Resolving Reduplicative Opacity in Malay Nasal Spreading: Argument for Base-Reduplicant Correspondence Theory*

Jian-Leat Siah[a], Sam Zukoff[a], Feng-fan Hsieh[b]
[a]*UCLA*, [b]*National Tsing Hua University*
8/27/25

1 Introduction

The interaction between nasal spreading and reduplication in Malay stands out as one of the most significant, and most dubious, patterns in the literature on reduplication and reduplicative theory. In the debates over the correct theory of reduplication, the empirical status of putative opaque reduplication-phonology interactions (Wilbur, 1973) has long taken center stage. These interactions largely fall into one of the two categories in (1). Theories differ in exactly which types and sub-types of these patterns they predict to exist. As such, understanding the empirical landscape is crucial in deciding between theories.

(1) Basic types of opaque reduplication-phonology interactions
   a. **Overapplication:** A process applies in a reduplicated word even though the environment is not met.
   b. **Underapplication:** A process does not apply in a reduplicated word even though the environment is met.

The case of nasal spreading in Malay reduplication, as reported by Onn (1976), appears to be the sole reported example of what Kiparsky (2010) termed "recopying" overapplication, a type of overapplication where the trigger for the process is contained within the reduplicant itself. It is widely recognized that such a pattern can only be accounted for when a mechanism like Base-Reduplicant (BR) Correspondence is incorporated into the theory (McCarthy & Prince, 1995; Ahmad, 2005). Proponents of alternative theories of reduplication — which reject BR correspondence, and seek to generate some of its results instead through some version of serialism — have questioned the veracity of this data (Inkelas & Zoll, 2005; Kiparsky, 2010; McCarthy et al., 2012). To our knowledge, the pattern has not been phonetically documented, with the recent exception of Siah et al. (2025), whose study was limited in terms of sample size and data scope.

This paper reports a full-scale acoustic study with a larger pool of subjects, aimed at investigating the interaction between nasal spreading and reduplication in Malay. Our findings confirm the existence of the putative recopying pattern reported by Onn (1976), but as part of a system of free variation not previously identified. To account for the newly observed variability,

* The first author conducted the experiment, analyzed the data, and implemented the computational model. The second author developed the theoretical analysis and its connection with the theoretical literature. The third author collaborated on writing and on other aspects of the ongoing research project. Thanks also to Bruce Hayes, Martin Krämer, Megha Sundara, Kie Zuraw, members of the UCLA Phonology Seminar, the audience at AMP 2024, and the audience at the Fresno State Linguistics colloquium. All mistakes and bad ideas are of our own doing. Online supplementary materials can be found at: https://osf.io/q2hw9/?view_only=bdb428b810df43afa8d1ccb5ea1ad0c2.

we revise McCarthy & Prince (1995)'s Base-Reduplicant Correspondence Theory (BRCT) analysis, in two stages.

First, using a categorical idealization of the continuous phonetic data, we propose a traditional Optimality Theoretic [OT] (Prince & Smolensky, 1993/2004) analysis of the variation using partially-ordered constraints (Anttila, 1998). This establishes the appropriateness of the BRCT approach and the specific constraints involved. However, it fails to fully capture various aspects of the phonetic distribution, including the relative frequency of the various candidates and correspondence effects at the phonetic level.

To better capture these results, we subsequently implement the analysis from a generative phonetics perspective (e.g., Flemming, 2001; Katz, 2010; Braver, 2013, 2019; Flemming & Cho, 2017; Lefkowitz, 2017; Hayes & Schuh 2019). This analysis quantizes the phonetic space, recasts constraint violations in terms of numerical values, and assigns constraint weights and probabilities within a Maximum Entropy Harmonic Grammar [MaxEnt] (Goldwater & Johnson, 2003; Hayes & Wilson, 2008; Hayes, 2022). The conclusions which can be drawn from the generative phonetic analysis accord completely with those of the more abstract OT analysis, confirming the existence of the recopying pattern as the majority output for the interaction between nasal spreading and reduplication in Malay.

These results bolster the argument for the inclusion of BR correspondence in reduplicative theory, whether framed in traditional abstract terms or more concrete phonetic terms. This paper also demonstrates the utility of integrating traditional phonological analysis with generative phonetic modeling, enabling each to inform the other.

The remainder of the paper is structured as follows. Section 2 reviews the nasal spreading pattern in Malay and its interaction with reduplication. Section 3 and Section 4 outline the methodology of the acoustic study and present experimental results, respectively. In Section 5, we propose our revised OT analysis of the observed variation in BRCT, building on McCarthy & Prince (1995). Section 6 introduces our generative phonetic modeling implemented in MaxEnt. Section 7 concludes the paper.

2 Background

Malay (*Bahasa Melayu*) is a Western Malayo-Polynesian language in the Austronesian language family (Blust, 2013:30–32)[1]. It is primarily spoken in Malaysia, Brunei, Singapore and Indonesia. This paper focuses on the variety of Malay spoken in Peninsular Malaysia (i.e., West Malaysia).

In Malay, nasalization on vowels and glides is fully allophonic, resulting in iterative rightward nasal spreading. The details of the allophonic distribution are stated in (2). Examples are given in (3) below, including evidence both from static distributions ((3)a–d) and alternations ((3)e,f).

(2) Allophonic distribution of nasalization in Malay
   a. Nasal stops trigger iterative rightward spreading of nasalization onto vowels and glides.
   b. Spreading is blocked by supralaryngeal consonants (e.g., *k*, *s*, *r*, etc.).
   c. All other vowels/glides surface as oral.

[1] Throughout this paper, all Malay examples are presented using the International Phonetic Alphabet (IPA) wherever possible. The native consonant inventory of Malay includes /p b t d tʃ dʒ k g s h m n ɲ ŋ r j w l (ʔ)/, while the vowel system comprises six phonemes: /i u e o a ə/ (Clynes and Deterding 2011; Nomoto and Soh 2019). The following orthography-to-phoneme correspondences may be helpful: <c> = /tʃ/, <j> = /dʒ/, <ny> = /ɲ/, <ng> = /ŋ/, and <y> = /j/. The grapheme <e> can represent either /e/ or /ə/.

(3) Distribution of nasalized vowels/glides in Malay (Onn, 1976:69–70)[2]

a.	'to drink'	[mĩnõm]	
b.	'to eat'	[mãkan]	(*[mãkãn], *[makan])
c.	'to rise'	[baŋõn]	(*[bãŋõn], *[baŋon])
d.	'to be luxurious'	[mẽw̃ãh̃]	(← /mewah/)
e.	'supervision'	[pəŋãw̃ãsan]	(← /pəŋ-awas-an/)
f.	'central focus'	[pən$_{1,2}$ə̃ŋãh̃ãn]	(← /pəŋ$_1$-t$_2$əŋah-an/)

Malay has a variety of reduplication patterns (Onn, 1976:104–107, 152–182, Ahmad, 2005:137–179). We focus on cases of total reduplication without further affixation, which can be used to indicate a variety of categories, including plurals, reciprocals, repetitive action, and intensification (Onn, 1976:105). The question of interest is how nasal spreading interacts with total reduplication. The kind of root that will be most probative of this question is a root like /waŋi/ (→ [waŋĩ]) 'fragrant', which has a trigger of nasal spreading (/ŋ/) in its second syllable and a target of nasal spreading (/wa/) in its first syllable. Putting aside our expectations about nasal spreading, when total copying produces a string *waŋi-waŋi*, there are (at least) four conceivable ways of (not) applying nasalization to the various spans of potential undergoers, as listed in (4).[3] These four options fully cross nasalizing the reduplicant-initial span ((4)C,D) vs. leaving it oral ((4)A,B), and nasalizing the base-initial span ((4)B,C) vs. leaving it oral ((4)A,D). We treat the lefthand constituent as the <u>reduplicant</u>, though nothing substantive changes if we assume the opposite.

(4) Potential outputs

		BASE (righthand member) Oral	Nasal
REDUPLICANT (lefthand member)	**Nasal**	D. [<u>w̃ãŋĩ</u>-waŋĩ]	C. [<u>w̃ãŋĩ</u>-w̃ãŋĩ]
	Oral	A. [<u>waŋĩ</u>-waŋĩ]	B. [<u>waŋĩ</u>-w̃ãŋĩ]

Output A [<u>waŋĩ</u>-waŋĩ] keeps both the base-initial and reduplicant-initial spans oral. This constitutes UNDERAPPLICATION of nasal spreading, because there is no spread across juncture, even though the context for further spreading is met. Output B [<u>waŋĩ</u>-w̃ãŋĩ] nasalizes the base-initial span but keeps the reduplicant-initial span oral. This constitutes NORMAL APPLICATION of nasal spreading, as it fully obeys the typical allophonic distribution seen outside of reduplication. Output C [<u>w̃ãŋĩ</u>-w̃ãŋĩ] nasalizes both the base-initial and reduplicant-initial spans. This constitutes OVERAPPLICATION of nasal spreading, because the reduplicant-initial span is nasal even though there is no local trigger. This is a special kind of overapplication, termed "recopying" by Kiparsky (2010:3), for reasons which we will detail below. The special status of this potential output is the reason why this pattern holds such importance. Output D [<u>w̃ãŋĩ</u>-waŋĩ] keeps the base-initial span

[2] We assume that glottal consonants /h,ʔ/ are undergoers, but an analysis which treats them as transparent segments is also feasible. When we refer to "glides", we mean to include /h, ʔ/.

[3] The unusual order of the labels A–D prefigures the presentation of our analysis results in Section 5.

oral while nasalizing the reduplicant-initial span. This is a pathological output, since nasalization appears in only the wrong place; that is, there is nasalization in a position where there is no local trigger of nasalization, and orality in a position where there is a local trigger.

According to Onn (1976:180), the actual output in Malay is Option C [w̃ãŋĩ-w̃ãŋĩ], the recopying overapplication output. Further examples are given in (5):

(5) Nasalization in Malay reduplication (Onn, 1976:180, Ahmad, 2005:157)

		Root in isolation	Reduplicated form
a.	‘fragrant/(intensified)’	[waŋĩ]	[w̃ãŋĩ-w̃ãŋĩ]
b.	‘germ/germs’	[hamə̃]	[h̃ãmə̃-h̃ãmə̃]
c.	‘reverie/ambition’	[aŋãn]	[ãŋãn-ãŋãn]
d.	‘wind/unconfirmed news’	[aŋẽn]	[ãŋẽn-ãŋẽn]
e.	‘termites’	[anãȷ̃]	[ãnãȷ̃-ãnãȷ̃]
f.	‘to look down upon’	[hinə̃]	[h̃ĩnə̃-h̃ĩnə̃]
g.	‘purple’	[uŋũ]	[ũŋũ-ʔ̃ũŋũ]
h.	‘henna’	[inãȷ̃]	[ĩnãȷ̃-ĩnãȷ̃]

The reason this is so significant is that process ordering theories *cannot* derive this sort of recopying pattern. McCarthy & Prince (1995:43–46) show that no ordering of nasalization and copying can derive recopying overapplication. If copying happens *before* nasalization, as shown in the derivation in (6), the normal allophonic distribution will apply to the full word-form, yielding NORMAL APPLICATION (Output (4)B above). Persistent nasalization, where nasalization happens both before and after copying, would also derive normal application. Alternatively, if copying happens *after* nasalization, as shown in the derivation in (7), the base-initial span will fail to nasalize despite surfacing with a nasal trigger to its left. This is an opaque UNDERAPPLICATION interaction (Output (4)A above).

(6) Copy > Nasalization = NORMAL

Input	/ RED-waŋi /
Rule 1: Copy	**waŋi**-waŋi
Rule 2: Nasalization	waŋ**ĩ-w̃ãŋĩ**
Output:	[waŋĩ-w̃ãŋĩ]

(7) Nasalization > Copy = UNDER

Input	/ RED-waŋi /
Rule 1: Nasalization	RED-waŋ**ĩ**
Rule 2: Copy	**waŋĩ**-waŋĩ
Output:	[waŋĩ-waŋĩ]

While both of these types of interactions are attested in the typology (arguably, at least, in the case of underapplication), these do not accord with the reported output for this pattern in Malay, which is RECOPYING OVERAPPLICATION. The moniker “recopying” stems from the idea that it could be derived by a further step of “copying” in an order like (6), where the nasality of the base is somehow transferred back onto the reduplicant without any concomitant segmental copying, as shown in (8).

(8) Recopying with process ordering

Input	/ RED-waŋi /
Rule 1: Copy	**waŋi**-waŋi
Rule 2: Nasalization	waŋ**ĩ-w̃ãŋĩ**
Rule 3: “Recopy” the nasalization	**w̃ã**ŋĩ-w̃ãŋĩ
Output:	[w̃ãŋĩ-w̃ãŋĩ]

Such an operation is not recognized in any process-ordering-based serial/derivational theory, nor is it clear what would motivate the addition of such a process beyond capturing this datapoint. What is needed is a representational device that creates a persistent link between base and reduplicant which can promote identity above and beyond surface-apparent locally-motivated process application. Among serial rule-based theories, this sort of link might be instantiated by the looped representations used by Raimy (2000:16–18, 2011:2398–2399) or the linked representations used by Frampton (2009). But among constraint-based theories, the only such representational device that has been proposed in the literature is Base-Reduplicant [BR] Correspondence (McCarthy & Prince, 1995). Base-Reduplicant Correspondence Theory [BRCT], typically couched in a parallel OT framework[4], introduces a relation between corresponding output segments in the base and the reduplicant, over which faithfulness constraints act to promote identity, subject to ranking. McCarthy & Prince (1995) demonstrate that a BR-correspondence analysis can capture the Malay recopying pattern.

Without such a mechanism, serial/derivational theories that rely on process ordering thus undergenerate the RECOPYING OVERAPPLICATION pattern, if it truly exists, whereas BRCT properly generates it. Proponents of such theories have thus questioned whether this data has been accurately reported. If the pattern is actually erroneous, then the argument is reversed: BRCT overgenerates, while alternative theories properly fail to predict recopying.

To our knowledge, prior to Siah et al. (2025)’s pilot study, there was no published phonetic data addressing this question.[5] We conducted an acoustic experiment to determine which output or outputs are actually attested in Malay.

3 Experiment

3.1 Participants

Thirty native speakers of Malay (22 females, 8 males) from Peninsular Malaysia participated in an in-person production study. They were recruited via Facebook and WhatsApp, where a digital flyer was circulated. At the start of the study, each participant completed a brief demographic questionnaire in Malay, reporting their age as well as their language and dialect background. The mean age of the participants was 33 years (SD = 9). Twelve participants reported speaking another Malay dialect in addition to Standard Malay — most commonly northern or northeastern varieties,

[4] Though see Yang (2023) for a proposal adding BR-correspondence to Harmonic Serialism with Serial Template Satisfaction (McCarthy et al., 2012).

[5] Kiparsky (2010) reports the results of small-scale impressionistic fieldwork where he does find evidence of recopying, but he does not include any phonetic evidence.

such as Kelantan (6), Kedah (3), or Perak Malay (2). Slightly more than two-thirds of the participants identified English as their second language, and one participant reported speaking Javanese. None of the participants had any speech, hearing, or language disorders.

3.2 Stimuli

The target Malay words used in the present study are given in (9), along with their English glosses.

(9) Target roots

MEDIAL C	ROOT-INITIAL SPAN					
	wa-		***ha-***		***a-***	
Nasal	/waŋi/	'fragrant'	/hama/	'pest'	/aŋan/	'dream'
Glide	/wajaŋ/	'movie'	/hajwan/	'animal'	/ajam/	'chicken'
	/waham/	'delusion'	/hajaŋ/	'wavering'	/awaŋ/	'atmosphere'
Blocker	/waduŋ/	'pickaxe'	/habuŋ/	'basket'	/abaŋ/	'brother'
	/waran/	'warrant'	/harum/	'fragrant'	/aram/	'twilight'
	/waruŋ/	'shop'	/haruŋ/	'navigable sea'	/ariŋ/	'false daisy'

All target words were disyllabic[6] and featured a nasal consonant in the second syllable (with no oral consonant to its right). The first syllables consisted of one of three different strings that undergo nasal spreading (*wa-*, *ha-*, or *a-*), with six words of each type. The medial consonant could be the nasal trigger itself (e.g., [wa**ŋ**i] 'fragrant'), an undergoer (e.g., [a**j**am] 'chicken'), or a blocker (e.g., [a**b**aŋ] 'brother'), as indicated in the first column. Each target word was embedded in the carrier sentence in (10). The carrier sentence was carefully constructed to avoid any nasal segment, minimizing potential extraneous nasalization due to context outside of the target word.

(10) *Sila tulis ____ satu kali lagi.*
"Please write ____ one time again."

3.3 Procedure

Participants produced each target word once in its unreduplicated form and three times in its reduplicated form. All trials were presented in randomized order using PowerPoint slides. This design yielded a total of 1,620 reduplicated tokens (30 participants × 18 target words × 3 repetitions) for further analysis. Data collection was conducted in a quiet room using a Shure SM10A head-mounted microphone connected to a Focusrite Scarlett USB audio interface. Recordings were made using *Audacity* at a sampling rate of 44,100 Hz. Each recording session lasted approximately 15 minutes per participant, and each was compensated $4.50 for their time.

All tokens were imported into Praat (Boersma & Weenink, 2024) for segmentation and annotation. The first vowel of each target word was segmented and annotated using Praat TextGrids. Vowel boundaries were identified based on the onset and offset of a clearly defined

[6] Native Malay words have a syllable structure of (C)V(C), with optional onset and coda consonants (Clynes and Deterding 2011; Nomoto and Soh 2019). Complex onset and coda consonants only occur in loanwords. More than 90% of roots in the native lexicon are disyllabic (Adelaar 1992).

second formant (F2) in the spectrogram, along with an abrupt change in amplitude in the waveform. Three vowel contexts were distinguished, as outlined in (11): a_U represents the vowel in the unreduplicated context (i.e., the root word), while a_R and a_B refer to the equivalent vowel in the reduplicant and the base, respectively, in the reduplicated word. Figure 1. The waveform (top) and spectrogram (middle) of the unreduplicated form [aram] (left) and the reduplicated form [aram-aram] (right). The letters "U", "R", and "B" in the last tier indicate unreduplicated, reduplicant, and base, respectively. provides an example of an output after segmentation and annotation, using the unreduplicated and reduplicated form of [aram] 'twilight' as an illustration.

(11) Contexts

i. a_U : In unreduplicated context /$\boxed{a_U}$ram/

ii. a_R : In reduplicated context in the first constituent (the reduplicant) /$\boxed{a_R}$ram-$a_B$ram/

iii. a_B : In reduplicated context in the second constituent (the base) /$a_R$ram-$\boxed{a_B}$ram/

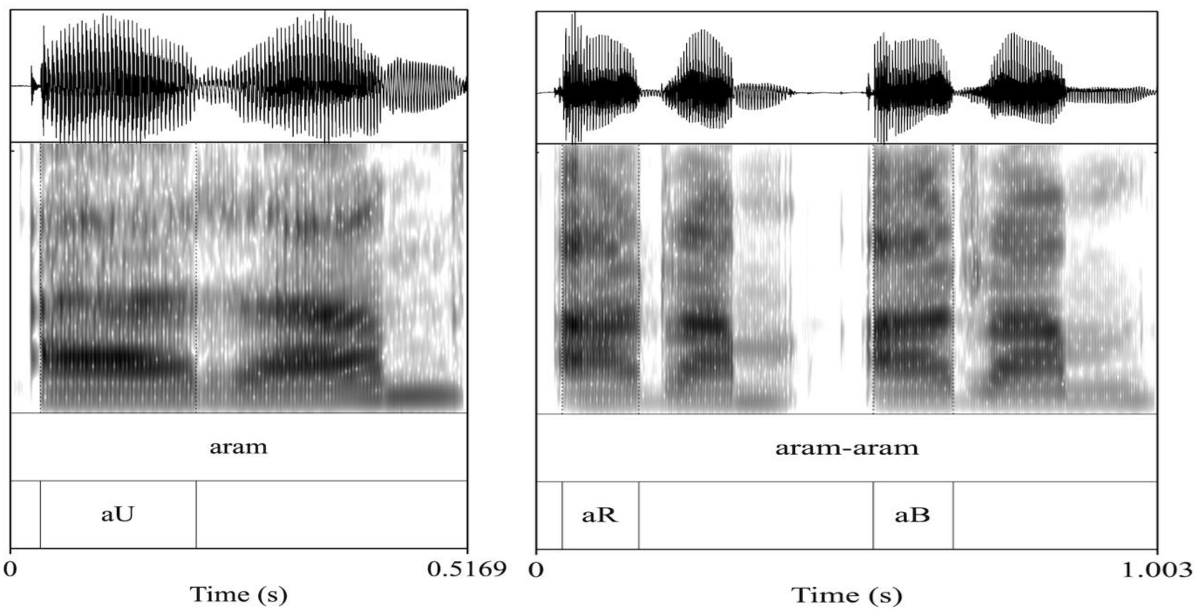

Figure 1. The waveform (top) and spectrogram (middle) of the unreduplicated form [aram] (left) and the reduplicated form [aram-aram] (right). The letters "U", "R", and "B" in the last tier indicate unreduplicated, reduplicant, and base, respectively.

3.4 Measurement

Nasality on the target vowels was measured using A1−P0, which quantifies the amplitude difference between the harmonic of the first formant (A1) and the low-frequency nasal peak (P0) (Chen, 1997; Styler, 2017). A1−P0 is the most commonly used acoustic measure of nasality in the literature on vowel nasality (Scarborough, 2013; Zellou & Tamminga, 2014; Garellek et al., 2016; Cho et al., 2017; Zellou, 2017), and is inversely correlated with the degree of nasalization, that is, larger A1−P0 values indicate less nasalization and vice versa. However, the robustness of this measure diminishes in vowels with low first formant frequency (e.g., high vowels) because both A1 and P0 might be associated with the same harmonic. To minimize the risk of misidentification of oral and nasal peaks, the present study exclusively examined target words whose first vowel was [a], i.e., a low vowel with high first formant frequency.[7]

A1−P0 values were automatically extracted at the temporal midpoint of the [a] vowel in the first syllable of the target words/constituents, using a Praat script (Styler, 2017). Tokens flagged as errors in the output were excluded from further analysis (81 out of 1,620 tokens; 5% exclusion rate). For each participant, nasalization was normalized by subtracting the mean A1−P0 value in the unreduplicated context across all forms ($\bar{a}_U$) from the A1−P0 value in the reduplicated context (a_R and a_B), and then dividing the result by the standard deviation of A1−P0 value in the unreduplicated context (σ_U). This procedure is illustrated in

Figure 2. Illustration of the procedure for calculating normalized A1−P0 values. below. The normalized measure estimates how oral or nasal the vowels in the reduplicated context are relative to their unreduplicated baseline, which is assumed to be oral due to the absence of a preceding nasal trigger.

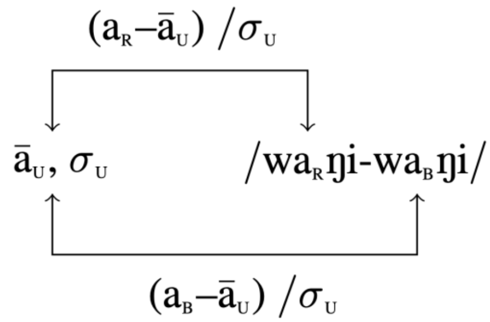

Figure 2. Illustration of the procedure for calculating normalized A1−P0 values.

4 Results

Figure 3. Scatter plot of aggregated results using normalized A1−P0 values. The x-axis and y-axis indicate the nasality of the base (a_R) and the reduplicant (a_B), respectively. Smaller A1−P0 values indicate greater nasality with respect to the oral baseline (i.e., $\bar{a}_U$). Note that both axes have been reversed, such that greater nasality appears higher and further to the right in the plot. visualizes the aggregated results across speakers as a scatter plot. Upper quadrants (quadrants I and II) indicate nasalization of the reduplicant-initial span; right-hand quadrants (quadrants I and IV) represent nasalization of the base-initial spans. We can now see that, in reality, three out of the four quadrants (corresponding to our four potential outputs from (4) above) are substantially populated, while the fourth (quadrant II) is sparsely populated. Of greatest interest to the present investigation is the robust attestation of points in quadrant I (548 tokens; 35.6%). These datapoints represent productions where *both* the reduplicant-initial span and the base-initial span have been nasalized. This is precisely our recopying overapplication output ((4)C) [w̃ãŋĩ-w̃ãŋĩ]. This comports with Onn (1976)'s original claim. However, we also see substantial attestation of the underapplication output in quadrant III (475 tokens; 30.9%) and the normal application output in quadrant IV (384 tokens; 25%). Interestingly, the output which is the most sparsely attested is quadrant II (132 tokens; 8.6%), which we identified earlier as pathological as it nasalizes only the wrong initial

[7] This is not to say that we believe the pattern to be limited to roots with low vowels. See the examples in ((5)f–h) for roots with high vowels in the relevant position.

span. This affirms that we are dealing with a case of output variation. Moreover, two additional gradient tendencies in the experimental results merit mention: (i) data points clustered near the origin, indicating that extreme deviations from the oral baseline were relatively rare; and (ii) a correspondence effect at the phonetic level, whereby the orality or nasality of the base and reduplicant tended to match (Pearson's correlation coefficient, $r = 0.50$).

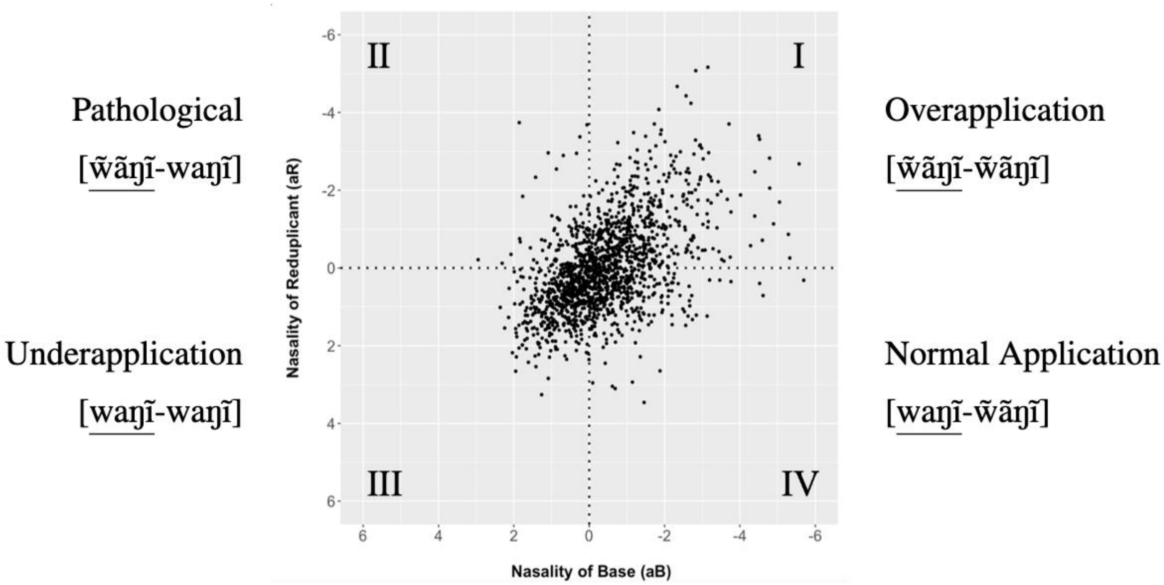

Figure 3. Scatter plot of aggregated results using normalized A1−P0 values. The x-axis and y-axis indicate the nasality of the base (a_R) and the reduplicant (a_B), respectively. Smaller A1−P0 values indicate greater nasality with respect to the oral baseline (i.e., $\bar{a}_U$). Note that both axes have been reversed, such that greater nasality appears higher and further to the right in the plot.

The variation observed in Figure 3. Scatter plot of aggregated results using normalized A1−P0 values. The x-axis and y-axis indicate the nasality of the base (a_R) and the reduplicant (a_B), respectively. Smaller A1−P0 values indicate greater nasality with respect to the oral baseline (i.e., $\bar{a}_U$). Note that both axes have been reversed, such that greater nasality appears higher and further to the right in the plot. appears to be truly free variation, as further illustrated in Figure 4 and Figure 5 below. Figure 4 breaks down the same dataset by speaker, showing the distribution of tokens across the four quadrants in proportions. Except for speaker F22, all participants produced a range of outputs that include over-, under-, normal, and pathological application of nasal spreading. Among these, pathological application was the least frequent across each member of the subject pool.

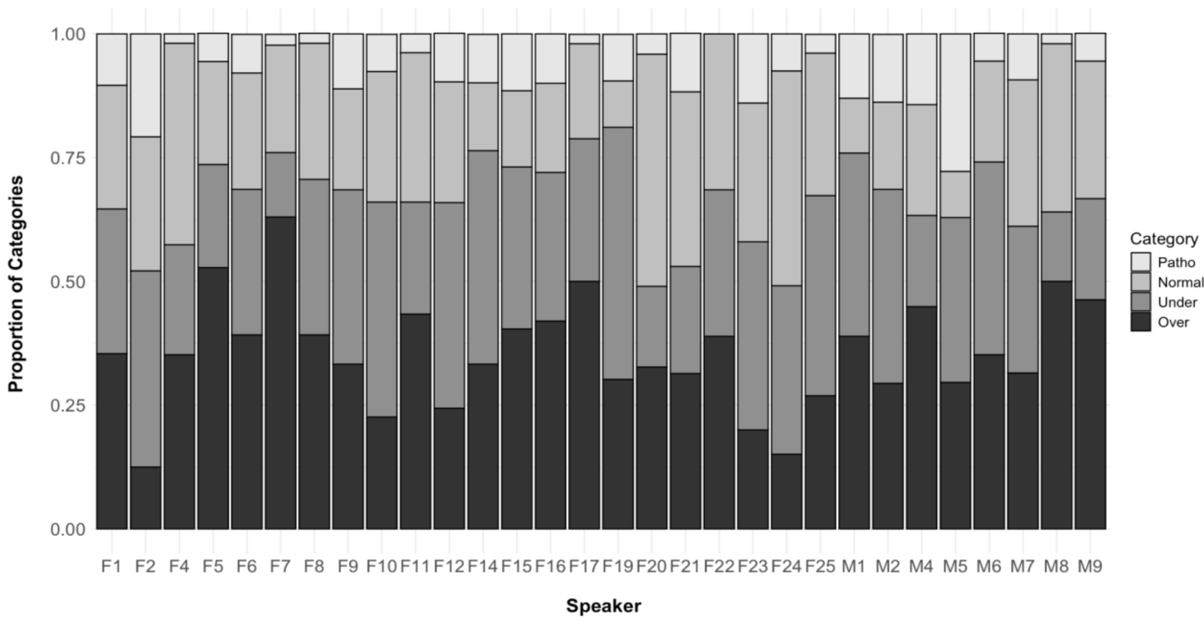

Figure 4. Proportion of output categories by speaker

A similar pattern emerges when the data is broken down by word, as shown in Figure 5. Here too, all reduplicated words (except for /aŋan-aŋan/ 'reverie') exhibit variable outputs that span all four patterns. Despite this broad variability, several noteworthy trends can be observed. First, initial glottal fricative [h] tends to favor the overapplication output. This is especially evident in words where the medial consonant is an undergoer of nasal spreading (i.e., /hajwan/ and /hajaŋ/). The affinity between glottal articulation and nasality (both synchronically and diachronically) has been dubbed *rhinoglottophilia* by Matisoff (1975), due to their similar acoustic effects on neighboring vowels, such as an increased formant bandwidth, the presence of anti-resonances in the spectrum, and an overall reduction in vowel amplitude (see Ohala 1975).

Second, overapplication is also more likely when the nasal trigger immediately follows the target vowel [a] (i.e., /aŋan/, /hama/, and /waŋi/). This increase in the overapplication output is likely attributable to regressive nasalization due to coarticulation. However, regressive coarticulatory nasalization alone cannot account for overapplication in words such as /aram/ and /abaŋ/, where an intervening blocker would prevent any local nasal coarticulation. Nasality in these reduplicated words must stem from other sources, which we argue is due to the recopying of nasal spreading during reduplication.

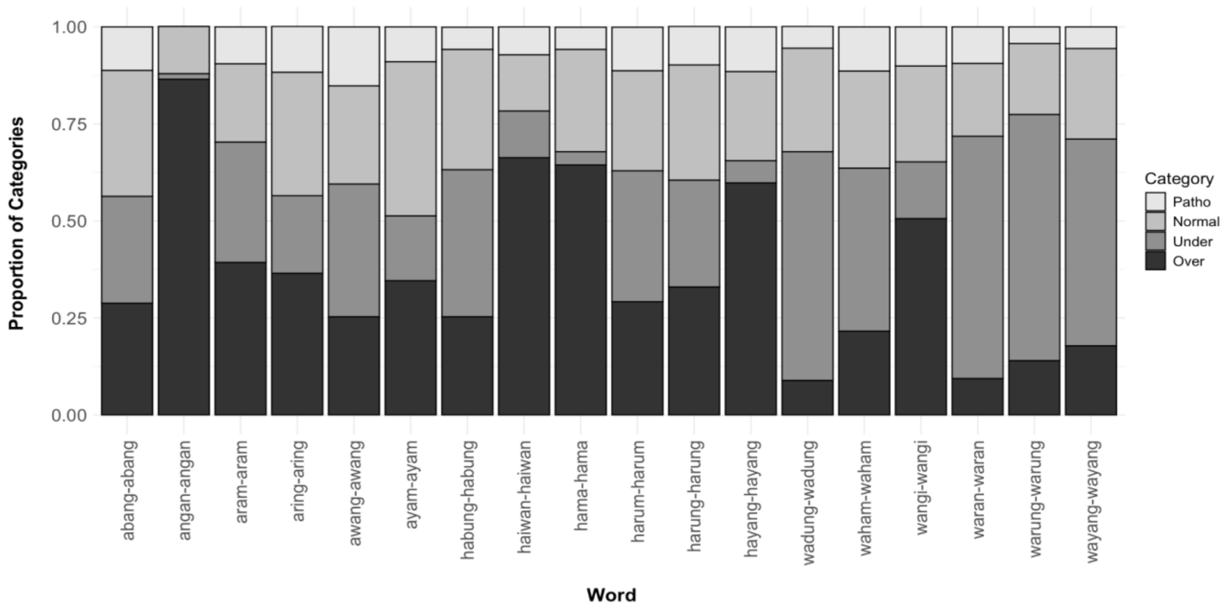

Figure 5. Proportion of output categories by word

4.1 Interim Summary

Our experimental results reveal that the distribution of nasality in reduplication in Malay is more complex than previously reported. Specifically, three different types of outputs — normal application, underapplication, and overapplication — are all substantially attested. Notably, these are the three types of outputs predicted by the factorial typology of constraint interactions in Base-Reduplicant Correspondence Theory [BRCT] (McCarthy & Prince, 1995). This suggests that a BRCT analysis, supplemented with a mechanism for deriving variation, is appropriate for the data. Such an analysis, revising the one proposed by McCarthy & Prince (1995), will be presented in Section 5.

However, the existence of a fourth possible output (quadrant II = Output ((4)D) [w̃ãŋĩ-waŋĩ]), albeit infrequent, challenges the sufficiency of a simple BRCT analysis set within classical Optimality Theory (Prince & Smolensky, 1993/2004). This is because, under the BRCT analysis to be proposed, this output is harmonically bounded by two of the other three attested outputs. To account for this additional output, and to capture the gradient patterns observed in the acoustic results, we develop a Maximum Entropy Harmonic Grammar (MaxEnt) model in Section 6, building on the constraints established in the standard BRCT analysis. Since our goal is to capture the fine-grained details of the phonetic distribution, our MaxEnt model will make use of scalar evaluation of phonetically-defined constraint violations, following the generative phonetic approach of Flemming (2001), Katz (2010), Braver (2013, 2019), Flemming & Cho (2017), Lefkowitz (2017), and Hayes & Schuh (2019).

5 Theoretical Analysis

In this section, we model the experimental results using Optimality Theory ([OT]; Prince & Smolensky, 1993/2004). As discussed in Section 2, the recopying overapplication output is only derivable in a theory with BR correspondence or some similar mechanism. For this reason, we model reduplication and its interaction with phonology with BRCT (McCarthy & Prince, 1995 [M&P]). To account for the variable outputs, we will propose a grammar with partially-ordered constraints (Anttila, 1998).

The structure of this analysis is as follows. First we will present M&P's OT analysis of the basic allophonic spreading pattern. Then we will give our adaptation of M&P's BRCT analysis of the reduplication pattern, supplemented with Output-Output/Base-Derivative correspondence (Benua, 1997) and partially-ordered constraints to account for the newly observed variation.

5.1 OT Analysis of Allophonic Nasal Spreading

M&P (p. 42) derive the allophonic distribution using the constraints in (12), ranked as in (13). This analysis derives the correct result for the static distribution of nasality (14), as well as cases involving iterative spreading and alternations (15).

(12) Constraints for the allophonic distribution

a. ***NV (*[+nas][−nas,−cons]):** Assign a violation * for each non-nasal vowel or glide which immediately follows a nasal(ized) segment.

b. ***Ṽ (*[+nas,−cons]):** Assign a violation * for each nasalized vowel or glide.

c. **IDENT[±nas]-IO:** Assign a violation * for each segment whose output value of [±nasal] does not match its input value.

(13) **Ranking:** *NV ≫ *Ṽ ≫ IDENT[±nas]-IO

The ranking of *NV above *Ṽ yields nasal vowels/glides after nasals: ((14)b) ≻ ((14)a,c). The ranking of *Ṽ above IDENT[±nas]-IO in turn ensures that all other vowels/glides are oral: ((14)b) ≻ ((14)d).[8]

(14) Static distribution of nasalization in Malay (with maximally unfaithful input)

/makãn/	*NV	*Ṽ	IDENT[±nas]-IO
a. makan	*!		*
b. ☞ mãkan		*	**
c. makãn	*!	*	
d. mãkãn		**!	*

The same ranking causes nasality to spread iteratively when there is an extended vowel/glide span (15). In /pəŋ-awas-an/, the prefixal /ŋ/ induces a *NV violation when concatenated with the root (15a). *NV ≫ *Ṽ prefers nasalizing the root-initial /a/: (15d) ≻ (15a). But to fully alleviate the *NV violation, the whole span must be nasalized, ruling out candidates like (15b,c) which nasalize only part of the span. Spreading terminates at the [+consonantal] segment /s/, because the string [ãs] does not violate *NV (*[+nas][−nas,−cons]). This means that the /a/ following the /s/ doesn't nasalize: (15d) ≻ (15e). Likewise, since the /ə/ of the prefix does not follow a nasal consonant, it does not nasalize either: (15d) ≻ (15f). This is worth making explicit, because this is the same kind of position where we find nasalization in the recopying overapplication output in reduplication.

[8] We need a high-ranked constraint against denasalization of nasal stops (e.g., IDENT[±nas]/[+cons]-IO) to prevent trigger effacement, i.e., /makan/ → *[bakan]. We omit this to simplify the analysis.

(15) Iterative nasal spreading in Malay

/pəŋ-awas-an/	*NV	*Ṽ	IDENT[±nas]-IO
a. pəŋ-awas-an	*!		
b. pəŋ-ãwas-an	*!	*	*
c. pəŋ-ãw̃as-an	*!	**	**
d. ☞ pəŋ-ãw̃ãs-an		***	***
e. pəŋ-ãw̃ãs-ãn		****!	****
f. pə̃ŋ-ãw̃ãs-an		****!	****

5.2 BRCT Analysis of Nasalization in Reduplication: Preliminaries

5.2.1 Assumptions/claims about correspondence

To generate recopying overapplication in reduplication, we need to adopt Base-Reduplicant Correspondence Theory [BRCT]. BRCT posits that a correspondence relation (BR) holds between the output base and the output reduplicant. Faithfulness constraints act over this relation to encourage similarity between the base and the reduplicant.

(16) Base-Reduplicant Correspondence Theory (McCarthy & Prince 1995:4)

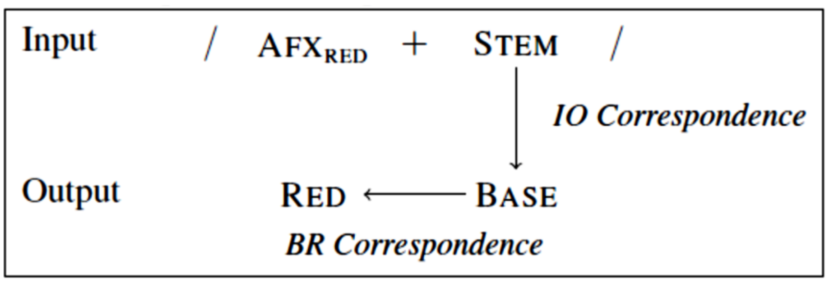

To capture the underapplication and normal application outputs, we adopt an additional component of Correspondence Theory: Output-Output / Base-Derivative (BD) correspondence (Benua, 1995, 1997; Burzio, 1996; Kenstowicz, 1996; Kager, 1999, *et seq*.). In this approach, a morphologically complex derivative corresponds to, and may be faithful to, its morphological base, i.e., the output of its stem in isolation. This is a parallelist alternative for capturing cyclic effects. Insofar as it applies to the current pattern, we can therefore understand underapplication and normal application in this case as cyclic effects.

(17) Base-Derivative Correspondence (cf. Benua, 1997:7)

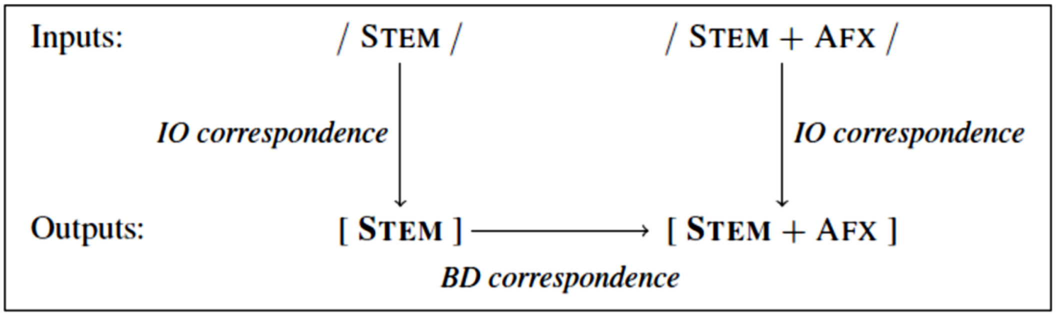

One subsidiary (and, as far as we know, novel) claim that we make is that the BD correspondence relation in Malay reduplication (indicated with red superscript numerical correspondence indices in (18) below) must hold between the morphological base (the unreduplicated output root) and *both* (i) the reduplicative base, and (ii) the reduplicant. The motivation for this claim will be explained below.

(18) Correspondence relations in Malay reduplication (nasalization suppressed)

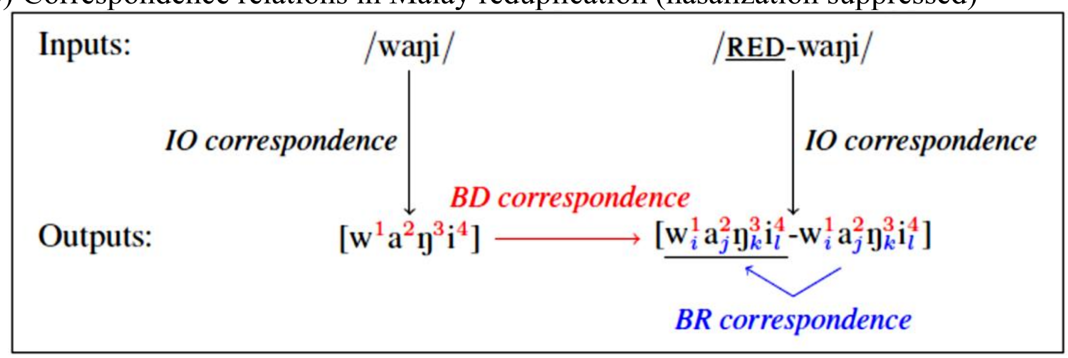

5.2.2 The candidates

To proceed from the continuous acoustic distribution outlined in Section 4 to an OT analysis with discrete candidates, we will collapse the data from Figure 3 above, repeated here as Figure 6, back into four candidates corresponding to the four quadrants. These four candidates match the four outputs contemplated earlier in (4), repeated in (19) below. Figure 3/Figure 6 demonstrated that three of these four outputs in free variation were substantially attested, while one output — Output ((4)D) [<u>w̃ãŋĩ</u>-waŋĩ] of quadrant II — was sparsely attested. As we will see shortly, the current BRCT analysis is not capable of selecting this output as the winning candidate due to harmonic bounding. For this reason, for the moment, we make the simplifying assumption that this candidate indeed ought not be selected as a winner from the perspective of the BRCT analysis. We will return to the status of this candidate in greater detail in Section 6.

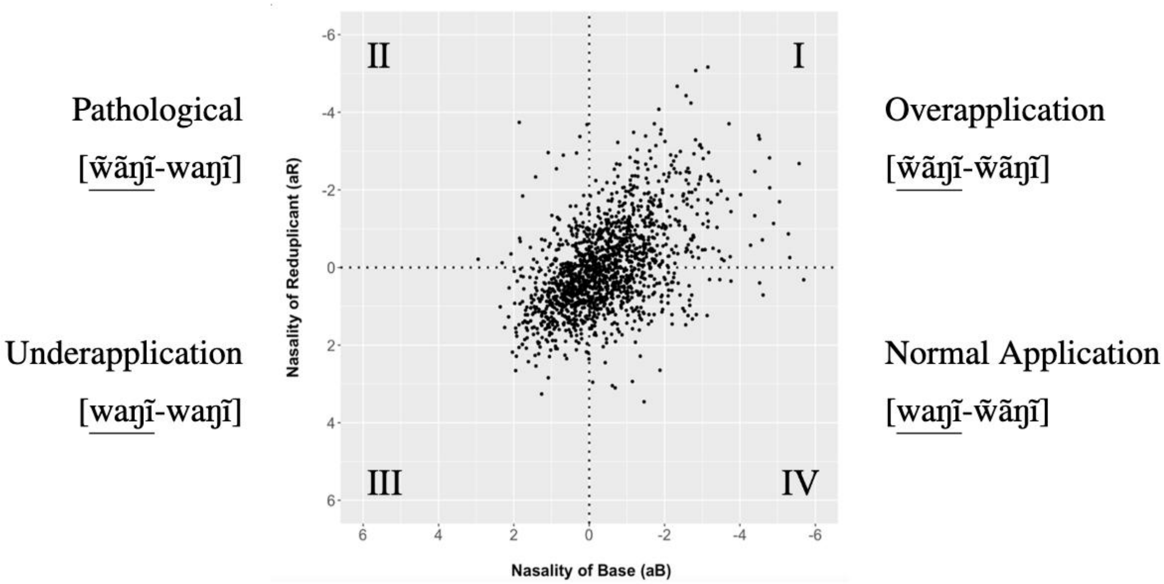

Figure 6. Scatter plot of aggregated results using normalized A1−P0 values (repeated from Figure 3).

(19) Outputs in free variation (see (4) above)

a. **Output A = Quadrant III:** [waŋĩ-waŋĩ] UNDERAPPLICATION
Nasalize just the /i/'s following the [ŋ]'s

b. **Output B = Quadrant IV:** [waŋĩ-w̃ãŋĩ] NORMAL APPLICATION
Do iterative nasal spreading like normal

c. **Output C = Quadrant I:** [w̃ãŋĩ-w̃ãŋĩ] OVERAPPLICATION ("recopying")
Do iterative nasal spreading like normal, nasalize the reduplicant-initial span to match base

d. ***Output D = Quadrant II:** [w̃ãŋĩ-waŋĩ] PATHOLOGICAL APPLICATION
Nasalize only the wrong initial span

5.2.3 The constraints and rankings

To capture the variation between Outputs A, B, and C, the only new constraints that we need to add to M&P's basic analysis of allophonic nasal spreading are the IDENT[±nasal] constraints defined over these two new correspondence relations ((20), (21)). M&P employ IDENT[±nasal]-BR in their analysis, which is sufficient to derive recopying overapplication. We add IDENT[±nasal]-BD, as it is necessary to account for the newly identified underapplication and normal application outputs.

(20) **IDENT[±nas]-BR**: Assign a violation * for each segment in the reduplicant whose value of [±nasal] does not match its correspondent in its reduplicative base.

(21) **IDENT[±nas]-BD:** Assign a violation * for each segment in the derivative whose value of [±nasal] does not match its correspondent in its morphological base.

The observed variation is derived via a variable ranking (cf. Anttila, 1998) between IDENT[±nasal]-BD, IDENT[±nasal]-BR, and *NV (all ranked above *Ṽ and IDENT[±nas]-IO), as shown in (22).

The factorial resolution of the three mutually unranked top constraints leads to our three attested outcomes.

(22) **Ranking**

IDENT[±nasal]-BD ~ IDENT[±nasal]-BR ~ *NV

*Ṽ

IDENT[±nas]-IO

5.3 BRCT Analysis of Nasalization in Reduplication: Analysis of Variation

5.3.1 Violation profiles

The tableau in (23) shows the violation profile of the relevant candidate outputs:

(23) Variable outputs of nasal spreading in Malay reduplication

MORPH BASE: [waŋĩ]						
INPUT: /RED, waŋi/			IDENT[±nas]-BD	IDENT[±nas]-BR	*NV	*Ṽ
a.	☞ waŋĩ-waŋĩ	UNDER			* (ĩ-w)	**
b.	☞ waŋĩ-w̃ãŋĩ	NORMAL	** (w̃,ã)	** (w,a)		****
c.	☞ w̃ãŋĩ-w̃ãŋĩ	OVER	**** (w̃,ã,w̃,ã)			******
d.	w̃ãŋĩ-waŋĩ	PATHO	** (w̃,ã)	** (w,a)	* (ĩ-w)	****

Candidates ((23)a–c) are our three frequently attested outputs in the experimental results. These are the three candidates which will win under some set of resolutions of the variable ranking in (22). The pathological candidate ((23)d), where the reduplicant-initial span is nasalized (with no local trigger) but the base-initial span is not (despite having a local trigger), is harmonically bounded by ((23)a) and ((23)b). This is not a completely desired result, as this candidate was attested in the experimental results. However, its substantially low frequency of attestation suggests that it is reasonable to treat it as a non-winner at this level of abstraction. As mentioned, we will treat this as a non-winner for the purposes of this BRCT analysis and return to this issue in Section 6.

To streamline the presentation, we have omitted other losing candidates which are phonotactically non-viable, namely, candidates that fail to nasalize either of the /i/'s, or that nasalize partial spans. The reader can verify that such candidates are substantially harmonically bounded by the winning candidates.

5.3.2 Underapplication

Underapplication is derived by the ranking condition in (24). The underapplication output (25a) maintains the isolation form of the root (satisfying IDENT[±nas]-BD) in both base and reduplicant at the expense of a nasal-oral sequence at the juncture (violating *NV). Other outputs nasalize segments which were oral in the morphological base, fatally violating IDENT[±nas]-BD.[9]

(24) **Underapplication ranking:** IDENT[±nas]-BD ≫ *NV[10]

(25) Variable realization: UNDERAPPLICATION[11]

MORPH BASE: [waŋĩ]			IDENT[±nas]-BD	IDENT[±nas]-BR	*NV	*Ṽ
INPUT: /RED, waŋi/						
a. ☞	waŋĩ-waŋĩ	UNDER			* (ĩ-w)	**
b.	waŋĩ-w̃ãŋĩ	NORMAL	*!* (w̃,ã)	*!* (w,a)		****
c.	w̃ãŋĩ-w̃ãŋĩ	OVER	*!*** (w̃,ã,w̃,ã)			******
d.	w̃ãŋĩ-waŋĩ	PATHO	*!* (w̃,ã)	*!* (w,a)	* (ĩ-w)	****

5.3.3 Normal application

Normal application is derived by the ranking conditions in (26). Normal application ((27)b) occurs when it is least important to maintain identity between base and reduplicant (i.e., it is tolerable to violate IDENT[±nasal]-BR). *NV must rank highest, so as to force spreading across the juncture, ruling out the underapplication candidate ((27)a). In order to distinguish between the normal application candidate ((27)b) and the overapplication candidate ((27)c), it needs to be the case that not just the reduplicative base portion of the output, but also the reduplicant stands in BD correspondence with the morphological base. We find this claim especially reasonable in this case given that we are dealing with total reduplication, such that the two constituents are indistinguishable from each other and from the morphological base. Under these assumptions about the correspondence relations, the normal application candidate has 2 fewer violations of IDENT[±nas]-BD than does the overapplication candidate, allowing the ranking of IDENT[±nas]-BD over IDENT[±nas]-BR to select normal application. In the absence of this assumption, normal application would not be derivable with the current constraints.[12]

[9] Bruce Hayes suggests to us that we might be able to view BD-faithfulness constraints as general purpose "blockers" for triggering reduplicative underapplication. M&P (p. 92) demonstrate that underapplication can be derived in BRCT when a higher-ranked constraint (a "blocker") rules out an otherwise ideal overapplication candidate. We agree with Hayes's assessment, and plan to investigate this further in future work.

[10] When this ranking condition holds, the relative ranking of IDENT [±nas]-BR with respect to the other two constraints is not relevant. Due to the graphical limitations of a tableau, the tableau in (25) summarizes over only two of the three possible ranking resolutions consistent with IDENT[±nas]-BD ≫ *NV. The other total ranking that derives underapplication is IDENT[±nas]-BD ≫ *NV ≫ IDENT[±nas]-BR.

[11] The notation with the three solid lines is meant to indicate the resolution of an "underlyingly" variable ranking.

[12] Martin Krämer points out to us that the same result would obtain if IDENT[±nas]-BD and IDENT[±nas]-BR were unordered with respect to not just *NV but also *Ṽ, while *NV is still ordered above *Ṽ. This would alleviate the need to posit BD correspondence between morphological base and reduplicant, because NORMAL APPLICATION would

(26) **Normal application ranking:** *NV ≫ IDENT[±nas]-BD ≫ IDENT[±nas]-BR

(27) Variable realization: NORMAL APPLICATION

MORPH BASE: [waŋĩ] / INPUT: /RED, waŋi/			*NV	IDENT[±nas]-BD	IDENT[±nas]-BR	*Ṽ
a.	waŋĩ-waŋĩ	UNDER	*! (ĩ-w)			**
b. ☞	waŋĩ-w̃ãŋĩ	NORMAL		** (w̃,ã)	** (w,a)	****
c.	w̃ãŋĩ-w̃ãŋĩ	OVER		***!* (w̃,ã,w̃,ã)		******
d.	w̃ãŋĩ-waŋĩ	PATHO	*! (ĩ-w)	** (w̃,ã)	** (w,a)	****

5.3.4 Recopying overapplication

Lastly, recopying overapplication is derived by the ranking conditions in (28). Recopying overapplication ((29)c) occurs when maintaining BD-identity is least important. Realizing nasalization on the reduplicant-initial span both eliminates all nasal-oral sequences *and* maintains BR identity. This comes at the expense of diverging from the root in isolation. Crucially, this nasalization is not locally triggered (i.e., not directly spurred by a *NV violation in the reduplicant), but rather motivated through correspondence and faithfulness to the other output constituent, driven by the BR faithfulness constraint IDENT[±nas]-BR. As noted already by M&P, this interaction is derivable only in a framework with BR correspondence or some equivalent mechanism.

(28) **Overapplication ranking:** *NV, IDENT[±nas]-BR ≫ IDENT[±nas]-BD

(29) Variable realization: OVERAPPLICATION

MORPH BASE: [waŋĩ] / INPUT: /RED, waŋi/			*NV	IDENT[±nas]-BR	IDENT[±nas]-BD	*Ṽ
a.	waŋĩ-waŋĩ	UNDER	*! (ĩ-w)			**
b.	waŋĩ-w̃ãŋĩ	NORMAL		*!* (w,a)	** (w̃,ã)	****
c. ☞	w̃ãŋĩ-w̃ãŋĩ	OVER			**** (w̃,ã,w̃,ã)	******
d.	w̃ãŋĩ-waŋĩ	PATHO	*! (ĩ-w)	*!* (w,a)	** (w̃,ã)	****

5.4 Interim Summary

In broad strokes, this revision of M&P's OT analysis is quite successful at capturing the data: through the inclusion of BD correspondence and partially-ordered constraints, we successfully

result when the ranking fragment {*NV ≫ *Ṽ} is ranked above both IDENT[±nas]-BD and IDENT[±nas]-BR. Our modeling results suggest that our analysis provides a better fit to the data, but this remains a viable alternative solution.

derive all and only the three robustly attested outputs. This confirms that BR correspondence is an appropriate and necessary component of the analysis of reduplication.

Nevertheless, this analysis does require certain idealizations, and also inadequately captures various fine details of the phonetic distribution. First, it fails to ascribe any frequency to the infrequent yet attested pathological output due to harmonic bounding. Second, it does not replicate the distributional patterns observed in the experimental results, nor the relative frequency of candidates. According to the R-volume hypothesis (Bane & Riggle, 2008; Riggle, 2010), the typological frequency of a pattern is predicted to correlate with the number of constraint rankings that generate it: the greater the number of compatible rankings, the more frequent the pattern should be. Based on the current analysis, the predicted frequency distribution in descending order is: underapplication (3) > overapplication (2) > normal application (1) > pathological application (0), with the number of compatible rankings given in parentheses. However, this prediction does not align with the experimental findings, which show overapplication as the most frequently attested output.

Lastly, the current analysis misses several gradient tendencies observed in the experimental data (cf. Figure 3). First, the majority of data points clustered near the origin, suggesting that extreme deviations from the oral baseline were relatively rare. Second, there was a noticeable trend for the vowels in the reduplicant and base to exhibit similar degrees of orality or nasality, indicating a possible base-reduplicant correspondence effect at the phonetic level. To address these limitations, the following section develops a Maximum Entropy Harmonic Grammar model [MaxEnt] (Goldwater & Johnson, 2003; Hayes & Wilson, 2008; Hayes, 2022) that incorporates a generative phonetic approach, offering a more nuanced account of the experimental findings.

This is not to say that the results of the OT BRCT analysis ought to be completely discounted. As we will outline below, it is a translation of the correspondence theory and its BR- and BD-faithfulness constraints into equivalent, phonetically-defined constraints that allows us to capture these fine-grained phonetic details. This correlation is what permits reasoning back and forth between the two representational levels of analysis in the manner advocated in this paper.

6 MaxEnt Modeling

6.1 MaxEnt Preliminaries

As a form of Harmonic Grammar (Smolensky & Legendre, 2006; Pater, 2009), MaxEnt uses constraints that are numerically weighted. These weights (w) are multiplied by the constraint violations (C) of an input-output mapping and then summed to get its harmony score (H), i.e., $H(x,y) = \sum_i w_i\, C_i(x,y)$. The harmony score is related to the conditional probability $Pr(y \mid x)$ of an output y given an input x as described in (30):

(30) $Pr(y \mid x) = \frac{exp(-\sum_i w_i C_i(x,y))}{Z},\ where\ Z\ =\ \sum_{y\,\in\,Y(x)} exp(-\sum_i w_i\, C_i(x,y))$

First, the harmony score is negated. This is to ensure that candidates with higher harmony score (i.e., more severe constraint violations, assuming that constraints are positively weighted) are assigned lower probabilities. The next step is to raise e to this negated harmony score and divide this value by Z, where Z is a normalizing term that sums the exponentiated harmony scores for all output candidates $y \in Y(x)$ for a given input x.

Unlike classical Optimality Theory (OT), which selects a single optimal output through strict constraint ranking, MaxEnt assigns a probability distribution over all possible candidates for a given input. When constraint weights differ substantially, MaxEnt closely approximates classical OT: the optimal candidate receives a probability near 1, while all others are assigned vanishingly low probabilities. However, when constraint weights are more similar to each other, the resulting distribution is less skewed, allowing for variation among outputs. Importantly, harmonically bounded candidates can still receive non-zero probabilities, though they will never be the most probable (Hayes & Kaplan, 2023).

6.2 Generative Phonetics

The division of labor between phonology and phonetics remains a complex issue (see Cohn & Huffman 2014 and references therein). Owing to the strikingly similar characteristics of categorical and gradient patterns, there have been attempts to model them in a unified framework using Harmonic Grammar (Flemming, 2001; Katz, 2010; Braver, 2013, 2019; Flemming & Cho, 2017). Building on this tradition, the current model adopts a MaxEnt-based approach to capture gradient phonetic behavior, following Lefkowitz (2017) and Hayes & Schuh (2019).

6.2.1 Discretizing Candidate Set

The candidate set for the MaxEnt model was constructed by discretizing the two-dimensional space of normalized A1−P0 values shown in Figure 3. Specifically, the phonetic space of nasality was divided into a 12 × 12 grid of equal-sized bins, resulting in 144 discrete candidates. Each bin corresponds to a unique pairing of nasality values in the base and reduplicant vowels, with the center of each bin representing a candidate in the MaxEnt model's input. The outcome of this discretization is visualized in the heat map given in Figure 7, where darker shading indicates a higher frequency of observed data points. Non-zero grid cells are also labeled by their counts.

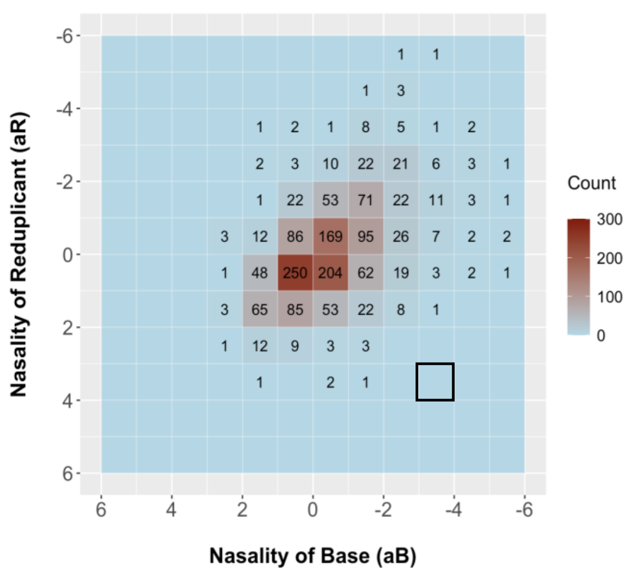

Figure 7. Heat map of the experimental results. Darker shading indicates higher frequency of observed tokens, with non-zero grid cells labeled by their counts. The highlighted cell represents the candidate located at (−3.5, +3.5).

As an illustration, the highlighted cell in the heat map corresponds to the candidate (−3.5, +3.5), where the base vowel is 3.5 units more nasal than the oral baseline and the reduplicant vowel is 3.5 units more oral. The heat map also helps visualize key phonetic tendencies in the data: (i) a higher frequency of data points near the origin (i.e., the oral baseline), and (ii) a phonetic correspondence effect, where vowels in the reduplicant and the base tended to exhibit similar degrees of nasality.

6.2.2 Gradientizing Constraint Violations

The next step involves computing constraint violations for each candidate in the discretized phonetic space. The constraint set employed here is nearly identical to that used in the classical OT analysis presented in Section 5.[13] However, in the current model, violations are assigned in a gradient rather than categorical fashion. Constraints are defined in terms of normalized A1−P0 values: a_R and a_B denote the degree of orality/nasality for the reduplicant and base vowels after normalization, respectively. The revised definitions of the constraints and the method of assessing violations are given below.

(31) **IDENT[±nas]-BD:** Penalize deviation from the oral baseline in both the reduplicant and base. Take the absolute magnitude and sum the values.[14]
Violation = $| a_B | + | a_R |$

(32) **IDENT[±nas]-BR**: Penalize dissimilarity in orality/nasality between the base and reduplicant. Take the absolute magnitude.
Violation = $| a_B - a_R |$

(33) ***Ṽ (*[+nas,−cons]):** Penalize any vowel that is more nasal than the oral baseline (i.e., any negative number). Take the absolute magnitude and sum the values.
Violation = $| \min(0, a_B) | + | \min(0, a_R) |$

(34) ***NV (*[+nas][−nas,−cons]):** Penalize discrepancies in nasality between a vowel and a preceding nasal segment, which is taken to be the most negative value in the dataset (−6). Take the absolute magnitude. Given the context under consideration, this constraint applies exclusively to a_B, the second constituent in a reduplicative construction.
Violation = $| -6 - a_B |$

The tableau in (35) illustrates how gradient constraint violations are assessed, using one representative candidate from each quadrant. Take candidate ((35)b), the normal application candidate (−1.5, +0.5), as an example. This candidate incurs a violation score of 2 for the constraint

[13] The constraint IDENT[±nas]-IO is excluded from the current modeling for two main reasons. First, it is ranked at the bottom of the hierarchy in the classical OT analysis, and would therefore be expected to receive a weight close to zero in a weighted grammar. Second, its violation profile is difficult to define due to the principle of *Richness of the Base* (Prince and Smolensky 2004; Smolensky 1996), which assumes that all possible inputs are equally available and thus renders input-output mappings less informative for learning the weight of IDENT[±nas]-IO.
[14] Squared differences are more commonly used in modeling work within this tradition. We did implement an alternative version of our model using squared differences. Although this model achieved a comparable fit to the data, the resulting constraint weights were intractably small, making interpretation difficult. For this reason, we adopt absolute differences in the present modeling simulation.

IDENT[±nas]-BD: the base vowel is 1.5 units more *nasal* than the oral baseline, and the reduplicant vowel is 0.5 units more *oral* than the baseline, resulting in a total violation of |−1.5| + |+0.5| = 2. For IDENT[±nas]-BR, the violation score is likewise 2, reflecting the absolute difference in orality/nasality between the base and reduplicant vowels: |−1.5 − (+0.5)| = 2. Moreover, the candidate also violates *NV by 4.5 units. This constraint penalizes base vowels following a nasal segment that are not nasalized enough, operationalized here as a penalty proportional to the distance from −6 (i.e., the most negative value in the dataset). The base vowel's value of −1.5 thus yields a violation of |−6 − (−1.5)| = 4.5. Lastly, the candidate violates *Ṽ by 1.5 units, as only vowels more nasal than the oral baseline violate this constraint. No penalty is assigned for vowels that are more oral than the baseline (e.g., the reduplicant in this case). Similar calculations apply to the other candidates in Tableau (35) and the entire candidate set.

(35) Gradient violation profiles for four candidates, each representing one quadrant of the phonetic space. The values in the frequency column indicate the number of data points falling within the corresponding grid cell.

MORPH BASE: $\tilde{a}_u$		Freq.	IDENT[±nas]-BD	IDENT[±nas]-BR	*NV	*Ṽ
INPUT: /a_B, a_R/						
a. (+5.5, +3.5)	UNDER	0	9	2	11.5	0
b. (−1.5, +0.5)	NORMAL	62	2	2	4.5	1.5
c. (−5.5, −2.5)	OVER	1	8	3	0.5	8
d. (+4.5, −1.5)	PATH	0	6	6	10.5	1.5

6.3 Learning Task and Simulation

Given the set of constraints presented in Tableau (35) and the training data derived from discretizing the phonetic space of nasality (as described in Section 6.2.1), the learning objective of a MaxEnt model is to determine the constraint weights that maximize the likelihood of the observed training data, denoted as *L(D)* (Goldwater & Johnson, 2003:113). This likelihood is calculated as the product of the conditional probabilities assigned by the model to each individual training data point, as shown in Equation (36), where *N* is the total number of training data points and (x_i, y_i) represents the i^{th} observed input-output mapping. In practice, however, this product can result in extremely small values when *N* is large, making computation intractable. To address this, it is standard to use the *log* likelihood instead. Since the logarithm of a product is equal to the sum of the logarithms, Equation (36) can be reformulated as Equation (37), which expresses the log likelihood as the sum of the log probabilities of individual data points. The learning simulation was done in R (R Core Team 2024) using the `maxent.ot` package (Mayer et al., 2024).

(36) $L(D) = \prod_{i=1}^{N} Pr(y_i \mid x_i)$ (37) $log\ L(D) = \sum_{i=1}^{N} log\ Pr(y_i \mid x_i)$

6.4 Simulation Outcome

The simulation outcome is shown in Figure 8 below. The left panel of the figure compares the predicted probabilities generated by the MaxEnt model with the observed probabilities derived from the experimental data. Each point represents a candidate, with the size of the point proportional to the logarithm of the number of data points in the corresponding grid. Candidates that fall on the diagonal line indicate a perfect match between the model's predictions and observed outcomes. The right panel presents a heat map generated based on the model's predictions.

Overall, the MaxEnt model achieves a strong fit to the experimental data, explaining 91% of the variance ($r^2 = 0.91$).[15,16] Moreover, the heat map generated from the model's predictions successfully captures the gradient patterns observed in the data, including the clustering of data points around the origin and the phonetic correspondence effect between base and reduplicant vowels.

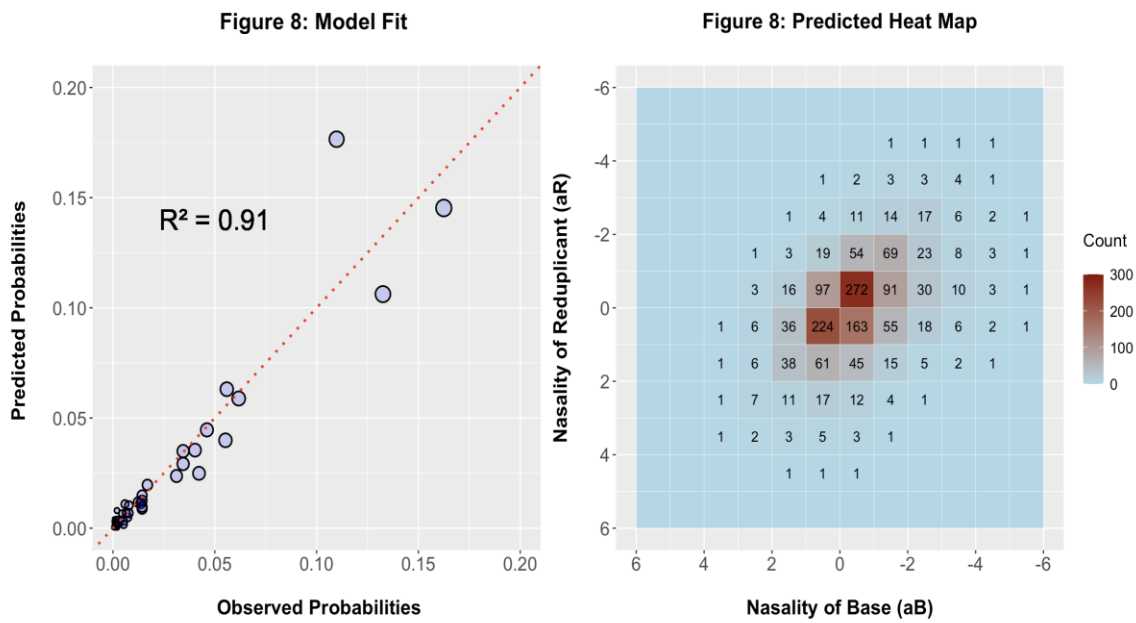

Figure 8. (Left) Observed vs. predicted probabilities. (Right) Heat map generated from the model's predictions.

A key question for the present analysis is whether the MaxEnt model can accurately predict the pathological candidates that are harmonically bounded under a classical OT analysis. To assess this, we examined the model's largest prediction errors. Table (38) lists candidates whose prediction error — calculated as the difference between predicted and observed probabilities — exceeds ±0.01. Positive values indicate overprediction, while negative values indicate underprediction. Notably, none of the large errors involve candidates from the pathological quadrant, suggesting that the model performs well in predicting these forms. Most discrepancies occur near the origin, where the model tends to overestimate the probabilities of candidates from the overapplication quadrant, and underestimate those from the normal and underapplication

[15] To guard against overfitting, we conducted five-fold cross-validation using the `cross_validate()` function from the `maxent.ot` package. The data were randomly divided into five equal slices. For each fold, one slice was held out while the model was trained on the remaining four. We then computed the log likelihood of the held-out data using the trained model. This procedure was repeated until each slice served as the held-out set once. The mean log likelihoods for the training and held-out data were highly similar (−3.074 vs. −3.078), indicating that the model generalizes well and does not overfit the experimental data.

[16] Modeling simulations were also conducted at the individual participant level, yielding an average r^2 value of 0.73 (SD = 0.11). Additional details can be found in the online supplementary materials.

quadrants. This pattern suggests a regularization effect, whereby the model exaggerates the distinctions among the non-pathological quadrants (Over > Under > Normal), in accordance with their overall relative frequency.

(38) Candidates whose error terms are greater than ±0.01

Candidate	Category	Frequency	Predicted	Observed	Error	Error Type
(−0.5, −0.5)	OVER	169	0.177	0.110	+0.067	Overprediction
(−0.5, +0.5)	NORMAL	204	0.106	0.132	−0.026	Underprediction
(+0.5, +0.5)	UNDER	250	0.145	0.162	−0.017	Underprediction
(+0.5, +1.5)	UNDER	85	0.040	0.055	−0.015	Underprediction
(+1.5, +1.5)	UNDER	65	0.025	0.042	−0.017	Underprediction

In addition to providing a good overall fit, the constraint weights learned by the MaxEnt model closely mirror the constraint ranking established in the classical OT analysis, reproduced in Figure 9 below. First, the constraints grouped in the highest stratum receive comparable weights — an essential property for generating the variable outputs observed in the experimental data. Second, the ranking that accounts for the allophonic pattern of nasal spreading in Malay (*NV ≫ *Ṽ ≫ IDENT[±nas]-IO) is also preserved under the weighted constraint grammar.[17]

Constraint	Weight
IDENT[±nas]-BD	0.62
IDENT[±nas]-BR	0.67
*NV	0.52
*Ṽ	0.33
IDENT[±nas]-IO	0

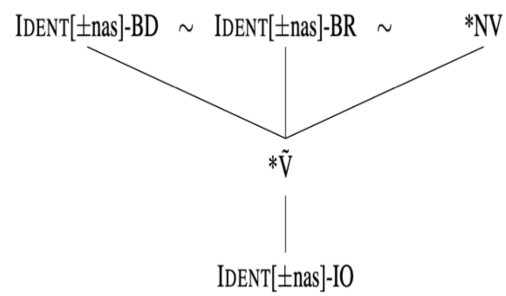

Figure 9. Constraint weights derived from the MaxEnt model (left), compared with the constraint ranking from the classical OT analysis presented in Section 5 (right).

A central claim of this paper is the necessity of using base-reduplicant correspondence constraints (i.e., IDENT[±nas]-BR) in modeling the experimental results. To evaluate the contribution of IDENT[±nas]-BR, we fitted a simplified MaxEnt model that excluded this constraint. A likelihood ratio test comparing the full and reduced models revealed a significant improvement in model fit when IDENT[±nas]-BR was included ($\chi^2(1) = 472.70$, $p < .001$). The base-derivative correspondence constraint IDENT[±nas]-BD was found to be similarly essential: a

[17] Notably, the weights of the constraints required to produce the allophonic pattern are not widely apart. This suggests that the allophonic pattern may be less categorical than previously assumed. Specifically, the model predicts that vowels in forms such as /wa/ — which lack a nasal trigger — may still exhibit some degree of spontaneous nasalization. Conversely, vowels in forms like /ma/, which contain a nasal trigger, may not show full nasalization. Kate Mooney (p.c., November 1, 2024) has noted that this gradient pattern appears to be impressionistically true in Indonesian, a language that exhibits a similar nasal spreading system, with anti-allophonic nasalization perhaps developing sociolinguistic salience for certain populations. We plan to investigate this further in future work.

reduced model excluding this constraint significantly underperformed compared to the full model ($\chi^2(1) = 293.16$, $p < .001$)[18]. We interpret these results as strong support for a theory of reduplication that incorporates different kinds of correspondence relations, regardless of the level of representation or analysis.

7 Conclusion

Our acoustic experiment revealed an intricate interaction between nasal spreading and reduplication in Malay, repeated below in (39), arranged in decreasing order of frequency of attestation.

(39) Outputs of free variation
 a. **Output C:** [w̃ãŋĩ-w̃ãŋĩ] OVERAPPLICATION
 b. **Output A:** [waŋĩ-waŋĩ] UNDERAPPLICATION
 c. **Output B:** [waŋĩ-w̃ãŋĩ] NORMAL APPLICATION
 d. **Output D:** [w̃ãŋĩ-waŋĩ] PATHOLOGICAL APPLICATION

While our findings present a more nuanced picture than originally reported by Onn (1976), they crucially confirm the existence of the recopying overapplication pattern, a phenomenon that has drawn considerable attention in the theoretical literature. We demonstrated that an extension of McCarthy & Prince’s original Base-Reduplicant Correspondence Theory (BRCT) analysis can account for the three robustly observed outputs within a grammar featuring variable constraint rankings. Additionally, we showed that a Maximum Entropy Harmonic Grammar model couched in a generative phonetic approach can produce the harmonically bounded candidate with the pathological application of nasal spreading, while also capturing the gradient patterns observed in the dataset (e.g., clustering of data points around the origin and phonetic correspondence effects between base and reduplicant vowels).

Recopying overapplication is a significant phenomenon for reduplicative theory because it cannot be derived serially (McCarthy & Prince, 1995:43–46) in the absence of a representational mechanism equivalent to BR-correspondence. This problem holds of most if not all recent constraint-based serial theories, developed explicitly as alternatives to BRCT. This includes Morphological Doubling Theory (Inkelas & Zoll, 2005:221, n. 18), Reduplication in Stratal OT (Kiparsky, 2010:3–4), and Serial Template Satisfaction in Harmonic Serialism (McCarthy et al., 2012:203), whose proponents have contested the existence of the recopying overapplication data in Malay, alongside acknowledgements that the pattern cannot be generated by their theory.

On the other hand, this literature unanimously agrees that BRCT can derive recopying overapplication, because of its use of BR-correspondence. It seems possible that certain other theories that make use of similar representational devices, such as the looped representations used by Raimy (2000:16–18, 2011:2398–2399) or the linked representations used by Frampton (2009), may be able to derive this pattern, even though they are serial and ruled-based. This makes clear that it is BR-correspondence that is the crucial element in deriving these results, rather than parallelism *per se*. However, it is yet to be shown how a Raimy- or Frampton-style framework could handle the variable and gradient patterns found in the experimental results.

[18] Additional model comparison metrics that balance model fit and complexity such as Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) converge on the same conclusion. Full model specifications and evaluation metrics are available in the online supplementary materials.

The aforementioned opponents of BRCT have used the dubious attestation of the Malay nasal spreading pattern, and thus the potential non-existence of a true recopying overapplication pattern, as an argument against BRCT on the basis of overgeneration, as it would predict a pattern that was not attested in the world's languages. Now that we have shown that the pattern does exist, the argument reverses: these non-BRCT alternatives suffer from an undergeneration problem, failing to derive an attested pattern. Our verification of the Malay recopying overapplication pattern is therefore a strong argument against these alternatives, and in favor of BRCT, as it is the only theory which embraces BR-correspondence.

This paper also makes clear that this result holds regardless of the level of abstractness at which the pattern is being considered. At the phonological level, BR (and BD) correspondence was crucial for deriving the simplified pattern of variation extracted from the experimental results. Likewise, at the phonetic level, an alternative implementation of the concept of BR (and BD) correspondence in terms of discretized phonetic measures yielded complementary results, capturing details appropriate to that level of representation. This dual approach to phonetic and phonological explanation and theorizing expands the types of data which may bear on complex and thorny problems in these domains. It also serves as a reminder that detailed phonetic evidence may be crucial in resolving long-standing phonological questions.

References

Adelaar, Alexander K. (1992). *Proto-Malayic: The reconstruction of its phonology and parts of its lexicon and morphology*. Depepartment of Linguistics, Research School of Pacific Studies, the Australian National University.

Ahmad, Zaharani. (2005). *The phonology-morphology interface in Malay: an optimality theoretic account*. Canberra: Asia-Pacific Linguistics, Research School of Pacific and Asian Studies, Australian National University.

Anttila, Arto Tapani. (1998). *Variation in Finnish phonology and morphology*. Stanford University.

Bane, Max, and Jason Riggle. (2008). Three correlates of the typological frequency of quantity-insensitive stress systems. *SIGMORPHON 2008 - 10th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology, Proceedings of the Workshop*, 29–38.

Benua, Laura. (1995). Identity Effects in Morphological Truncation. *UMass Occasional Papers in Linguistics 18: Papers in OT*, *1995*(1), 1–61.

Benua, Laura. (1997). *Transderivational identity: Phonological relations between words*. University of Massachusetts Amherst.

Blust, Robert. (2013). *The Austronesian languages*. Canberra: Asia-Pacific Linguistics, Research School of Pacific and Asian Studies, Australian National University.

Boersma, Paul., and David Weenink. (2024). *Praat: doing phonetics by computer*.

Braver, Aaron. (2013). *Degrees of incompleteness in neutralization: Paradigm uniformity in a phonetics with weighted constraints*. Rutgers The State University of New Jersey.

Braver, Aaron. (2019). Modelling incomplete neutralisation with weighted phonetic constraints. *Phonology*, Vol. 36, pp. 1–36.

Burzio, Luigi. (1996). Surface Constraints Versus Underlying Representation. *Current Trends in Phonology: Models and Methods*, 97–122.

Chen, Marilyn Y. (1997). Acoustic correlates of English and French nasalized vowels. *The Journal of the Acoustical Society of America*, *102*(4), 2360–2370.

Cho, Taehong, Daejin Kim, and Sahyang Kim. (2017). Prosodically-conditioned fine-tuning of coarticulatory vowel nasalization in English. *Journal of Phonetics*, *64*, 71–89.

Clynes, Adrian, and David Deterding. (2011). Standard Malay (Brunei). *Journal of the International Phonetic Association*, Vol. 41, pp. 259–268.

Cohn, Abigail C., and Marie K. Huffman. (2014). Interface between phonology and phonetics. *Oxford University Press*.

Flemming, Edward. (2001). Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology*, *18*(1), 7–44.

Flemming, Edward, and Hyesun Cho. (2017). The phonetic specification of contour tones: Evidence from the Mandarin rising tone. *Phonology*, *34*(1), 1–40.

Frampton, John. (2009). *Distributed Reduplication*. Linguistic Inquiry Monograph. MIT Press.

Garellek, Mark, Amanda Ritchart, and Jianjing Kuang. (2016). Breathy voice during nasality: A cross-linguistic study. *Journal of Phonetics*, *59*, 110–121.

Goldwater, Sharon, and Mark Johnson. (2003). Learning OT Constraint Rankings Using a Maximum Entropy Model. *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*, 111–120.

Hayes, Bruce. (2022). Deriving the Wug-Shaped Curve: A Criterion for Assessing Formal Theories of Linguistic Variation. *Annual Review of Linguistics*, Vol. 8, pp. 473–494.

Hayes, Bruce, and Aaron Kaplan. (2023). Zero-Weighted Constraints in Noisy Harmonic Grammar. *Linguistic Inquiry*, 1–14.

Hayes, Bruce, and Russell G. Schuh. (2019). Phonological analysis: Metrical structure and sung rhythm of the hausa rajaz. *Language*, *95*(2), e253–e299.

Hayes, Bruce, and Colin Wilson. (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, *39*(3), 379–440.

Inkelas, Sharon, and Cheryl Zoll, C. (2005). *Reduplication : doubling in morphology* (Vol. 106). Cambridge University Press.

Kager, René. (1999). Surface Opacity of Metrical Structure in Optimality Theory. In *The derivational residue in phonological Optimality Theory* (pp. 207–246).

Katz, Jonah. (2010). *Compression effects, perceptual asymmetries, and the grammar of timing*. Massachusetts Institute of Technology.

Kenstowicz, Michael. (1996). Base-Identity and Uniform Exponence: Alternatives to Cyclicity. *Current Trends in Phonology: Models and Methods*, *1*, 363–393.

Kiparsky, Paul. (2010). Reduplication in Stratal OT. In *Reality Exploration and Discovery: Pattern Interaction in Language & Life*.

Lefkowitz, Lee Michael. (2017). *Maxent harmonic grammars and phonetic duration*. University of California, Los Angeles.

Matisoff, James A. (1975). Rhinoglottophilia: the mysterious connection between nasality and glottality. *Nasálfest: Papers from a Symposium on Nasals and Nasalization*, 265–287. Stanford: Language Universals Project, Department of Linguistics, Stanford University.

Mayer, Connor, Kie Zuraw, and Adeline Tan. (2024). maxent.ot: A package for doing Maximum Entropy Optimality Theory in R. *Phonological Data and Analysis*, *6*(4), 1–44.

McCarthy, John. J., Wendell Kimper, and Kevin Mullin. (2012). Reduplication in Harmonic Serialism. *Morphology*, *22*(2), 173–232.

McCarthy, John. J., and Alan Prince. (1995). Faithfulness and Reduplicative Identity. *University of Massachusetts Occasional Papers* 18, GLSA, Amherst, pp. 249–384.

Nomoto, Hiroki, and Hooi Ling Soh. (2019). Malay. In *The Mainland Southeast Asia linguistic*

area (pp. 475–522).
Ohala, John J. (1975). Phonetic explanations for nasal sound patterns. *Nasálfest: Papers from a Symposium on Nasals and Nasalization*, 289–316. Stanford: Language Universals Project, Department of Linguistics, Stanford University.
Onn, Farid Mohamed. (1976). *Aspects of Malay Phonology and Morphology*. University of Illinois, Urbana-Champaign.
Pater, Joe. (2009). Weighted constraints in generative linguistics. *Cognitive Science*, *33*(6), 999–1035.
Prince, Alan, and Paul Smolensky. (1993/2004). Optimality Theory: Constraint interaction in generative grammar. *Optimality Theory in phonology: A reader*, 1-71.
Raimy, Eric. (2000). *The Phonology and Morphology of Reduplication* (Vol. 52). Walter de Gruyter.
Raimy, Eric. (2011). Reduplication. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell Companion to Phonology* (pp. 2418–2448).
Riggle, Jason. (2010). Sampling rankings. *Rutgers Optimality Archive*, *1075*.
Scarborough, Rebecca. (2013). Neighborhood-conditioned patterns in phonetic detail: Relating coarticulation and hyperarticulation. *Journal of Phonetics*, *41*(6), 491–508.
Siah, Jian-Leat, Sam Zukoff and Feng-fan Hsieh. (2025). Reduplicative Opacity in Malay Revisited: Preliminary Phonetic Evidence for Variable "Recopying" and BRCT. *Proceedings of the Annual Meetings on Phonology*.
Smolensky, Paul. (1996). The initial state and ‘richness of the base’ in Optimality Theory. *Rutgers Optimality Archive*, *293*.
Smolensky, Paul, and Géraldine Legendre. (2006). *The harmonic mind: From neural computation to optimality-theoretic grammar (Cognitive architecture) Vol. 1*. MIT press.
Styler, Will. (2017). On the acoustical features of vowel nasality in English and French. *The Journal of the Acoustical Society of America*, *142*(4), 2469–2482.
Wilbur, Ronnie Bring. (1973). *The Phonology of Reduplication*. University of Illinois at Urbana-Champaign.
Yang, Yifan. (2023). Rapa nui: A case for correspondence in reduplication. *Linguistic inquiry*, *54*(2), 395-412.
Zellou, Georgia. (2017). Individual differences in the production of nasal coarticulation and perceptual compensation. *Journal of Phonetics*, *61*, 13–29.
Zellou, Georgia, and Meredith Tamminga. (2014). Nasal coarticulation changes over time in Philadelphia English. *Journal of Phonetics*, *47*, 18–35.