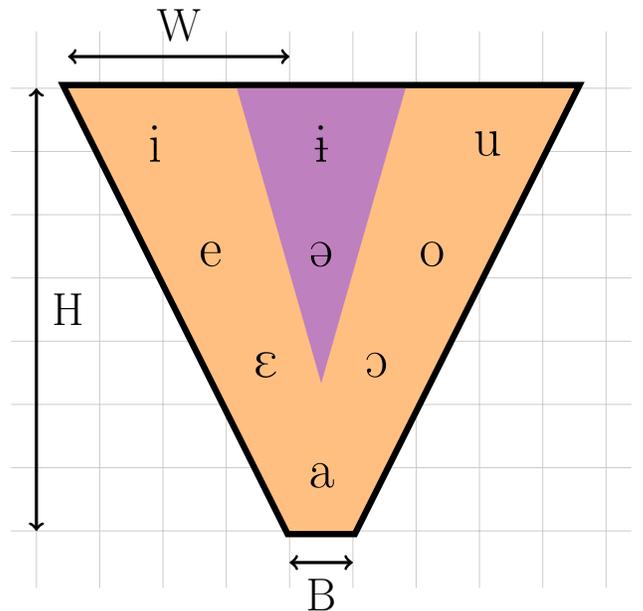


FIG. 11. Symmetric trapezoidal vowel space used in our model, together with the IPA symbols and approximate positions of the nine most common short vowels (see table I). Orange region is the *periphery* of vowel space, purple is the *interior*. Speakers of English may be surprised by the absence of low back vowels /ɑ/ (7 %) or /ɒ/ (4%) which appear in words like (pot). These are rather rare worldwide.

origin of coordinates, then their positions are

$$x = \frac{B}{2} \pm \frac{kW}{4} \quad (40)$$

$$y = \frac{H}{8} + \frac{kH}{4} \quad (41)$$

with $k \in \{0, 1, 2, 3\}$, with interior vowels /ə, i/ at $(B/2, 5H/8)$ and $(B/2, 7H/8)$.

In our model, repulsion between phonemes occurs as

IPA Symbol	Example word	Frequency
i (i:)	GenAm USA: ⟨bee⟩[bi:]	92% (32%)
u (u:)	GenAm USA: ⟨shoe⟩[ʃu:]	88% (29%)
a (a:)	RP UK: ⟨now⟩[nau]	86% (30%)
e (e:)	RP UK: ⟨bay⟩[bei]	61% (21%)
o (o:)	GenAm USA: ⟨go⟩[gou]	60% (21%)
ɛ (ɛ:)	GenAm USA: ⟨bet⟩[bet]	37% (11%)
ɔ (ɔ:)	RP UK: ⟨bore⟩[bɔə]	35% (10%)
ə (ə:)	RP UK: ⟨bear⟩[beə]	22% (4%)
ɪ (ɪ:)	S. African Eng: ⟨lip⟩[lip]	16% (1%)

TABLE I. The nine most common short vowels and their long versions. Frequency column shows the percentage of languages in the PHOIBLE database [53] which contain each vowel. Frequencies (and IPA symbols) of long forms are in brackets. GenAM and RP stand for “General American” and “Received Pronunciation”.

a result of cloud overlap, rather than from any notion of *contrast*. We therefore use formant cloud shapes to estimate the dimensions of our trapezium. It is clear both from Figure 1 and formant inventories for different languages [47] that clusters of formant data representing different phonemes within a single language, and the same phoneme across different languages, vary in size and shape. We characterise shape and position as follows. Given a large formant data set $\{F_{1i}, F_{2i}\}_{i=1}^N$, we let $S(X)$ denote the set of formants which were uttered for phoneme X . We define the F_1 mean and radius (standard deviation) of phoneme X to be

$$\mu_1(X) = \frac{1}{|S(X)|} \sum_{i \in S(X)} F_{1i} \quad (42)$$

$$\sigma_1(X) = \left(\frac{1}{|S(X)|} \sum_{i \in S(X)} F_{1i}^2 \right) - \mu_1^2(X) \quad (43)$$

with similar expressions for the F_2 mean and radius. Because of vocal tract size, average formant values for male and female speakers differ systematically. This elongates formant clouds at the population level. Since listeners subconsciously normalize for such differences [25], we consider only one sex: female. In Figure 6 the F_1 radius of phoneme clusters increases (approximately linearly) with increasing F_1 . If $\sigma_1(X) = \sigma_2(X)$ we say that X is spherical. Provided phonemes which are close enough to interact experience approximately the same systematic shape variation with position in formant space, then a transformation of this space which makes all phonemes spherical will have little effect on their overlaps. This assumption means we can model all frame clouds as spherical with the same cloud radius in this transformed space.

When frame and phoneme clouds are spherical, vowel arrangements which are near equilibrium with respect to repulsive forces are strongly affected by the aspect ratio

of the transformed vowel space

$$\text{AR} = \frac{\text{Max width}}{\text{Height}}. \quad (44)$$

We may estimate this ratio using the formant data in Figure 1. We define a *standardised distance* between the high vowels /i/ and /u/, which is approximately the same in both formant and transformed space

$$\Delta(u, i) := \frac{\mu_2(u) - \mu_2(i)}{\bar{\sigma}_2} \quad (45)$$

$$= \frac{1891}{227} \quad (46)$$

$$\approx 8.3 \quad (47)$$

where $\bar{\sigma}_2$ is the average radius of u and i. A similar calculation for front vowels gives

$$\Delta(\text{æ}, i) = \frac{\mu_1(\text{æ}) - \mu_1(i)}{\bar{\sigma}_1} \quad (48)$$

$$= \frac{580}{93} \quad (49)$$

$$\approx 6.2 \quad (50)$$

where $\bar{\sigma}_2$ is the average radius of the front vowels /i, ɪ, ε/, and /æ/. Since /æ/ is not the lowest vowel, but lies between /a/ and /ε/, then the true standardized height of vowel space is greater than this. According to the standard vowel chart /æ/ lies midway between /a/ and /ε/ so we approximate $\Delta(a, u) = 6\Delta(\text{æ}, u)/5$. Using our two standardized distances we can estimate the aspect ratio

$$\text{AR} \approx \frac{\Delta(u, i)}{\Delta(a, i)} = \frac{8.3}{7.4} \approx \frac{8}{7} \quad (51)$$

We take $H = 7$ and $B + 2W = 8$ so that a unit phoneme cloud radius would be consistent with the Figure 1 data. We set $B = 1$ to accommodate the lowest vowel, giving the vowel space dimensions shown in Figure 11 which we use for all simulations. Because the phenomenological model generates shortened interaction range we set the frame cloud radius to be $\sigma = 1.5$ in simulations, which generates interactions of comparable range to the exemplar version with unit radius (see Figure 9). We model the effect of vowel space boundaries using a repulsive force perpendicular to each boundary, of magnitude

$$|\mathbf{f}_i^{\text{bou}}| = \frac{E}{2} \left[1 + \tanh \left(\frac{|\Delta \mathbf{x}_i|}{w} \right) \right] \quad (52)$$

where $|\Delta \mathbf{x}|$ is the distance from the boundary, E is the maximum magnitude and w is the width of the boundary. We set $E = 10, w = 0.2$ in all simulations.

B. Properties of real vowel systems

Our sources of statistical information on vowel system properties are PHOIBLE [53], an online repository of

Category Repr.	Freq. of Repr.	Category Description	Members of Cat.	Frequ. of Cat.
i	96.1 %	High Front	i,ɪ,y,ʏ	99.4 %
u	92.4 %	High Back	u,ʊ	99.0%
a	91.4 %	Low	a,ɑ,æ,ɶ,ɛ	94.5%
e	74.2 %	Upper Mid Front	e,ø	74.8%
o	74.1 %	Upper Mid Back	o	74.1 %
ɛ	39.0 %	Lower Mid Front	ɛ,œ	40.7 %
ɔ	37.5 %	Lower Mid Back	ɔ,ʌ,ɒ	40.1 %
ə	23.9 %	Central	ə,ɤ,ɘ,ɜ,ɞ	29.1 %
ɨ	17.4 %	High Mid	ɨ,ɯ,ɥ,θ	26.9 %

TABLE II. The nine vowel categories used in our analysis. “Frequency of Representation” is the percentage of surveys in the PHOIBLE repository [53] containing the category representative (after stripping modifications). “Frequency of Category” is the percentage of surveys containing at least one category member. Categories members were selected based on their typical formant value proximity to the category representative.

phonological inventory data (containing 2186 languages); Ian Maddieson’s “Patterns of Sounds” [12], based on a representative sample of the world’s languages, contained in the UCLA Phonological Segment Database (UPSID), and Crothers’ vowel system typology [52], which sought to find a simplified classification of vowel systems.

Both Crothers and Maddieson use the idea of vowel quality, allowing them to identify sets of similar sounds as equivalent. We mimic this approach by identifying every vowel sound as a member of one of nine categories, each represented by one of the most common sounds shown in Figure 11. All of the most common systems with fewer than ten vowels identified by Crothers also used only these nine sounds [52]. The phonetic alphabet is equipped with an extensive notation for recording subtle modifications of the 28 symbols on the standard IPA vowel chart. For example surveys recorded in the PHOIBLE repository together contain 1094 symbols representing vowels. We firstly removed all modifying marks (producing 28 symbols) then used the mapping in Table II to assign category membership. For example /æ:/ is first stripped of its length mark, and then assigned to the category /a/. Our category assignment was based on estimated overlap in formant space, using typical formant values for the 28 IPA vowels. The frequencies in Table II also illustrate the percentage representation of each category amongst the surveys in the PHOIBLE repository. It is interesting to note the very high degree of front-back symmetry in these representations, which is reflected in our choice of system geometry.

Our categorisation scheme generates a *typology*: a classification of the world’s languages into groups according (in this case) to the structure of their vowel system. Of the $2^9 = 512$ possible system types in our scheme, only 106 are attested out of 3020 surveys (note: the number

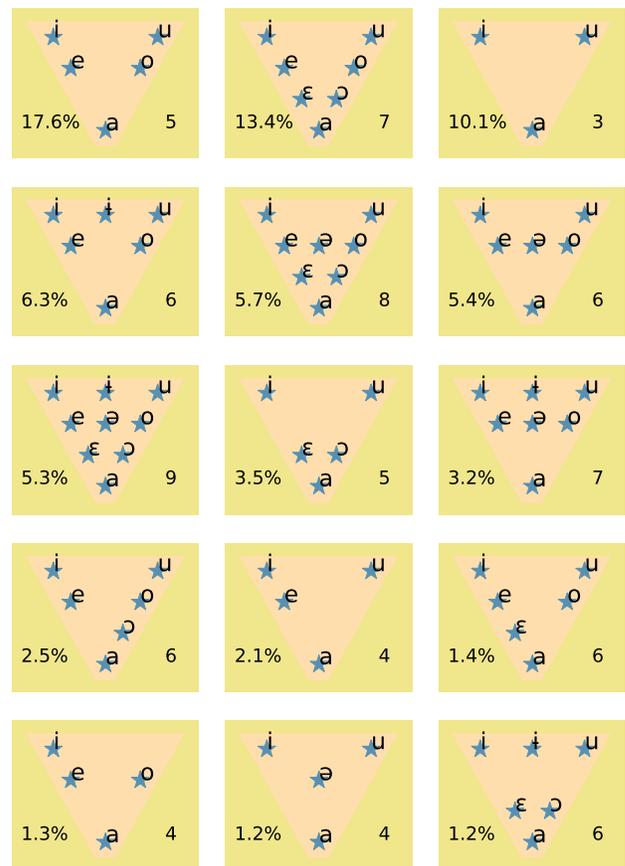


FIG. 12. The fifteen most common vowel system types following categorisation of PHOIBLE repository vowel data. Percentages give frequency amongst 3020 surveys, and integers give system size.

of surveys exceeds the number of languages due to dialectal variation [54]). This suggests that the dynamical processes which generate vowel systems are attracted to certain equilibria, and that many possible types are inherently unstable or unattractive to speakers. Figure 12 shows the fifteen most common vowel systems in the world’s languages, accounting for over 80% of the surveys. Eleven have perfect front-back symmetry, and of the four that don’t, two form a symmetrical pair. Similar results were obtained by Crothers [52], with eight of the eleven most common systems in his typology being symmetric.

The relative frequency of different system sizes (cardinalities) has also been of interest to linguists [12, 14, 52]. According to both Crothers and Maddieson, five vowel systems are more common than any other, with /a,e,i,o,u/ the most common of all. Our typology reproduces this result. However there is no unique solution to the problem of determining the relative frequencies of different vowel system sizes. Human societies in their “natural” pre-industrial state have low geographical connectivity, supporting many closely related languages

Num.Vowels	Maddieson	Crothers	PHOIBLE
2	-	-	0.2 %
3	5.4%	11.1 %	10.3%
4	8.5%	10.6%	6.8%
5	30.9%	30.8 %	22.7%
6	18.9%	19.2 %	16.5 %
7	14.8%	13.5 %	16.8 %
8	5.4%	4.3 %	7.5 %
9	7.9 %	7.2 %	8.5 %
10,11,...	8.1 %	3.4 %	10.6 %

TABLE III. Percentage of languages by their vowel system sizes (number of qualities) estimated by Maddieson [12], Crothers [52] and using our typology, based on PHOIBLE [53].

and dialects in relatively small areas. A nation state destroys this variation, replacing it with a *national language*. Regions such as Papua New Guinea, Central and West Africa and Australia therefore contribute a disproportionately high number of languages. This problem is typically addressed by sampling uniformly from language family groups [12] and from geographical areas [52]. Ambiguities can also result from the need to define vowel quality. Despite this, in Table III we see that our results (all PHOIBLE surveys with no adjustments for language size or family) are in broad agreement with those of Maddieson and Crothers.

The final statistical property that we consider is correlation between different sounds (two-point functions, in statistical mechanics). Given a category X , we define the indicator function that it is present in survey ℓ

$$S_X(\ell) = \begin{cases} 1 & \text{if } X \in \ell \\ 0 & \text{if } X \notin \ell. \end{cases} \quad (53)$$

The indicator is a binary variable, and the correlation between two categories is

$$\phi_{XY} := \frac{\langle S_X S_Y \rangle - \langle S_X \rangle \langle S_Y \rangle}{\sqrt{(\langle S_Y^2 \rangle - \langle S_X \rangle^2)(\langle S_Y^2 \rangle - \langle S_Y \rangle^2)}} \quad (54)$$

where $\langle \cdot \rangle$ denotes the average over all surveys. In statistics this correlation is called the *phi coefficient*, but it is also equal to the Pearson Correlation Coefficient. Table IV shows these correlations calculated from PHOIBLE, as well as simulated values. The strongest correlations are between front and back vowels of the same height: that is, if we have one of a front-back pair at given height, then we are likely to have the other. After these, there are weaker correlations between the mid central category and its surroundings, and in particular between the high and mid central categories i and ϵ . We will discuss the significance of this relationship in section VI.

Rank	Pair	ϕ (PHOIBLE)	Pair	ϕ (Model)
1	$\text{ɔ, } \epsilon$	0.74	i, o	0.47
2	$\star e, o$	0.71	i, e	0.42
3	i, u	0.28	$\star e, o$	0.32
4	$\star \text{ə}, i$	0.19	$\star \text{ə}, \epsilon$	0.32
5	$\text{ə}, \text{ɔ}$	0.18	$e, \text{ɔ}$	0.29
6	$\star \text{ə}, \epsilon$	0.15	$\star \epsilon, o$	0.23
7	$\text{ə}, e$	0.15	ϵ, i	0.23
8	$u, \text{ɔ}$	0.15	a, e	0.20
9	$\star \text{ə}, o$	0.15	$\star \text{ə}, o$	0.16
10	$\star o, \epsilon$	0.15	$\star \text{ə}, i$	0.15

TABLE IV. The ten largest correlation (ϕ) coefficients between categories in the PHOIBLE database and the model. Starred pairs appear in both lists. Model parameters $\sigma = 1.5, \gamma = 0.5, D = 0.025, n = 150$.

C. Model behaviour and cross-linguistic comparison

To investigate the behaviour of our model, and compare to real systems, we generate frame distributions beginning from randomized starting configurations where frames form a constant intensity Poisson point process [55] within vowel space. We then run the model for sufficient time (25 lifetimes) for a set of phoneme clusters to form, and to settle into a quasi stationary state. Stabilisation consists of frame movement in response to repulsive interactions, and occasional *mergers* between phonemes. More rarely, we see *splitting* events, where groups of frames spontaneously break away from a phoneme. When only one or two frames break away, we call this *evaporation*. In order to determine what vowel qualities our model has produced we perform a mean shift clustering [56] of the locations of frames once a stable state is reached. The mean shift algorithm locates the peaks of a kernel density estimate [57] of the distribution of frames within vowel space. Peaks of this density represent the centres of phonemic clusters. Having identified the locations of our phonemes, we then assign each to one of our nine categories (Table II) by proximity to the positions of their representative sounds, shown in Figure 11. In some cases we find that two simulated peaks are assigned to the same quality: we interpret this as a primitive form of allophonic variation, and count both peaks as a single phoneme. Figure 13 shows a superposition of the phoneme locations computed by this process. There are dense clusters of phonemes at the vertices of the space, corresponding to the three most common vowel qualities $/a/, /i/, /u/$.

The parameter with greatest influence on how many vowels form is the template range α . We can obtain an approximate lower bound on this parameter by considering the threshold for (first) formant frequency discrimination in normal speech, which is $\Delta F_1 \approx 50 \pm 10$ (Hz) [58], where $F_1 \in [235, 850]$ [59]. In standardised coor-

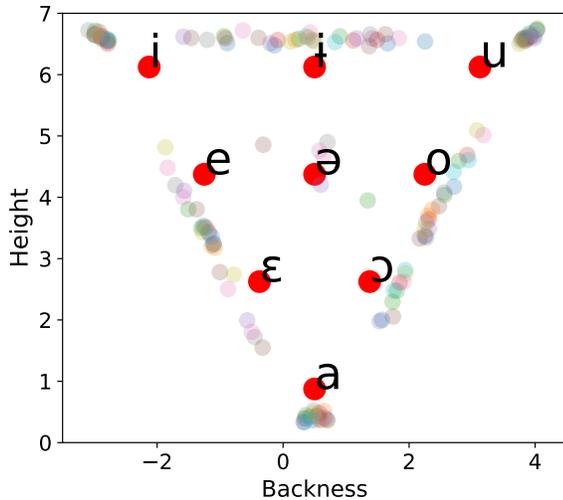


FIG. 13. Red dots show positions of standard vowels used for classification. Each coloured dot shows the position of a single phoneme obtained by mean shift cluster analysis of quasi-stable state of $n = 150$ frame system. Parameter values: $D = 0.025, \sigma = 1.5, \gamma = 0.5, \alpha = 1$. Plot superimposes 20 different simulated vowel systems.

dinates (dividing by $\bar{\sigma}_1$) this corresponds to a template range $\alpha_{\min} \approx 0.5$. It is likely that the acoustic proximity that defines two sounds as equivalent will vary between languages, because phoneme clouds and their separations vary between languages [47]. Evidence for this is provided by vowel-to-vowel co-articulation [60] where the articulatory requirements (configuration of the vocal articulators) for a phoneme are anticipated during the production of a previous phoneme. Such interactions between phonemes at different positions in a word lead to greater variability in how the same phoneme is pronounced in different words. For example, in [60], utterances of the form /apV/ where V is a vowel which gives the *context* of /a/ were analysed to discover the extent to which the choice of V altered the average formants of /a/ in three Bantu languages, Ndebele, Shona and Sotho. Whereas /a/ is relatively isolated in the vowel systems of Ndebele and Shona (/i,e,a,o,u/) it is relatively crowded in Sotho (/i,e,ε,a,ɔ,o,u/). Formant experiments showed that coarticulatory effects on /a/ were considerably greater in Ndebele and Shona than in Sotho. That is, a greater range of sounds were used as if they were the same phoneme in the languages with fewer vowels. In our model, such languages would have a greater template range. Additional factors used for contrast beyond the first two formants may also make precision in these formants less important, therefore altering α . We explored this effect (Figure 14) by generating the mean number of phonemes for a series of template ranges. From this we see that our estimate for α_{\min} is consistent with the observation that very few (3%) of all languages have more

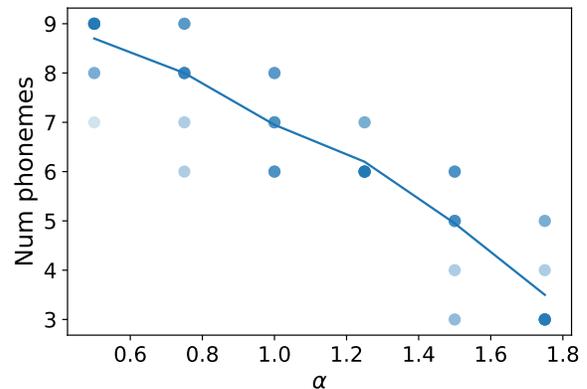


FIG. 14. Average number of spontaneously formed phonemes vs. template range, for a system of $n = 150$ frames. Parameter values: $D = 0.025, \sigma = 1.5, \gamma = 0.5$. Dots show underlying data with darker shading indicating repeated data points.

than 9 vowel qualities [52].

At high template range we find that most systems have only three vowels, typically /i/, /a/, /u/, consistent with observed systems. Almost no real languages have fewer than three vowel qualities. (Table III).

From Figure 14 we see that with a template range in the interval $1 \leq \alpha \leq 1.5$ we obtain vowel system sizes in the range $\{3, \dots, 8\}$, capturing $\approx 84\%$ of the size variability across world languages [12]. However, α is defined exogenous to the model, and we do not have means to estimate its distribution, other than by comparison to the distribution of vowel system sizes. In order to compare the relative frequencies with which different vowel phonemes appear in our model to their frequencies in real languages, we use Maddieson's data in Table III and [12]. We tabulate all 120 simulated vowel systems used to create Figure 14, and then compute the empirical probability mass function over the nine vowels for each size of system. We can then compute relative vowel frequencies over all system sizes as an average over our fixed size mass functions, weighted by the frequencies of each mass function in the world's languages. The results are given in table V. The predictions of the model typically lie within $\approx 10\%$ of the observed values except for the high and central vowels /i/ and /ə/. As with other dispersion theories, our model predicts that a high central vowel should appear with high probability, when in real languages it is rather rare. We return to this problem in section VI.

Table VI shows that apart from small height variations in the mid vowels /e, ε/ and /o, ɔ/ our model, the data from PHOIBLE and the Crothers study agree on the most common vowel systems. As we noted earlier, these height variations are subject to the interpretation of individual linguists, and should not be thought of as corresponding to a precisely defined formant interval. In less common systems there is more disagreement between

Category	Obs. Freq.	Pred. Freq.	Quantal Pred.
i	99%	98%	93%
u	99%	97%	93%
a	94%	94%	87%
e	75%	64%	67%
o	74%	56%	68%
ɛ	41%	36%	39%
ɔ	40%	44%	36%
ə	29%	11%	1%
ɨ	27%	64%	36%

TABLE V. Predictions (compared to data from table II) for the frequencies of the nine most common vowel qualities, with and without quantal effects. Simulation parameters: $n = 150$, $\sigma = 1.5$, $\gamma = 0.5$, $D = 0.025$. In the quantal case we have parameters $Q = 2$, $\omega_x = 2$, $\omega_y = 1$.

#	Rnk	PHOIBLE	%	Crothers	%	Model	%
3	1	a,i,u	93	a,i,u	100	a,i,u	62
3	2	a,i,o	2	-		ɛ,i,u	17
3	3	a,e,o	2	-		ɔ,i,u	12
4	1	a,e,i,u	28	a,ɛ,i,u	59	a,e,i,u	24
4	2	a,e,i,o	17	a,i,u,i	41	a,o,i,u	24
4	3	a,i,ə,u	16	-		a,ɛ,i,u	16
5	1	a,e,i,o,u	67	a,ɛ,i,ɔ,u	86	a,ɛ,i,o,u	28
5	2	a,ɛ,i,ɔ,u	13	a,ɛ,i,o,i	8	a,e,i,o,u	17
5	3	a,ɛ,i,o,u	4	-		a,i,o,u,i	13
6	1	a,e,i,o,u,i	31	a,ɛ,i,ɔ,u,i	73	a,e,i,o,u,i	55
6	2	a,e,i,o,u,ə	27	e,i,o,u,ɛ,ɔ	18	a,e,i,ɔ,u,i	20
6	3	a,e,i,o,u,ɔ	12	-		a,ɛ,i,o,u,i	18
7	1	a,e,i,o,u,ɛ,ɔ	64	a,e,i,o,u,ə,i	50	a,e,i,o,u,ɔ,i	41
7	2	a,e,i,o,u,ə,i	15	a,e,i,o,u,ɛ,ɔ	46	a,e,i,o,u,ɛ,i	36
7	3	a,e,i,o,ɔ,ə	4			a,ɛ,i,ɔ,u,ə,i	8

TABLE VI. The most common vowel systems of each size in PHOIBLE [53], Crothers [52] and the model. Frequency (%) columns show the percentage of vowel systems of the given size which have the given form. Simulation parameters: $n = 150$, $\sigma = 1.5$, $\gamma = 0.5$, $D = 0.025$.

the three typologies. In particular, the 7 and 8 vowel systems generated by the model are unrealistically likely to possess a high central vowel. We conclude that while the model matches the broad characteristics of real systems, there are details which it fails to match. At the same time, the question of which typology is correct may not have an answer.

We now consider the effect of functional load. The *functional load hypothesis*, first proposed by Gilliéron (1918) [48, 61], is that the probability of phoneme loss is inversely related to the amount of “work” done by the phoneme in identifying words. One simple measure of this work is the number of minimal pairs that a phoneme

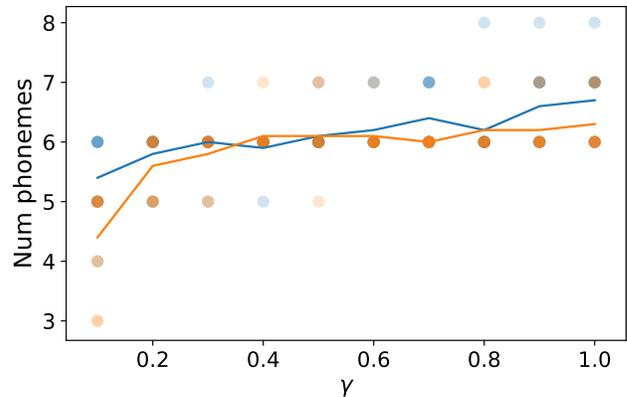


FIG. 15. Average number of spontaneously formed phonemes vs. functional load, for a system of $n = 150$ frames. Parameter values: $\sigma = 1.5$, $\alpha = 1.25$. Blue curve $D = 0.025$. Orange curve $D = 0.05$.

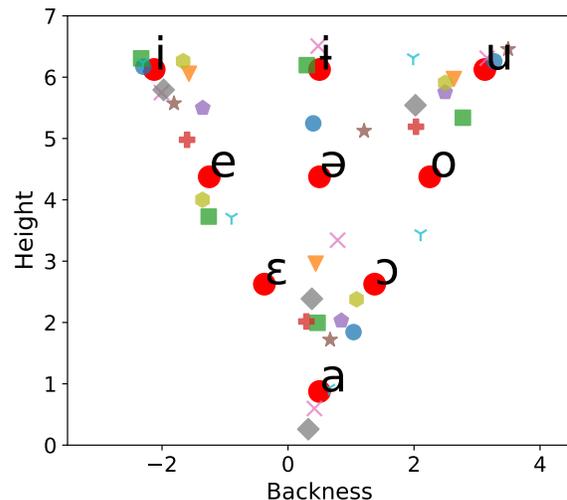


FIG. 16. Phoneme locations for ten simulations of $n = 150$ using low functional load ($\gamma = 0.1$). Other parameter values: $\sigma = 1.5$, $\alpha = 1.25$, $D = 0.05$. Each set of symbols represent the phonemes from a single simulation.

distinguishes, and a recent cross-linguistic corpus study [48], has shown that phonemes which define more minimal pairs are less likely to merge. In our model, phoneme merger can occur if the peripheries of two clusters are closer than the template range. Peripheral frames from the two clusters are attracted, pulling the clusters progressively closer. Merger is more likely if inter-phoneme forces are weaker, and from the definition of the repulsive force (16), we see that reducing functional load γ weakens them. Since vowel system formation consists of sequential cluster merging (and splitting) then we expect low functional load to result in smaller phoneme inventories. Figure 16 shows the strength of this effect in our

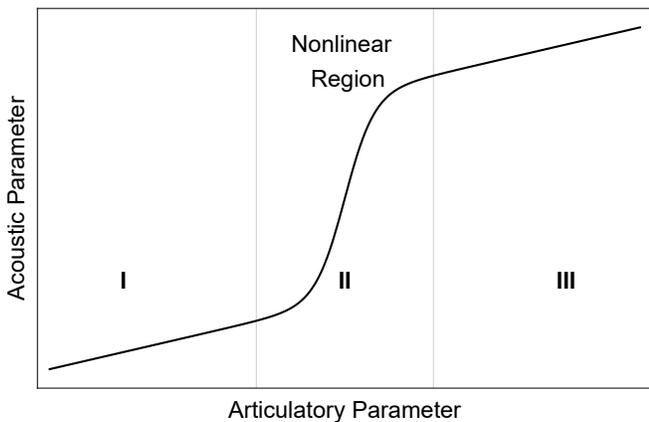


FIG. 17. Nonlinear relationship between articulatory and acoustic parameters. The nonlinear region II divides acoustic/articulatory space into two *quantal regions* I and III, where speech output is less sensitive to changes in articulation.

model for two levels of stochastic noise $D \in \{0.025, 0.05\}$. We find that the effects on system size (and therefore merger probability) are systematic but weak. With lower noise, the extremes of functional load produce an average difference of one phoneme, and for higher noise the difference is larger, but seen only at very low functional load where weak repulsive forces between phonemes allow some new systems to form (Figure 15). For example we have three different four vowel systems containing a mid central vowel, matching the 14th most common system in PHOIBLE (Figure 12). We also see that doubling the stochastic diffusion D , has a subtle but systematic effect on system size. Stochastic effects are required to bring two phonemes within merger range, so higher diffusivity results in more mergers, and therefore smaller systems.

VI. EFFECT OF THE SUBGLOTTAL RESONANCE

Our model, in its simplest form, may be viewed as a form of dispersion theory [13, 19]. It explains the placement of vowel phonemes in terms of a force which acts to maximize the acoustic distances between them. Patterns of sounds for which these forces are in equilibrium, or near it, are more likely to be observed in the model, helping us understand why certain vowel sounds, and combinations of sounds, are more common than others. An alternative explanation, *quantal theory* [4, 11, 43, 62], begins from the observation that the relationships between articulatory configurations and acoustic outputs of the human vocal apparatus contain pronounced nonlinearities where small changes in articulatory parameters can generate relatively large changes in acoustic output (Figure 17). The central idea of quantal theory is that these non-linearities quantize acoustic space into separate stable regions where the effects of changing articulatory parameters are small and predictable. Within these

regions speakers can more reliably produce a desired output, increasing the efficiency of communication. It is argued that these *quantal regions* define the inventory of sounds used in human languages. Of particular interest to us is the relation between F_2 and tongue “backness” which has been experimentally observed in the acoustic signals of diphthongs - continuous sounds which begin as one vowel and end as another. When pronouncing back-front diphthongs (for example /aɪ/), F_2 increases with time as the tongue moves forward. At $\approx 1400\text{Hz}$ F_2 jumps rapidly by around 50 – 300Hz [62]. This is caused by a coupling between the oral and subglottal cavities, and occurs near the second subglottal resonance (see Figure 10). The typical magnitude and location of the jump is predictable using a simple two-tube acoustic model of the cavities [62].

Because quantal effects arise via the map \mathbf{g} from articulatory to acoustic space, we temporarily switch our attention to the distribution of utterances in articulatory space. We consider a two dimensional articulatory domain, with two dimensions *backness* (y_1) and *height* (y_2). It appears that listeners mentally compensate for the effects on formant values of age and sex so that they perceive utterances of the same phoneme by a small and a large person as essentially the same sound [63]. We therefore assume that both articulatory coordinates $\mathbf{y} = (y_1, y_2)$ and their acoustic counterparts $\mathbf{x} = (x_1, x_2)$ are normalised, so we can think of all speakers as being the same size. Because the utterances of each frame in acoustic space form a cloud ($\psi_i(\mathbf{x})$), so must the articulatory parameters which generated them. In the definition of our model (section III), we assumed that when uttering the sound in frame i , speakers on average used articulatory parameters

$$\mathbf{y}_i := \mathbf{g}^{-1}(\mathbf{x}_i) \quad (55)$$

with normally distributed variations. Here \mathbf{g}^{-1} is the inverse of the articulatory to acoustic map. We write the *articulatory cloud*

$$\phi_i(\mathbf{y}) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(y_1 - y_{i1})^2 + (y_2 - y_{i2})^2}{2\sigma^2}\right) \quad (56)$$

$$= \frac{\exp\left(-\frac{(y_1 - y_{i1})^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma} \times \frac{\exp\left(-\frac{(y_2 - y_{i2})^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma} \quad (57)$$

$$:= \phi_{i1}(y_1)\phi_{i2}(y_2) \quad (58)$$

where we have assumed that the units of articulatory parameters are chosen so that the cloud radii, σ , in the two spaces are the same in those (quantal) regions where the map \mathbf{g} is affine. Since the jump in F_2 is generated by front-back movement of the tongue, parameterised by y_1 , we can write the map \mathbf{g} as

$$\mathbf{g}(\mathbf{y}) = \begin{bmatrix} g_1(y_1) \\ a + y_2 \end{bmatrix} \quad (59)$$

where $g_1(y_1)$ is a nonlinear function which captures the effect of the subglottal resonance. To determine the

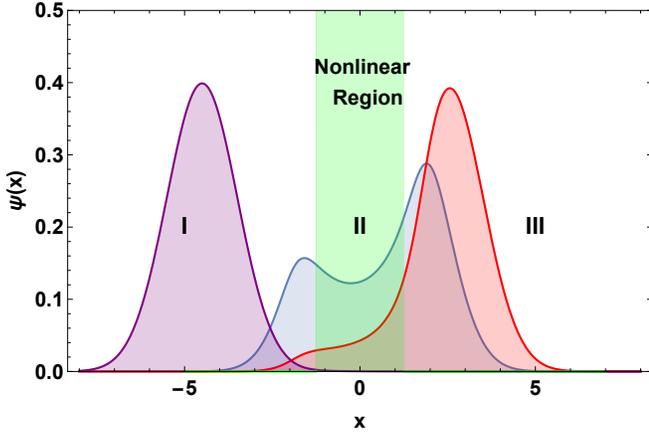


FIG. 18. Distortion of the acoustic frame/phoneme cloud as a subglottal resonance at the origin is approached (red and green distributions). Each cloud is normal in the articulatory parameter (with $\sigma = 1$), and away from the nonlinear region it remains normal (purple cloud). The parameters of g_1 , given by equation (65) are $j = 1$, $w = 0.5$.

shape of the acoustic cloud corresponding to $\phi_i(\mathbf{y})$ we let (Y_1, Y_2) be random variables drawn from ϕ_i . The corresponding acoustic variables are then

$$X_1 = g_1(Y_1) \quad (60)$$

$$X_2 = a + Y_2. \quad (61)$$

Because we have assumed that the articulatory cloud is spherical, having covariance matrix $\Sigma = \sigma \mathbf{I}$, then $Y_1 \perp Y_2$ and also $X_1 \perp X_2$ so the cloud density in acoustic space also factorizes

$$\psi_i(\mathbf{x}) = \psi_{i1}(x_1)\psi_{i2}(x_2). \quad (62)$$

Because the map \mathbf{g} is linear in the height variable, the acoustic height distribution is simply

$$\psi_{i2}(x_2) = \frac{\exp\left(-\frac{(x_2 - x_{i2})^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma}. \quad (63)$$

Changing random variables [55] in the nonlinear case gives the acoustic backness distribution

$$\psi_{i1}(x_1) = \left| \frac{d}{dx} g_1^{-1}(x_1) \right| \phi_{i1}(g_1^{-1}(x_1)). \quad (64)$$

In order to compute this distribution we use the following explicit form for g_1 (also used to generate Figure 17)

$$g_1(y_1) = y_1 + j \tanh(x/w), \quad (65)$$

which is a linear function throughout most of its domain, but with a step of height $2j$ and width w , centred at the origin. As $w \rightarrow 0$, the step becomes a sharp discontinuity. Figure 18 shows the effect of the resonance on the acoustic backness distribution. In the quantal

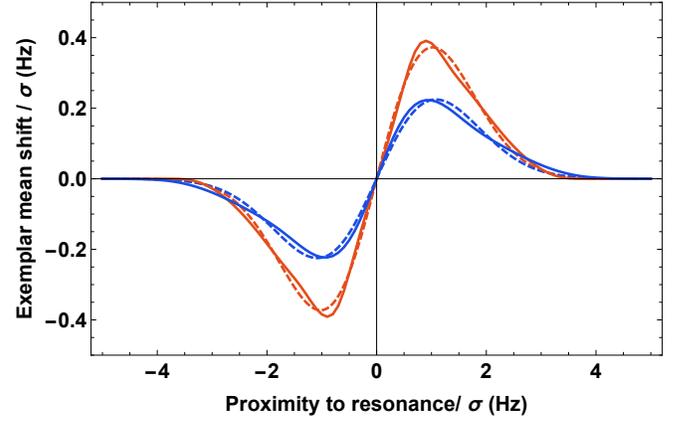


FIG. 19. Solid blue curve shows shift in exemplar mean with threshold $R = 2\sigma$, generated by jump in g_1 of magnitude $j = 1$ and width $w = 0.5$, defined in equation (65). Solid red curve shows exemplar shift for a sharper discontinuity ($w = 0.2$). Dashed curves show least squares fits of approximate quantal force function (68). Fitting parameters: Blue: $Q = 0.34$, $\omega_x = 1.53$, Red: $Q = 0.60$, $\omega_x = 1.45$. In both cases we have assumed we are at the top of vowel space of $x_2 = x_2^*$.

regions, away from the nonlinearity, the acoustic cloud remains normal. However, as the nonlinear region is approached, some utterances have articulatory parameters which cross the jump point, creating sounds with extreme acoustic characteristics compared to the average acoustic value of the phoneme. If unusual sounds are rejected by learners, as we have assumed when defining the acceptance probability (23), then the mean of the set of accepted sounds will shift. For a single phoneme composed of co-located frames, assuming that repulsion effects have had sufficient time to isolate it, the acceptance probability for one of its frames, i , becomes approximately $p_i(\mathbf{x}) = \mathbf{1}_{|\mathbf{x} - \mathbf{x}_i| < R}$, where R is the threshold beyond which sounds are rejected for being too unusual. The backness component of the exemplar mean is therefore

$$\bar{x}_{i1} = \frac{\int_{-R}^R x_1 \psi_{i1}(x_1) dx_1}{\int_{-R}^R \psi_{i1}(x_1) dx_1}. \quad (66)$$

The shift of the exemplar mean away from the mean position of the phoneme, defines an additional *quantal force*

$$\mathbf{f}_i^{\text{qua}} = \begin{bmatrix} \bar{x}_{i1} - x_{i1} \\ 0 \end{bmatrix}, \quad (67)$$

the backness component of which is plotted in Figure 19. Here we see that the discontinuity drives phonemes away from the subglottal resonance. We have also fitted an approximate quantal force curve with backness component

$$f^{\text{qua}} = Q(x_1 - x_1^*) \exp\left(-\frac{(x_1 - x_1^*)^2}{\omega_x^2} - \frac{x_2 - x_2^*}{\omega_y}\right), \quad (68)$$

where x_1^* is the position of the resonance in acoustic space and x_2^* is highest point in acoustic space. The height

dependent term in the exponent accounts for the fact that the quantal force requires a wide range of backness ($\approx 6\sigma$) in order to operate (Figure 18), and so we only expect to see its effects near the top of vowel space. From Fig 19 we see that this force curve gives a close match to the numerical calculation of the force. Since we do not know the exact form of the nonlinear map g_1 , and because computing the distorted clouds and exemplar mean requires numerical integration, we take (68) as a phenomenological definition of the force induced by the subglottal resonance. Combing this force with the attractive and repulsive components of our phenomenological model we obtain the quantal model with inter-frame force

$$\mathbf{f}_i = \mathbf{f}_i^{\text{att}} + \mathbf{f}_i^{\text{rep}} + \mathbf{f}_i^{\text{qua}}. \quad (69)$$

In the exemplar dynamics picture, the inter-frame forces will also be altered by the nonlinear map \mathbf{g} , because clouds become stretched as they approach the resonance, increasing the effective interaction range across it. Because the repulsive effect of the resonance will push phonemes away until their clouds no longer cross it, we assume that enhanced cross-resonance phonemic repulsion may be self consistently neglected in the steady state.

To examine the effects of the subglottal resonance we re-estimate the relative frequencies of each vowel quality in our model when a quantal force is present. In order to estimate the parameters of this force, we note from [62] that the F_2 jump can be up to 300Hz , which in standard coordinates gives $j \approx 1.5$. The width, w , of the jump region depends on the properties of the vocal tract. A more sudden jump increases the magnitude, Q , of the quantal force. Taking $j \in [1, 1.5]$ and $w \in [0.01, 0.5]$ gives force parameters $Q \in [0.5, 2.5]$ and $\omega_x \in [1.3, 2]$ when $\sigma = 1.5$. We have selected $Q = 2, \omega_x = 2$. Table V shows the predicted frequencies of each vowel when quantal forces are included. Because quantal forces drive phonemes away from the high central position /i/ we find that the frequency of this phoneme reduces to approximately half its non-quantal value, close to its frequency in real systems. However, the quantal force also appears to remove the mid central vowel in nearly all the systems generated by our sample. This happens because, although initially many systems have a central vowel, it is only stable if repulsive forces from boundary vowels are in balance. If the high central vowel is removed, stability is lost, and the mid central vowel migrates up and out to the system edge. In real vowel systems the pair /i,ə/ have the fourth highest correlation (Table IV), suggesting that the existence of one creates conditions which make the other more likely to persist. However, the question of what mechanism gives rise to stable central vowels is unresolved by our model. One possibility is that these sounds have low functional load, and are therefore less strongly affected by repulsion. Alternatively, some languages exhibit vowel systems which are in flux [64], and are therefore not equilibria of our dynamics. We address the question of how stable systems can spontaneously

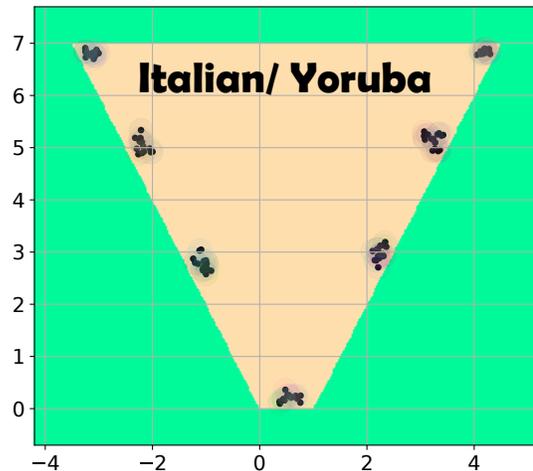


FIG. 20. A phonemic configuration matching the Yoruba and Italian languages [47] which is highly stable in the presence of a quantal force, but only marginally stable without. This is the second most common vowel system cross-linguistically, accounting for 13.4% of language surveys in the PHOIBLE database [53]. Parameter values $\alpha = 1, \sigma = 1.5, D = 0.025, \gamma = 0.5, Q = 2, \omega_x = 2, \omega_y = 2$.

change in section VIII.

The origins of quantal theory lie in a theory of sound patterns, developed by Chomsky, Halle and others [65], where each unit (segment) of speech is characterised by the presence or absence of a set of features [66]. For example, the height of a vowel is described by the two features $[\pm \text{high}]$ and $[\pm \text{low}]$, where \pm denotes the presence or absence of the feature. The back mid vowel /ɔ/ includes among its features $[-\text{high}, -\text{low}, +\text{back}]$. Phoneticians were motivated to search for a physical mechanism which could explain why it was possible to construct a successful phonological theory based on features like $[\pm \text{back}]$ when phonetically, backness appears to be a continuous variable [67]. The subglottal resonance provides a phonetic basis for the feature $[\pm \text{back}]$, because it divides vowel space into two stable regions. The relative scarcity of high and mid central vowels in real systems is consistent with this explanation, and formant based studies have shown that the resonance provides a reliable boundary between front and back vowels [67]. Within our model, the second most common vowel system seen in real languages (Figure 12), of which Italian and Yoruba are examples, requires quantal effects for its long term stability (Figure 20). Without this force, the upper mid vowels /e,o/ are only marginally stable, and if stochastic effects cause one of them to migrate away from the boundary of vowel space, then repulsion from the remaining boundary vowels will push it upwards to become /i/.

VII. ALLOPHONIC VARIATION

A phoneme is a set of sounds which are never *contrastive*. That is, if two sounds are part of the same phoneme then exchanging them within a frame cannot change the meaning of that frame. The allophones of a phoneme are subcategories of that phoneme which are acoustically distinguished from one another, and are used predictably in different contexts. That is, allophones are in *complementary* distribution with one another. Acoustic parameters are not intrinsically allophonic or phonemic. For example, in Australian English there are minimal pairs which differ only in vowel length as in the words ⟨cut⟩[kʌt] and ⟨cart⟩[kɑ:t] implying that /ʌ/ and /ɑ:/ are different phonemes, whereas in English received pronunciation (RP) lengthening is an allophonic variation which occurs, for example, when a vowel is followed by a voiced consonant.

We model allophonic variation by defining each acoustic parameter to be either phonemic or allophonic, noting that these definitions may spontaneously change. For simplicity, we consider two parameters: height, x (phonemic) and length, z (allophonic). The set of vowel sounds which are phonemically equivalent to the sound in frame i are then the templates $S_i = \{k \text{ s.t. } |x_i - x_k| < \alpha\}$. We also define another, similar set based on the allophonic parameter $P_i = \{k \text{ s.t. } |z_i - z_k| < \beta\}$ where β is an allophonic template range. The set of frames which are in both S_i and P_i are then allophonically equivalent to i . We write this set

$$A_i = S_i \cap P_i. \quad (70)$$

When frame i is uttered the probability density for the phonemic parameter is, as before, $\psi_i(x)$, whereas we write the density of the allophonic parameter $\chi_i(z)$ which we can assume is Gaussian but with a different variance to ψ_i . The density of phonemic templates for frame i is then defined exactly as in the standard version of the model

$$\widehat{\psi}_i(x) := \sum_{j \in S_i} (1 + \omega \delta_{ij}) f_j \psi_j(x) \quad (71)$$

whereas the density of allophonic templates is

$$\widehat{\chi}_i(z) := \sum_{j \in A_i} (1 + \omega \delta_{ij}) f_j \chi_j(z). \quad (72)$$

Phonemic and allophonic anti-template densities, $\widetilde{\psi}_i, \widetilde{\chi}_i$ are defined as sums of the summands above, but over S_i^c and $A_i^c \cap S_i$. As with the purely phonemic version of the model, we expect frames to be attracted to the mean of their phonemic templates \hat{x}_i and repelled from phonemic anti-templates. Similar arguments apply to interactions in the allophonic dimension, producing forces of identical form to (13) and (16), with the replacements $\psi_i(x) \rightarrow \chi_i(z)$ and $(S_i, S_i^c) \rightarrow (A_i, A_i^c \cap S_i)$. As a result, in this model we will not see interactions between allophones of

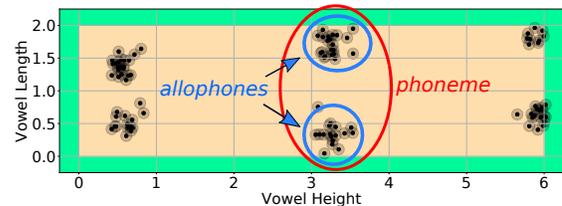


FIG. 21. Simulation of a vertical vowel system where vowel length is the allophonic parameter and $\alpha = 1, \beta = 0.5, \sigma_x = 1, \sigma_z = 0.5, D = 0.02, \gamma = 1$.

different phonemes: if a long allophone of /u/ gets longer then this will not affect a long allophone of /i/.

An example simulation is shown in Figure 21. Each set of allophones behave as their own sub-phonemic vowel system within a subspace parametrised by z_i . Using this example it is possible to see how a phonemic split might occur. Consider the central phoneme in Figure 21. The existence of the two allophones implies that there is some conditioning factor in the language (e.g. the voicing of a following consonant) which determines the phonemic environments in which each allophone should be used. If this factor disappears from the language then it will no longer be possible to predict which allophone to use in a given frame. The two allophones are then merely two different sounds in the language. More importantly, after the conditioning factor is lost, their functional load will increase because the contrast brought by the conditioning factor has gone. They may even distinguish minimal pairs. For this reason the allophonic parameter switches to being phonemic and we are left with two crowded phonemes, which will mutually repel, generating a phonemic split.

VIII. CHAIN SHIFTS, MOMENTUM AND WORD FREQUENCY EFFECTS

Beginning from a randomised initial state, our model will evolve over time into a stable system in which repulsive forces between phonemes are near equilibrium. Due to stochasticity it is in principle possible for one or more phonemes to move sufficiently far away from this equilibrium so that the system enters the basin of attraction of some other equilibrium configuration. A series of further phoneme movements will then ensue, until the new equilibrium is reached. Spontaneous sound changes from one stable state to another also occur in real languages, and if they involve multiple interacting sounds, they are referred to as *chain shifts* [35]. This terminology arises because multiple-vowel linguistic changes often occur in sequence with the movement of one vowel inducing another nearby to move, and so on. Chain shifts are traditionally divided into two classes: *push chains*, and *pull chains*, discussed below. Our model appears unrealistically stable when compared to real languages, generating

changes too infrequently. We will show below how sensitivity to an *age vector* [40–42] in the linguistic community can generate more realistic change processes.

A. Push and pull chains

One of the best known examples of a chain shift is *The Great Vowel Shift* which occurred in the English language between the time of Geoffrey Chaucer (1343-1400) and William Shakespeare (1564-1616) [68]. The shift affected the front and back long vowels of English, which moved upwards in vowel space, with the high vowels /i:/ and /u:/ shifting inwards and becoming diphthongs [68]

$$/i:/ \rightarrow /əɪ/ \quad (73)$$

$$/u:/ \rightarrow /əʊ/ \quad (74)$$

where the notation $A \rightarrow B$ indicates that after the shift, the frames which originally contained vowel A , contain vowel B . Two theoretical explanations exist for the mechanism which allowed this to happen. We use these alternatives to illustrate the difference between a push chain and a pull chain without commenting on which is more likely [17]. For simplicity, let us consider only the front vowels. A simplified summary of their movement is as follows

$$/a:/ \rightarrow /ɛ:/ \rightarrow /e:/ \rightarrow /i:/ \rightarrow /əɪ/. \quad (75)$$

This shift also involved a merger, so the final state of the frames using phonemes /e:/ and /ɛ:/ was /i:/. The *pull* chain explanation is that /i:/ changed first, creating a gap at the top of vowel space, into which the other lower vowels moved. That is, the *leading edge* of the chain moves first. The *push* chain explanation is that the back of the chain shifts first, overcrowding the sounds in front, and pushing them forwards in sequence.

A learning based explanation for push and pull (or *drag*) chains has been outlined by Labov [17], and an adapted form of this is summarised in Figure 22. Consider first the pull chain. In [17], the initial locations of the phonemes A and B (configuration 1) are described as stable, without need for the “mystery walls” that we have added. These walls represent other phonemes or the boundaries of vowel space. In our model the initial configuration would not be stable without these confining elements because the rejection of peripheral tokens would lead to a repulsion effect between A and B . If the wall confining A is removed then it will move away from B and B -tokens that were previously rejected start to be accepted, causing B to begin motion. The fact that the walls are needed in our model reveals a difference between it and Labov’s qualitative description of a pull chain. We require a *release of confinement* in order to produce a pull chain. That is, we must begin with a configuration where vowels are compressed; for example if there are four front vowels /e,ɛ,e,i/ held in place by repulsive forces from back vowels. If one of these, say /i/

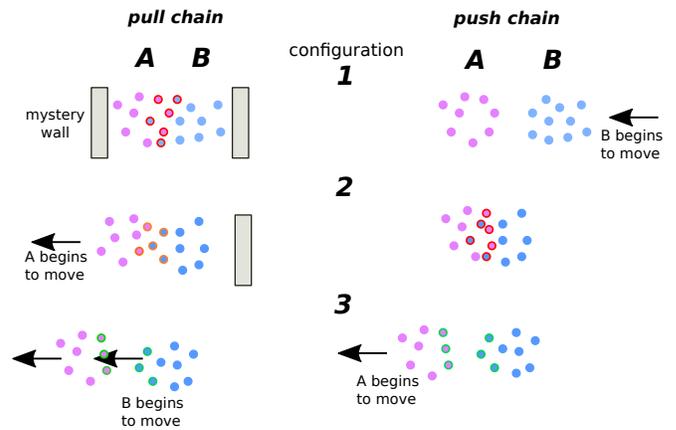


FIG. 22. The mechanics of push and pull chains, adapted from Labov [17]. Pink and blue dots are utterances of two different phonemes, A and B . Dots circled in red have a high probability of rejection as tokens used for learning. Orange circling indicates a lower probability of rejection. Green circle dots are unlikely to be rejected. Mystery walls are required to hold the phonemes together before the push chain begins.

“pops out” into the interior of vowel space then the repulsive forces between the remaining sounds will push them upwards to fill the front positions until inter-phonemic forces are in balance.

We now consider the push chain. The explanation in [17] begins with configuration 2 of Figure 22. As we have noted, this configuration is not stable in our model. We therefore assume that our two phonemes are initially separate and form part of some larger stable system. We then assume that something causes phoneme B to start moving (we propose a mechanism below). When B is sufficiently close to A , token rejection effects (repulsion) cause phoneme A to start moving away. The phoneme B is then said to have *pushed* phoneme A .

Within our description of vowel dynamics, the distinction between push and pull chains is to some extent unnecessary. Both process may be understood as the effects of repulsive interaction forces, but in different contexts: one in which compression is released, and one in which it is applied. In our view there is really no such thing as a pull chain; vowel shifts are all caused by pushing, and it is the start of a shift that determines whether it is referred to as a push or a pull.

In Figure 23 we have illustrated a pull chain in a one dimensional vowel space, which may be viewed as a toy model of the consequence of Labov’s *upper exit principle* [64], which we paraphrase as: “in chain shifting high peripheral vowels become non-peripheral.” In other words: high front vowels (/i,i:/) have a tendency to pop out of the top left corner of vowel space, as perhaps happened at the start of the Great Vowel Shift. We may view the one dimensional space in Figure 23 as a simple model of the front of two dimensional vowel space, in which phonemes are confined by the repulsion of other non-front vowels (see Figure 13). The four initial vowels in this space

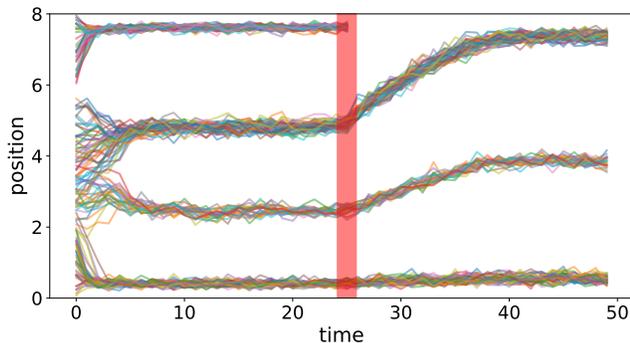


FIG. 23. An example of what is commonly called a pull or “drag” chain. Starting from randomized frame locations, four phonemes spontaneously formed. The highest is removed from the system at $t = 25$. Parameter values $\alpha = 1, \sigma = 1.5, \gamma = 1, D = 0.025$. One time unit = a single speaker’s expected life time. No momentum.

formed spontaneously, starting from randomised initial frame positions. At $t = 25$ lifetimes, we removed the highest vowel, mimicking a spontaneous shift of the form /i:/ \rightarrow /a/. Repulsive forces between the remaining three then pushed them upwards to fill out the empty space. The traditional interpretation would be that the movement of the top vowel *dragged* the lower vowels upward.

B. Momentum

Stochasticity of frame trajectories derives in our model from unpredictability in the language learning process, which at the population level is realised as a diffusion process for frame location. We will show in section IX that beyond a critical level of diffusivity, D_c , phonemic clusters cannot form. Assuming $D < D_c$ then the motion of the centroid of a phoneme consisting of N frames will undergo a diffusion process with coefficient $D^{\text{pho}} \propto N^{-1}$. The magnitude of D^{pho} also depends on the distribution of relative frame frequencies within the phoneme, with heavier tailed distributions leading to greater diffusivity, due to the dominance a smaller number of very common words. This purely diffusive behaviour is problematic from the point of view of empirical observations of sound change, which often self actuate before progressing monotonically (or nearly so) over a sustained period [42]. One explanation is that language learners are sensitive to the direction of change. This direction is observable from differences in language use between older and younger speakers, often referred to as the *age vector* of a linguistic feature, or its *momentum*. A number of studies have sought to model such effects [40–42] in terms of the relative frequency of a linguistic feature. Although their mathematical details differ, the essential idea is that when (new) speakers select their linguistic state, they are biased in the direction of the age vector.

It is straightforward to incorporate the momentum ef-

fect into our model. Working with a one dimensional acoustic space, where the position of the i th frame is given by $x_i(t)$, we define the linguistic memory of the community for frame i as a time average over its history

$$m_i(t) = \frac{1}{\tau} \int_{-\infty}^{\tau} x_i(s) e^{(s-t)/\tau} ds. \quad (76)$$

When $\tau = 1$ this memory is an average over the historical states of the community when each speaker was born. Differentiating with respect to time we obtain

$$\dot{m}_i(t) = \frac{1}{\tau} (x_i(t) - m_i(t)). \quad (77)$$

We define the difference between the current state and the memory as the *age vector*

$$\Delta_i(t) := x_i(t) - m_i(t). \quad (78)$$

We then define an additional *momentum response*

$$\psi(\Delta) = \theta \tanh\left(\frac{b\Delta}{\theta}\right) \quad (79)$$

which is added to the attractive (13) and repulsive (16) forces already in our model. This is the shift in the average sound learned by a new speaker, based on their tendency to emphasize “younger” forms of speech. An analogous response function appears in the frequency based models [40, 41], where it is termed a *prediction function* [40] or generates a *perceived frequency* [41]. The parameter b in (79) is the *momentum sensitivity* and θ is the *cut-off*, giving the maximum possible magnitude of the momentum driven rate of change. For $\Delta \ll \theta$ the momentum response is approximately linear $\psi(\Delta) \sim b\Delta$ as $\Delta \rightarrow 0$. The momentum model for a single frame may be written

$$dx_i = [f_i^{\text{att}} + f_i^{\text{rep}} + \psi(x_i - m_i)] dt + \sqrt{2D} dW_i \quad (80)$$

$$dm_i = \frac{1}{\tau} (x_i - m_i) dt. \quad (81)$$

We now analyse the stability of a single isolated phoneme consisting of N frames of equal frequency with self focus $\omega = 0$. Formally we consider the limit $\alpha \rightarrow \infty$, so the phoneme does not lose frames through evaporation. The template mean and the *template memory mean* are

$$\hat{x}(t) = \frac{1}{N} \sum_{i=1}^N x_i(t) = \langle x_i \rangle \quad (82)$$

$$\hat{m}(t) = \frac{1}{N} \sum_{i=1}^N m_i(t) = \langle m_i \rangle \quad (83)$$

where $\langle \cdot \rangle$ denotes the average over frames (the *phoneme average*). From the definition of the momentum model we have

$$d\hat{x} = \frac{1}{N} \sum_{i=1}^N [(\hat{x} - x_i) dt + \psi(x_i - m_i) dt + \sqrt{2D} dW_i] \quad (84)$$

$$= \langle \psi(\Delta_i) \rangle + \sqrt{\frac{2D}{N}} d\hat{W} \quad (85)$$

where \hat{W} is a standard Brownian motion. The analogous equation for the template memory mean is

$$d\hat{m} = \frac{1}{\tau}(\hat{x} - \hat{m})dt. \quad (86)$$

Subtracting (86) from (85) we obtain the age vector dynamics

$$d\hat{\Delta} = \left[\langle \psi(\Delta_i) \rangle - \frac{\hat{\Delta}}{\tau} \right] dt + \sqrt{\frac{2D}{N}} d\hat{W}. \quad (87)$$

Expanding the momentum response function about $\Delta_i = 0$ we obtain to order $\langle \Delta_i^3 \rangle$

$$d\hat{\Delta} = \left[\left(b - \frac{1}{\tau} \right) \langle \Delta_i \rangle - \frac{b^3}{3\theta^2} \langle \Delta_i^3 \rangle \right] dt + \sqrt{\frac{2D}{N}} d\hat{W}. \quad (88)$$

Writing $\Delta_i = \hat{\Delta} + \epsilon_i$ where ϵ_i is a zero mean random variable we have, if ϵ_i is also symmetric (so $\langle \epsilon_i^3 \rangle \approx 0$)

$$\langle \Delta_i^3 \rangle = \hat{\Delta}^3 + 3\hat{\Delta} \langle \epsilon_i^2 \rangle \quad (89)$$

so

$$d\hat{\Delta} = \left[\left(b - \frac{1}{\tau} - \frac{b^3}{\theta^2} \langle \epsilon_i^2 \rangle \right) \hat{\Delta} - \frac{b^3}{3\theta^2} \hat{\Delta}^3 \right] dt + \sqrt{\frac{2D}{N}} d\hat{W}. \quad (90)$$

For a phoneme with a large number of frames the drift term dominates the dynamics and we see that provided

$$b - \frac{1}{\tau} - \frac{b^3}{\theta^2} \langle \epsilon_i^2 \rangle < 0 \quad (91)$$

then the age vector has a stable fixed point at $\hat{\Delta} = 0$ so the phoneme will be subject only to a weak diffusion with $D^{\text{pho}} = D/N$. In this case we will not see any sustained movement in one direction. At a critical value of age vector sensitivity, the fixed point destabilizes and two stable fixed points appear at

$$\hat{\Delta} = \pm \sqrt{\frac{3\theta^2}{b^3} \left(b - \frac{1}{\tau} - \frac{b^3}{\theta^2} \langle \epsilon_i^2 \rangle \right)}. \quad (92)$$

The phoneme will select one of these at random, and then execute a sustained movement in that direction, until noise effects, interactions with other phonemes, or with the boundaries of vowel space cause it to change direction or return to stability. If sensitivity to the age vector fluctuates with time, or if its value is near the threshold for stability, then the phoneme may switch between periods of stability and instability. We note that it is in principle possible to measure the age vector within a speech community by considering differences in the vowel systems of old and young speakers, and to measure sensitivity, given sufficient longitudinal data.

To illustrate the effect of momentum we first consider the behaviour of a single phoneme driven entirely by diffusive dynamics, without any repulsive effects ($\gamma = 0$).

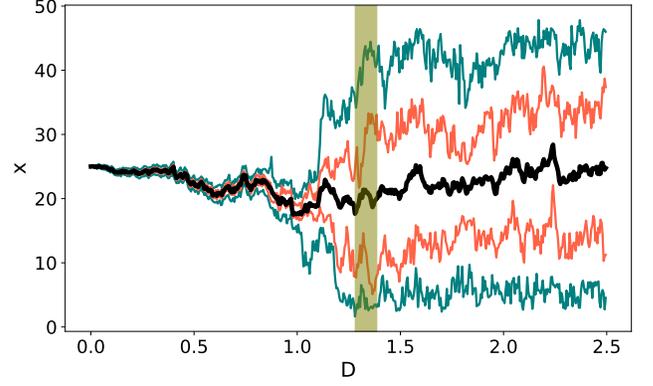


FIG. 24. Black line shows mean location of 100 frames, starting from position $x_i = 15$, as diffusion coefficient is increased from zero over $T = 5000$ lifetimes. System size $L = 30$. Red and blue-green lines show 70th (30th) and 90th (10th) percentile frame locations, which indicate distribution of frames within the system. Vertical olive band shows theoretical location of critical diffusion coefficient $D_c = \alpha^2/3$ where phoneme disintegrates. Other parameter values $\alpha = 2, \sigma = 1, \gamma = 0$.

We consider a one dimensional vowel space $x_i \in [0, L]$, beginning with all frames located at the same location $x = L/2$. Starting from zero diffusion coefficient, we gradually increase D over a long time interval ($T = 5 \times 10^3$ lifetimes), yielding the behaviour shown in Figure 24. While the the diffusion coefficient is less than the critical value D_c , the phoneme remains intact and of finite size. Diffusive changes in position during this phase of evolution generate a cumulative shift of ≈ 5 cloud radii over the first 2×10^3 lifetimes. Such a shift corresponds approximately to a vowel changing from “high” to “low” (e.g. /i/ \rightarrow /æ/). Taking a single lifetime as 50 years then this shift would take around ten thousand years to complete. This time scale is unrealistic. For example, the Great Vowel Shift took place over ≈ 300 years [68]. Moreover, the template range used in Figure 24 is at the upper end of realistic, allowing larger diffusion coefficients to be reached before phonemic destruction. As the destruction point is approached, the phoneme expands in size, before disintegrating entirely, leaving frames distributed approximately uniformly over the system. We conclude that diffusive dynamics alone is not sufficient to describe the pace of realistic sound changes.

We now consider the effect of momentum. In Figure 25 we have simulated a two vowel system where momentum sensitivity fluctuates over time. We model these fluctuations using Ornstein Uhlenbeck dynamics [32]

$$db_t = a(b^* - b_t)dt + \sigma_b dW_t \quad (93)$$

where b^* is the long run mean sensitivity, a is the reversion rate towards to this mean, and σ_b is the volatility of the sensitivity. We set $b^* = 0.9$, which is below the threshold for spontaneous sustained shifts. Initially the two phonemes are four cloud radii apart, producing very

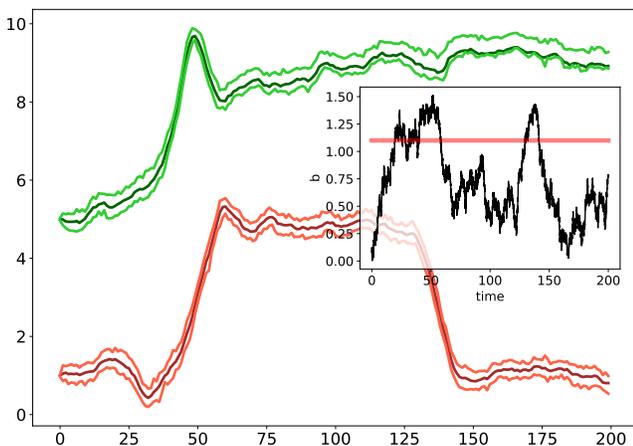


FIG. 25. Darker green and red lines show locations of two phonemes starting from positions $x = 1, 5$ in a system of size $L = 10$. Lighter lines give 90th (10th) percentile frame locations for these two phonemes, as an indication of size. Inset plot shows age vector sensitivity as a function of time. Horizontal red line shows estimated critical sensitivity $b \approx 1.1$. Model parameter values $\alpha = 1, \sigma = 1, \gamma = 1, D = 0.025, \theta = 0.5$. Parameters for mean reversion (Ornstein Uhlenbeck) process followed by momentum sensitivity are $b^* = 0.9, a = 0.05, \sigma_b = 0.1$.

weak interactions. When the sensitivity crosses the critical threshold, the lower vowel spontaneously begins rising, starting a push chain. The motion of the green vowel is halted by the boundary of vowel space, and the vowels enter a temporary equilibrium while the sensitivity again becomes sub-critical. A final move of the lower vowel is generated by a short lived supercritical period of sensitivity. We calculate the threshold b_c by first numerically estimating the variance of the age vector

$$\langle \epsilon_i \rangle \approx 0.07 \quad (94)$$

and then setting the right hand side of (91) to zero, and solving for $b_c \approx 1.1$. The two major shifts generated in this simulation took ≈ 10 lifetimes, corresponding to ≈ 500 years, in line with the Great Vowel Shift. For given τ the speed of the shift is controlled largely by the momentum response cut-off θ .

The one dimensional momentum model we have defined here is too simple to be directly comparable to the Great Vowel Shift, because in its current form it does not describe different lengths of vowel or diphthongs. However we can reproduce the essential properties of observed systems, serving a starting point for more sophisticated models which could be used to test hypotheses about the nature of historical vowel shifts, and to predict future changes.

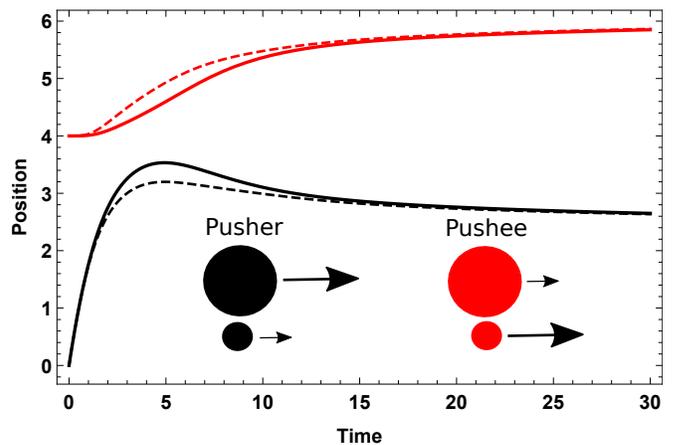


FIG. 26. The interaction of two phonemes (black and red), each with two frames. Solid lines: high frequency frames ($f = 0.45$), dashed: low frequency ($f = 0.05$). Parameter values $\sigma = 1, \omega = 5, \gamma = 1$. Initial age vector of black frame $\Delta = 2$. Momentum response parameters $b = 0.7, \theta \rightarrow \infty$.

C. Word frequency effects

Exemplar theory and empirical studies suggest [24] that high frequency frames change more slowly than low frequency frames when pushed by the movement of another vowel (a *push chain* [17]). The opposite occurs in the frames which are being pushed. We consider a system of two phonemes, each consisting of a one high and one low frequency frame. Initially the phonemes are separated by 4σ with one phoneme having an age vector, $\Delta > 0$ in the direction of the other, with sub-critical momentum sensitivity. This creates a simple two-frame push chain. Figure 26 shows the dynamics of the high and low frequency frames as they approach, in the noiseless limit $D \rightarrow 0$. We see that in the incident frame (the “pusher”) the high frequency frame moves faster, whereas in the “pushee” frame, the high frequency frame is slower to react. The effects are quite subtle, as they are in empirical data [69], and we note that the differing response of high and low frequency frames requires positive self focus. Without this the data used by learners is the same for all frames in a phoneme, so the dynamics of each frame are statistically identical. The magnitude of word frequency effects therefore provides a mechanism to infer the extra weight that listeners place on words as templates for the vowel sounds they contain, as compared to words which contain similar sounds. Finally we note an intuitive physical analogy. When $\omega > 0$, higher frequency frames behave as more massive particles which are less strongly influenced by the proximity of others.

IX. CONTINUUM LIMIT

So far we have considered languages with relatively small numbers of frames, each of which can be explicitly

simulated. We now consider the many-word limit where system dynamics may be described in terms of a continuum frame density. Let N be the total number of frames in our language and consider the limit $N \rightarrow \infty$ in one dimension, in which case we can define a frame density

$$\rho(x) = \lim_{\delta \rightarrow 0} (1/\delta) \sum_i f_i \mathbf{1}_{|x-x_i| < \delta/2}. \quad (95)$$

We suppress the time dependence of ρ for compactness. The set of frames which are templates for a frame located at x lie in the *template interval* $T_x := [x-\alpha, x+\alpha]$. In one dimension the total force on frame i , given the positions of the the other frames, may be written

$$\mu(x_i | \{x_j\}_{j \neq i}) = \hat{x}_i - x_i + \sum_{j \notin S_i} \frac{f_j \psi_j(x_i)(x_i - x_j)}{\gamma^{-1} \hat{\psi}_i(x_i) + \tilde{\psi}_i(x_i)}, \quad (96)$$

where we have used the symbol μ to avoid confusion with the frequency of frame i . In the limit $N \rightarrow \infty$ we have, for the attractive component of this force,

$$\hat{x}_i - x_i \rightarrow \frac{\int_{T_{x_i}} (u - x_i) \rho(u) du}{\int_{T_{x_i}} \rho(u) du} \quad (97)$$

and for the repulsive component

$$\begin{aligned} & \sum_{j \notin S_i} \frac{f_j \psi_j(x_i)(x_i - x_j)}{\gamma^{-1} \hat{\psi}_i(x_i) + \tilde{\psi}_i(x_i)} \\ & \rightarrow \frac{\rho(x) \int_{T_{x_i}^c} (x_i - u) \psi_u(x_i) \rho(u) du}{\gamma^{-1} \int_{T_{x_i}} \psi_u(x_i) \rho(u) du + \int_{T_{x_i}^c} \psi_u(x_i) \rho(u) du} \end{aligned} \quad (98)$$

where $\psi_u(x)$ is the $\mathcal{N}(u, \sigma^2)$ density; we consider the case $\sigma = 1$ and set the self focus $\omega = 0$ for simplicity. In the limit $N \rightarrow \infty$ we write the force on a frame at x as $\mu(x)$, with the dependence on ρ implicit. In this limit the density evolves deterministically provided the relative frequencies of all frames tend to zero as $N \rightarrow \infty$. Conditional on the density field, the location of frame i obeys the stochastic differential equation

$$dx_i = \mu(x_i) dt + \sqrt{2D} dW_i \quad (99)$$

and therefore the probability density function, $p_i(x, t)$, for its location obeys the Fokker Plank equation [33]

$$\partial_t p_i(x, t) = -\partial_x [\mu(x) p_i(x, t)] + D \partial_x^2 p_i(x, t). \quad (100)$$

To find the density field $\rho(x)$ we note that because it is deterministic

$$\rho(x) = \mathbb{E}(\rho(x)) \quad (101)$$

$$= \lim_{\delta \rightarrow 0} (1/\delta) \sum_i f_i \mathbb{E}(\mathbf{1}_{|x-x_i| < \delta/2}) \quad (102)$$

$$= \lim_{\delta \rightarrow 0} (1/\delta) \sum_i f_i p_i(x, t) \delta \quad (103)$$

$$= \sum_i f_i p_i(x, t). \quad (104)$$

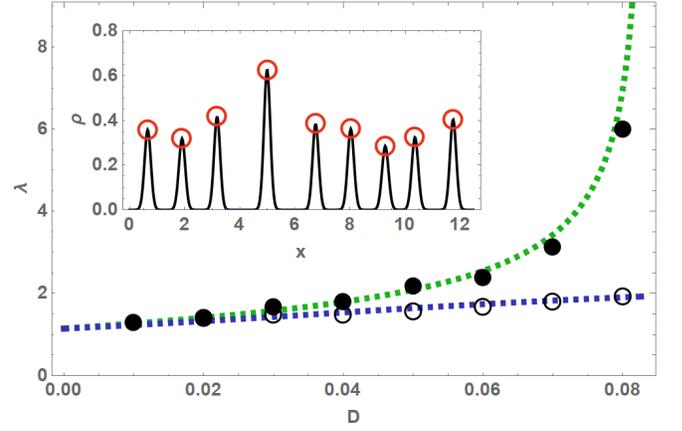


FIG. 27. Dots: estimated λ values (obtained by automated peak identification) when $\alpha = 0.5, \gamma = 0$. Circles: $\gamma = 1$. Initial fluctuations given by high frequency sine wave. Dashed lines: analytical predictions. Inset: solution to (105) for $\gamma = 0.5, D = 0.02$, giving $\lambda = 1.39$ cf. prediction $\lambda^* = 1.33$.

Multiplying both sides of (100) by f_i and summing over i we obtain

$$\begin{aligned} \partial_t \rho(x) = & D \partial_x^2 \rho(x) - \partial_x \left(\frac{\rho(x) \int_{T_x} (u - x) \rho(u) du}{\int_{T_x} \rho(u) du} \right) \\ & - \partial_x \left(\frac{\rho(x) \int_{T_x^c} (x - u) \psi_u(x) \rho(u) du}{\gamma^{-1} \int_{T_x} \psi_u(x) \rho(u) du + \int_{T_x^c} \psi_u(x) \rho(u) du} \right). \end{aligned} \quad (105)$$

This is the continuum evolution equation which defines our model in the limit of large numbers of words. We have made use of the fact that the contribution of individual frames to the density field is negligible so $\mu(x)$ is the same function for all frames.

Beginning from small fluctuations in frame density, solutions to equation (105) take the form of regularly positioned peaks representing distinct vowels (inset Figure 27). Expanding $\rho(u)$ to second order in the attractive term, and comparing to the magnitude of the diffusive flux, we obtain the lower bound $\alpha > \sqrt{3D}$ on template range for vowel formation. This relationship was tested in Figure 24.

We can calculate the typical number of spontaneously formed vowels in a system by finding the wave number at which density fluctuations are most susceptible to cluster formation. To achieve this we write equation (105) in terms of diffusive, attractive and repulsive fluxes

$$\partial_t \rho(x) = -\partial_x (j_D(x) + j_A(x) + j_R(x)) \quad (106)$$

and assume that our vowel system begins from a primordial state where the distribution of frames is approximately uniform with fluctuations which are spatially correlated only over very short ranges (no clusters of significant size). We write the state of the system at this early

stage

$$\rho(x) = c + \epsilon\eta(x, t) \quad (107)$$

$$= c + \frac{\epsilon}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{\eta}(\omega, t) e^{i\omega x} d\omega \quad (108)$$

where $c > 0$ is a constant and $\eta(x, t)$ is a zero-mean fluctuating field with Fourier transform $\hat{\eta}(\omega, t)$ and $\epsilon \ll 1$ is a small parameter measuring the magnitude of fluctuations. At small times t , the spatial correlation function

$$C(z, t) = \int_{-\infty}^{\infty} \eta(x, t) \eta(x + z, t) dx \quad (109)$$

decays rapidly with increasing $|z|$, and therefore has an energy spectrum $\sqrt{2\pi}|\hat{\eta}(\omega)|^2$ whose dominant contributions are from high wave numbers, ω [33]. As the system evolves and phonemic clusters form, the energy spectrum will begin to concentrate at lower wave numbers, and spatial correlations will decay more slowly. The peak, ω^* , of the energy spectrum gives the typical separation of clusters as

$$\lambda^* = \frac{2\pi}{\omega^*}. \quad (110)$$

To estimate the location of the peak we consider the behaviour of the diffusive, attractive and repulsive fluxes as the clusters begin to form. To lowest order in ϵ we have, from (105), as $\epsilon \rightarrow 0$

$$j_D(x) \sim -\epsilon\eta'(x) \quad (111)$$

$$j_A(x) \sim \frac{\epsilon}{2\alpha} \int_{S_x} (u - x)\eta(u) du \quad (112)$$

$$j_R(x) \sim \frac{\epsilon\gamma \int_{S_x^c} (x - u)\phi(x - u)\eta(u) du}{(1 - \gamma)\text{erf}(\alpha/\sqrt{2}) + \gamma} \quad (113)$$

where ϕ is the standard normal probability density function. Substituting the Fourier representation of η into these fluxes allows us to write the total flux $j = j_D + j_A + j_R$ as

$$j(x) = i\epsilon \int_{-\infty}^{\infty} \hat{\eta}(\omega, t) f(\omega) e^{i\omega x} d\omega \quad (114)$$

where

$$f(\omega) = -D\omega + \frac{\sin(\alpha\omega) - \alpha\omega \cos(\alpha\omega)}{\alpha\omega^2} - \frac{\gamma}{(1 - \gamma)\text{erf}(\frac{\alpha}{\sqrt{2}}) + \gamma} \int_{T_0^c} v\phi(v) \sin(\omega v) dv. \quad (115)$$

The final integral term may be evaluated in terms of error functions of a complex argument [70], but the integral representation is more compact and easier to interpret. Substituting expression (114) for the flux into the continuum evolution equation (106) we obtain the following

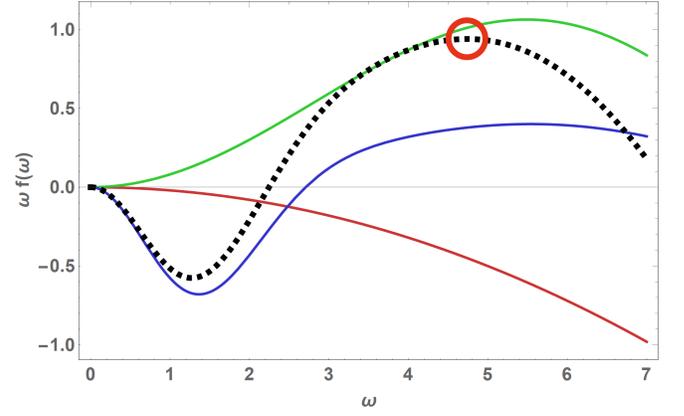


FIG. 28. Contributions to the net mode growth rate from Diffusion (red), Attraction (green) and Repulsion (blue), when $D = 0.02, \alpha = 0.5, \gamma = 1$. Black dotted curve shows overall mode growth rate $\omega f(\omega)$ with maximum circled.

ordinary differential equation for the transform of the fluctuations

$$\partial_t \hat{\eta}(\omega, t) = \epsilon \omega f(\omega) \hat{\eta}(\omega, t) \quad (116)$$

which has solution

$$\hat{\eta}(\omega, t) = \hat{\eta}(\omega, 0) \exp[\epsilon \omega f(\omega) t]. \quad (117)$$

The decoupling of Fourier modes which allowed this solution [71] is a consequence of the linearisation of the fluxes in (111), (112) and (113), so the solution is only valid for small fluctuations. However, because the locations and sizes of clusters (peaks in ρ) are decided early in the evolution of the system, the wavelength λ^* corresponding to the fastest growing wave number ω^* gives an accurate approximation to the wavelength of the final peak distribution. From (117) we see that fluctuations with wave numbers for which $f(\omega) < 0$ decay over time, whereas those for which $f(\omega) > 0$ grow, with the fastest growing wave number ω^* corresponding to the maximum of $\omega f(\omega)$

$$\omega^* = \arg \max_{\omega} \omega f(\omega). \quad (118)$$

This prediction is compared to simulations in Figure 27. In the small noise limit, $D \rightarrow 0$ we have $\lambda \rightarrow 2.29\alpha$, so we expect vowels to be separated by just over twice the maximum separation which speakers consider to be phonemically equivalent. The threshold for first formant frequency discrimination in normal speech is $\Delta F_1 \approx 50 \pm 10$ (Hz) [58], where $F_1 \in [235, 850]$ [59]. Taking ΔF_1 as a lower bound on α gives a theoretical maximum of six spontaneously formed vowel heights, consistent with observed systems [13].

In Figure 28 we have plotted the separate contributions from the diffusive, attractive and repulsive fluxes to the overall growth rate, $\omega f(\omega)$, of Fourier modes. Here we see that attractive forces cause all modes to

grow, corresponding to increasing amplitude fluctuations at all length scales as frames are pulled together producing higher density peaks. The diffusive flux reduces (smoothes out) fluctuations at all length scales, but has greatest effect on high frequency (short length scale) fluctuations because it penalizes the second derivative of the density distribution. The behaviour of the repulsive flux is more subtle: for low wave number fluctuations (large clusters) it has a negative effect: if the typical repulsion range is smaller than the size of a cluster then it will break up the cluster into smaller parts. At larger wave numbers (corresponding to smaller clusters) repulsion has a positive effect on growth: when the clusters are smaller, the dominant repulsion effect is a squeeze from near neighbours, increasing the density of all clusters.

X. DISCUSSION

We have presented a model of the evolution of vowel systems in which words behave as interacting particles diffusing in vowel space, with positions determined by the vowel sounds they contain. The interactions between particles are modelled mathematically as physical forces, but are derived based on our understanding of the language learning process. Our work builds on an already substantial collection of models of vowel system evolution [13, 14, 18–24, 27, 30], so it is important to make clear what we are adding. Because the mechanisms which control the evolution of sound inventories are not fully understood, existing models have invoked a considerable range of mechanisms in order to explain observations. Many of these mechanisms induce a combination of attractive and repulsive interactions between vowel sounds. Since the precise linguistic, social and cognitive processes which drive sound change remain unknown, it is difficult to justify anything other than a phenomenological approach. That is, one which is consistent with existing theories but does not rely on a precise scientific hypotheses regarding the underlying mechanisms. Under these conditions, simplicity is beneficial. By defining our model using the language of physics: forces and diffusion, we simplify the conceptual picture of vowel system dynamics, and allow for a more transparent and thorough analysis of how the system parameters control predictions. Moreover, these parameters may in some cases be directly measured or bounded (cloud radius, template range) or their values inferred from data (diffusion coefficient, momentum sensitivity).

The simplicity of our model definition has allowed us to combine quantal and dispersion theory into a single framework [11, 43, 67], to capture allophonic variation and elucidate the mechanism of phonemic splitting [35], to explore the effects of functional load [48], model self actuating sound changes [40–42, 64], provide a simplified picture of chain shifting in which the distinction between drag and push chains becomes redundant [17], explain empirically observed word frequency effects in

sound change [24], predict the critical level of stochasticity at which phonemes disintegrate, predict the maximum possible number of vowel heights a language can contain [12, 19], and to provide a simple picture of the process of phonemic merger. Its simplicity also allows for efficient simulation and mathematical analysis. Because it is derivable by considering the behaviour of individual speakers, it is in principle straightforward to extend the definition to interacting communities and social networks. Because the atomic constituents of the model are individual words, it is also in principle possible to use our approach to model specific cross-linguistic interactions such as borrowing, where new phonemes are created by the inclusion of foreign words into a language (e.g. from French to English [35]).

Our scientific conclusions are as follows. In its simplest form our model may be seen as a *dynamical* dispersion theory [31], where repulsive inter-phoneme forces drive vowel systems toward configurations which maximise contrast. By comparison to a large database of phonemic inventories, and by defining of our own system typology, we have shown that the model captures cross linguistic relative frequencies of different sounds to within $\approx 10\%$, with the exception of the high central vowel /i/. Moreover, for vowel systems up to cardinality six, the most common systems generated by the model are consistent with our typology and that of Crothers [52]. The over representation of the high central vowel is seen in other (non-dynamic) dispersion theories [13], and attempts have been made in the past to correct these predictions by defining sophisticated perceptual distance metrics [19] which effectively warp the shape and structure of vowel space. By their nature these metrics are difficult to rigorously derive and test. In contrast, the dynamics in our model is driven by phonemic overlap in acoustic space, which may be quantitatively measured using formants, fixing the aspect ratio of vowel space unambiguously.

Quantal Theory provides an alternative to dispersion theory for understanding the structure phonetic inventories. We have shown that the two theories combine in our simple framework, with quantal effects entering via the articulatory to acoustic map [11, 39, 43] which exhibits a pronounced non-linearity at the second subglottal resonance. We have shown that this induces a repulsive force away from the centre of vowel space, reducing the relative frequency of /i/, and stabilising the second most common vowel system (the seven vowel system of Italian and Yoruba). However, the quantal force also destabilises the mid central vowel, and we can only speculate as to the mechanisms which could counteract this effect. We note, however, that in real languages /i/ and /ə/ are positively correlated, suggesting that the presence of one may increase the stability of the other.

Moving beyond the over preponderance of /i/, examination of cross-linguistic correlations between phonemes (Table IV), and of the most common empirically observed vowel systems (Figure 12) reveals that vowels tend to oc-

cur in front-back pairs of the same height. This symmetry is present, but to a lesser extent, in our predictions. In dispersion theory, positive correlations between vowels of the same height occur because such configurations happen to maximize contrast. Our results suggest in real vowel systems there may be some additional mechanism which imposes this symmetry. We will address this point in further work.

Sporadic sound changes [72], often involving multiple vowels, are well documented [35, 64]. However, the long standing *actuation problem* remains unresolved: “Why do changes in a structural feature take place in a particular language at a given time, but not in other languages with the same feature, or in the same language at other times?” [73]. One possible resolution to the problem is *momentum based selection* [40–42, 74], according to which speakers react to features whose relative frequency has risen in the recent past, by further emphasising the use of these features. In the context of sound change, the analogue of an increase in frequency is a net direction of change with time, which may be realised in *apparent time* [35] as a difference between the behaviour of old and young speakers. Inspired by frequency based models [40–42] we have incorporated momentum based change in to our model, showing that when sensitivity to this change exceeds a critical threshold then long term sustained shifts in vowel systems can take place. We find that without such an effect natural variations in speaker’s articulatory behaviour, which can be substantial [75], and are captured by our diffusion coefficient D , are not sufficient to actuate sustained shifts over realistic time scales, whilst still preserving the integrity of phonemes (Figure 24). Our model of momentum based change generates

shifts over $O(10^1)$ lifetimes, with variations depending on the level of sensitivity, consistent with the durations of observed changes such as the Great Vowel Shift [68]. In the deterministic setting ($D = 0$), artificially activated push chains described by a system of ordinary differential equations, show that high frequency words behave as more massive particles, responding more slowly to the encroachment of a phoneme (Figure 26), provided that the self focus is positive; that is, when learning how to pronounce the vowel sound in a word, learners place more emphasis on the word itself, than on other words containing similar vowel sounds.

Our model is by no means a perfect description of vowel system structure and dynamics. It is best described as a *toy model*. However, due to its simplicity and flexibility we have been able to use it to study and understand a considerable variety of linguistic processes, providing a simple mathematical picture which we hope may be a useful tool for understanding the evolution of the sounds of languages. The model may provide a framework for future predictive modelling approaches which use large formant datasets of phonemic inventories to directly calibrate dynamical models.

ACKNOWLEDGMENTS

The authors are grateful to the Royal Society for an APEX award (2018-2020), funded by the Leverhulme trust.

-
- [1] P. Roach, *English Phonetics and Phonology* (Cambridge University Press, Cambridge, 2009).
- [2] D. Fogerty and D. Kewley-Port, “Perceptual contributions of the consonant-vowel boundary to sentence intelligibility,” *The Journal of the Acoustical Society of America* **126** (2009).
- [3] F. Chen, L. Wong, and E. Wong, “Assessing the perceptual contributions of vowels and consonants to mandarin sentence intelligibility,” *The Journal of the Acoustical Society of America* **134** (2013).
- [4] K. Stevens, *Acoustic phonetics* (MIT Press, Cambridge, MA, 2012).
- [5] <http://www.internationalphoneticassociation.org/content/ipa-chart> (2015).
- [6] G. E. Peterson and H. Barney, “Control methods used in a study of the vowels,” *Journal of the Acoustical Society of America* **24**, 175 (1952).
- [7] A. Takayuki, “Education in acoustics and speech science using vocal-tract models,” *The Journal of the Acoustical Society of America* **131**, 2444 (2011).
- [8] J. Wang, J. R. Green, A. Samal, and Y. Yunusovac, “Articulatory distinctiveness of vowels and consonants: A data-driven approach,” *Journal of Speech, Language, and Hearing Research* **56**, 1539 (2013).
- [9] J. Lee, S. Shaiman, and G. Weismer, “Relationship between tongue positions and formant frequencies in female speakers,” *Journal of the Acoustical Society of America* **139**, 426 (2016).
- [10] G. Birkhoff and S. MacLane, *A Survey of Modern Algebra* (A K Peters, Massachusetts, 1997).
- [11] K. N. Stevens and S. J. Keyser, “Quantal theory, enhancement and overlap,” *Journal of Phonetics* **38**, 10–19 (2010).
- [12] I. Maddieson, *Patterns of Sounds* (Cambridge University Press, Cambridge, 1984).
- [13] J. Liljencrants and B. Lindblom, “Numerical simulation of vowel quality systems: The role of perceptual contrast,” *Language* **48**, 839–862 (1972).
- [14] B. de Boer, “Evolution and self-organisation in vowel systems,” *Evolution of Communication* **3**, 79–103 (1999).
- [15] J. Burridge, “Spatial evolution of human dialects,” *Phys. Rev. X* **7**, 031008 (2017).
- [16] J. Milroy and L. Milroy, “Linguistic change, social network and speaker innovation,” *Journal of Linguistics* **21**, 339–384 (1985).
- [17] W. Labov, *Principles of Linguistic Change: Cognitive*

- and *Cultural Factors, Volume 3* (Wiley, Chichester, 2010).
- [18] N. Vallee, “The weight of phonetic substance in the structure of sound inventories,” *ZAS Papers in Linguistics* **28**, 145–168 (2002).
- [19] B. Lindblom, “Phonetic universals in vowel systems,” in *Experimental Phonology*, edited by J. Ohala and J. J. Jaeger (Academic Press, Orlando, 1986) pp. 13–44.
- [20] R. Caree, “Prediction of vowel systems using a deductive approach,” International Conference on Spoken Language Processing (ICSLP) **96**, 434–437 (1996).
- [21] B. de Boer, “Self-organization in vowel systems,” *Journal of Phonetics* **28**, 441–465 (2000).
- [22] J. Pierrehumbert, “Speech perception without speaker normalization: An exemplar model,” in *Talker Variability in Speech Processing*, edited by K. Johnson and J. W. Mullenix (Academic Press, 1997) pp. 145–165.
- [23] P. F. Tupper, “Exemplar dynamics and sound merger in language,” *SIAM J. Appl. Math.* **75**, 1469–1492 (2015).
- [24] S. Todda, J. B. Pierrehumbert, and J. Hay, “Word frequency effects in sound change as a consequence of perceptual asymmetries: An exemplar-based model,” *Cognition* **185**, 1 (2019).
- [25] K. J. B. Johnson, “Speech perception without speaker normalization: An exemplar model,” in *Talker Variability in Speech Processing*, edited by K. Johnson and J. W. Mullenix (Academic Press, 1997) pp. 145–165.
- [26] W. G. Moulton, “Dialect geography and the concept of phonological space,” *Word* **18**, 23–32 (1962).
- [27] J. L. Schwartz, L. J. Boe, N. Vallee, and C. Abry, “The dispersion-focalization theory of vowel systems,” *Journal of Phonetics* **25** (1997).
- [28] J. Pierrehumbert, “Exemplar dynamics: Word frequency, lenition and contrast,” in *Frequency and the Emergence of Linguistic Structure*, edited by J. L. Bybee and P. J. Hopper (John Benjamins, Amsterdam, 2001) p. 137.
- [29] R. M. Nosofsky, “Attention, similarity, and the identification-categorization relationship,” *Journal of Experimental Psychology: General* **115**, 39–57 (1986).
- [30] B. Goodman and P. F. Tupper, “Stability and fluctuations in a simple model of phonetic category change,” *SIAM J. Appl. Dyn. Syst.* **17** (2018).
- [31] B. Vaux and B. Samuels, “Explaining vowel systems: dispersion theory vs natural selection,” *The Linguistic Review* **32**, 573–599 (2015).
- [32] B. Oksendal, *Stochastic Differential Equations: An Introduction with Applications* (Springer, Berlin, 2010).
- [33] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences* (Springer, Berlin, 2009).
- [34] L. D. Landau and E. M. Lifshitz, *A course of Theoretical Physics: Statistical Physics: Volume 5* (Butterworth-Heinemann, Oxford, 1980).
- [35] W. Labov, *Principles of Linguistic Change: Internal Factors, Volume 1* (Wiley, Chichester, 1994).
- [36] R. J. Hunter, *Foundations of Colloid Science* (Oxford University Press, 2001).
- [37] A. Stradner, H. Sedgwick, F. Cardinaux, W. C. K. Poon, S. U. Egelhaaf, and P. Schurtenberger, “Statistical physics of language dynamics,” *Nature* **432** (2004).
- [38] S. Mossa, F. Sciortino, P. Tartaglia, and E. Zaccarelli, “Ground-state clusters for short-range attractive and long-range repulsive potentials,” *Langmuir* **20** (2004).
- [39] Y. Liu and Y. Xi, “Colloidal systems with a short-range attraction and long-range repulsion: Phase diagrams, structures, and dynamics,” *Current Opinion in Colloid and Interface Science* **39**, 123–136 (2019).
- [40] W. G. Mitchener, “A mathematical model of prediction-driven instability: How social structure can drive language change,” *Journal of Logic, Language and Information* **20**, 385–396 (2011).
- [41] K. Stadler, R. A. Blythe, K. Smith, and S. Kirby, “Momentum in language change,” *Language dynamics and change* **6**, 171–198 (2016).
- [42] W. G. Mitchener, “A stochastic model of language change through social structure and prediction driven instability,” *SIAM Journal of Applied Mathematics* **77**, 2272–2293–396 (2017).
- [43] M. Sonderegger, *Subglottal coupling and vowel space: An investigation in quantal theory (PhD Thesis)* (Massachusetts Institute of Technology, 2004).
- [44] K. L. Pike and Pike E. V., “Immediate constituents of mazateco symbols,” *International Journal of American Linguistics* **13** (1947).
- [45] Pierrehumbert, “Probabilistic phonology: Discrimination and robustness,” in *Probabilistic Linguistics*, edited by R. Bod, J. Hay, and S. Jannedy (MIT Press, Cambridge, MA, 2003) pp. 225–226.
- [46] Association phonetique internationale, “Expose des principes de l’association phonetique internationale,” *Le Maitre Phonetique* **19** (1900).
- [47] S. Disner, “Vowel quality: The relation between universal and language-specific factors,” *UCLA Working Papers in Phonetics* **58** (1983).
- [48] A. Wedel, A. Kaplan, and S. Jackson, “High functional load inhibits phonological contrast loss: A corpus study,” *Cognition* **128**, 179–186 (2013).
- [49] A. Martinet, “Function, structure, and sound change,” *Word* **8**, 1–32 (1952).
- [50] J. Hillenbrand, L. A. Getty, M. Clark, and K. Wheeler, “Acoustic characteristics of american english vowels,” *The Journal of the Acoustical Society of America* **97** (1995).
- [51] P. Adank, R. van Hout, and R. Smits, “An acoustic description of the vowels of northern and southern standard dutch,” *The Journal of the Acoustical Society of America* **116** (2004).
- [52] J. Crothers, “Phonetic universals in vowel systems,” in *Universals of human language. Vol 2: Phonology*, edited by J. H. Greenberg, C. A. Ferguson, and Moravcsik (Stanford University Press, Stanford, 1978) pp. 93–152.
- [53] Steven Moran and Daniel McCloy, eds., *PHOIBLE 2.0* (Max Planck Institute for the Science of Human History, Jena, 2019).
- [54] J.K. Chambers and Peter Trudgill, *Dialectology* (Cambridge University Press, 1998).
- [55] G. Grimmett and D. Stirzaker, *Probability And Random Processes* (Oxford University Press, Oxford, 2001).
- [56] K. Fukunaga and L. D. Hoststler, “The estimation of the gradient of a density function, with applications in pattern recognition,” *IEEE Transactions on Information Theory* **21** (1975).
- [57] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning* (Springer, New York, 2009).
- [58] C. Liu and D. Kewley-Port, “Vowel formant discrimination for high-fidelity speech,” *The Journal of the Acoustical Society of America* **116** (2004).
- [59] J.C. Catford, *A Practical Introduction to Phonetics* (Oxford University Press, 1988).

- [60] S. Y. Manuel, “The role of contrast in limiting vowel-to-vowel coarticulation in different languages,” *The Journal of the Acoustical Society of America* **88** (1990).
- [61] J. Gillieron, *Genealogie des Mots Qui Designent l’Abeille d’apres l’Atlas Linguistique de la France* (E. Champion, Paris, 1918).
- [62] X. Chi and M. Sonderegger, “Subglottal coupling and its influence on vowel formants,” *The Journal of the Acoustical Society of America* **122** (2007).
- [63] M. J. Sjerps and R. Smiljanic, “Compensation for vocal tract characteristics across native and non-native languages,” *Journal of Phonetics* **41** (2013).
- [64] W. Labov, S. Ash, and C. Boberg, *The Atlas of North American English: Phonetics, Phonology and Sound Change* (Mouton de Gruyter, 2006).
- [65] N. Chomsky and M. Halle, *The Sound Pattern of English* (Harper and Row, New York, 1968).
- [66] R. Jakobson and M. Halle, *Fundamentals of Language* (Mouton, The Hague, 1971).
- [67] S. Lulich, “Subglottal resonances and distinctive features,” *Journal of Phonetics* **38** (2010).
- [68] C. M. Millward, *A Biography of the English Language* (Holt Rinehart Wilson, Fort Worth, 1989).
- [69] J. B. Hay, J. B. Pierrehumber, A. J. Walker, and P. LaShell, “Tracking word frequency effects through 130 years of sound change,” *Cognition* **139** (2015).
- [70] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions* (Dover, New York, 1970).
- [71] J. Mathews and R. L. Walker, *Mathematical Methods of Physics* (Benjamin, New York, 1964).
- [72] W. Labov, *Principles of Linguistic Change: Social Factors, Volume 2* (Wiley, Chichester, 1994).
- [73] U. Weinreich, W. Labov, and M. T. Herzog, “Empirical foundations for a theory of language change,” in *Directions for historical linguistics: a symposium*, edited by W. Lehmann and Y. Malkiel (University of Texas Press, Austin Texas, 1968) pp. 96–195.
- [74] T. M. Gureckis and R. L. Goldstone, “How you named your child: Understanding the relationship between individual decision making and collective outcomes,” *Topics in Cognitive Science* **1** (2009).
- [75] L. Ellis and W. J. Hardcastle, “Categorical and gradient properties of assimilation in alveolar to velar sequences: evidence from epg and ema data,” *Journal of Phonetics* **30** (2002).