

Phonotactic Complexity across Dialects

Ryan Soh-Eun Shim*, Calvin Chang*, David R. Mortensen

Institute for Natural Language Processing, University of Stuttgart

Language Technologies Institute, Carnegie Mellon University

soh-eun.shim@ims.uni-stuttgart.de

{kalvinc, dmortens}@cs.cmu.edu

Abstract

Received wisdom in linguistic typology holds that if the structure of a language becomes more complex in one dimension, it will simplify in another, building on the assumption that all languages are equally complex (Joseph and Newmeyer, 2012). We study this claim on a micro-level, using a tightly-controlled sample of Dutch dialects (across 366 collection sites) and Min dialects (across 60 sites), which enables a more fair comparison across varieties. Even at the dialect level, we find empirical evidence for a tradeoff between word length and a computational measure of phonotactic complexity from a LSTM-based phone-level language model—a result previously documented only at the language level. A generalized additive model (GAM) shows that dialects with low phonotactic complexity concentrate around the capital regions, which we hypothesize to correspond to prior hypotheses that language varieties of greater or more diverse populations show reduced phonotactic complexity. We also experiment with incorporating the auxiliary task of predicting syllable constituency, but do not find an increase in the strength of the negative correlation observed.

Keywords: phonotactic complexity, linguistic niche hypothesis, compensation hypothesis

Code: <https://github.com/cmu-llab/phonotactic-complexity-across-dialects>

1. Introduction

Phonotactics refers to a set of language- or dialect-specific constraints on what constitutes a licit or illicit sound sequence. Phonotactic complexity refers to the variety of structures allowed at different positions within a syllable or word—essentially how unpredictably a language variety’s phonemes behave at different positions. If we view the possible upcoming phonemes as branches of a tree structure, where the root denotes the start of the word, the phonotactic complexity is the size of the tree. For instance, a language that requires VCV sequences to obey vowel harmony has lower phonotactic complexity than another VCV-language without this phonotactic constraint, since the rule reduces the number of possible vowels in each position (Pimentel et al., 2020).

Recent results have shown a moderate negative correlation between phonotactic complexity and average word length in a dataset of 1,016 basic concept words across 106 languages (Pimentel et al., 2020). A decrease in word length accompanies an increase in phonotactic complexity, which the compensation hypothesis (Hockett, 1958; Moran and Blasi, 2014) would argue occurs as a compensatory mechanism. However,

there has been comparably little work that examines this hypothesis in the context of dialects of the same language. Such research would provide a microcosmic view of complexity where most variables (e.g., areal influences, phylogenetic biases, variation in descriptive conventions) are held relatively constant. This addresses the problem of typological imbalance in Pimentel et al. (2020)’s study, where the NorthEuraLex dataset used (Dellert et al., 2020) favors Uralic and Indo-European languages. Examining dialects additionally holds much interest for dialectology, since geographical effects such as spatial autocorrelation can be observed, and areas of low or high linguistic complexity (in some dimension) can be explained. For example, geographic and demographic factors such as population size may (inversely) correlate with complexity (Lupyan and Dale, 2010; Dale and Lupyan, 2012). Indeed, our results provide empirical support for such a hypothesis.¹

We measure the phonotactic complexity of a dialect with Shannon entropy (bits per phoneme), estimated with Pimentel et al. (2020)’s LSTM-based phonotactic language model (Equation 1). While predicting the next phoneme accounts for the number of options at some position, their model only considers the order of segments in a word and fails to account for the interaction

¹Recently, Shcherbakova et al. (2023) do not find such an inverse correlation in a dataset of 1,314 languages (Skirgård et al., 2023), but their study examines different dimensions of linguistic complexity.

*These authors contributed equally.




Figure 1: Contour plot of GAM results using a thin plate regression spline, where the response variable is phonotactic complexity (left) and average word length (right) of Dutch dialects, and the predictors, their latitude and longitude. Colder colors (blue) indicate low complexity, while warmer colors (yellow) indicate higher values. The larger blue region on the phonotactic complexity plot is the region of Holland, which is correspondingly yellow in the word length plot.

between syllable structure and phonotactic constraints (subsection 4.2). Our contributions are as follows:

1. **Phonotactic tradeoffs across dialect continua:** We corroborate the negative correlation between word length and phonotactic complexity found in Pimentel et al. (2020) even when examined in our controlled setting of a large number of very similar language varieties. The correlation also holds across two typologically diverse languages.
2. **Social explanation of complexity distributions:** Areas of low phonotactic complexity concentrate around the capital regions, which supports the hypothesis that varieties with more diverse populations tend to have simpler phonology and morphology, potentially due to learnability constraints (Lupyan and Dale, 2010; Dale and Lupyan, 2012; Bentz and Winter, 2014).
3. **Supervised syllable constituency prediction:** We incorporate syllabification into a phonotactic language model by virtue of multi-task learning, and observe that knowledge of syllable constituency does not strengthen the negative correlation in our data.

2. Related Work

Pimentel et al. (2020) are not the first to model

phonotactics computationally. Hayes and Wilson (2008) use a maximum entropy grammar to assign scores of well-formedness to sequences of phonemes based on how many phonotactic constraints are violated. The model learns *SPE*-style (Chomsky and Halle, 1968) constraints from primary data.

Mayer and Nelson (2020) find that RNN-based phoneme-level language models can learn challenging phonotactic patterns, such as Finnish vowel harmony (a long-distance pattern). While they train on articulatory feature vectors, they do not incorporate syllable structure into their model. They hypothesize that an Elman RNN may suffice for the task of phonotactic modeling (though we use an LSTM).

Using Mayer and Nelson (2020)'s model, Kirby (2021a) finds that the phonotactic probability of a syllable is not affected by the placement of tone in the phonetic transcription. He does not train on transcriptions of words and instead trains on all possible syllables for 4 languages with syllables as tone-bearing units (Mandarin, Thai, Vietnamese, Cantonese) to maximize the probability of the syllable lexicon.

Steuer et al. (2023) find that LSTM-based phoneme-level language models can capture vowel harmony, a phonotactic constraint, in different languages. They quantify the degree of vowel harmony with an entropy-based measure derived from the likelihood but revise the language model objective to only predict the next vowel. In con-

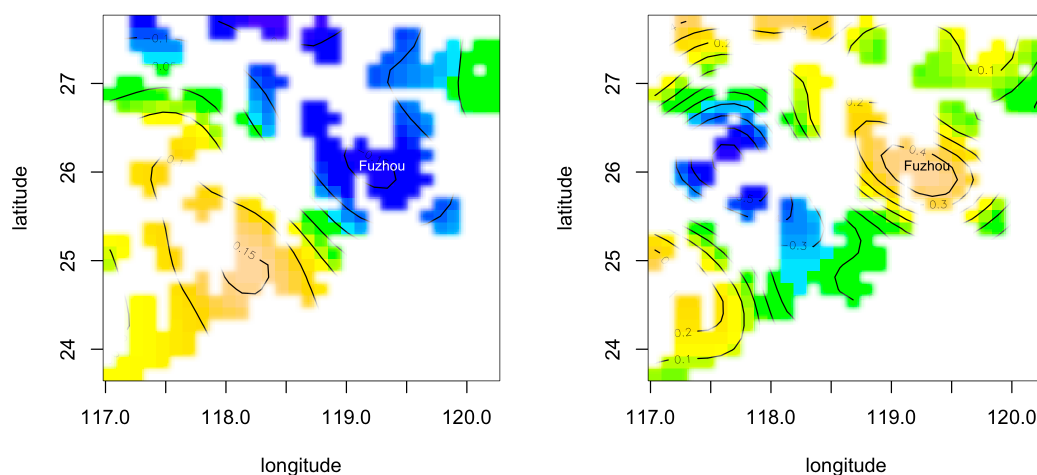


Figure 2: Contour plot of GAM results using a thin plate regression spline, where the response variable is phonotactic complexity (left) and average word length (right) of Min dialects, and the predictors, their latitude and longitude. Colder colors (blue) indicate low complexity, while warmer colors (yellow) indicate higher values. The city of Fuzhou is marked out in text in both figures.

trast, the language model we use is not restricted to a single phonotactic constraint.

Daland et al. (2011a) use a version of Hayes and Wilson (2008)'s phonotactic model supervised with syllable structure and find that models need to incorporate both sonority and syllabification to capture the phonotactic phenomenon of sonority projection. We implicitly incorporate sonority by incorporating syllable constituency obtained via sonority-based syllabification (subsection 4.3).

Regarding studies on geographical distributions of linguistic complexity, Wichmann and Holman (2023) find that across 1,267 languages, average word length exhibits the behavior typical of typological linguistic features, showing sensitivity to areal influence detectable over a 5000 km range. Average word length is also significantly inversely correlated with (log) population size when compared across six macroareas, but inconsequential when measured within each macroarea.

Our study extends this line of typological research in linguistic complexity by focusing instead on a large number of closely-related language varieties, which has the benefit of holding many variables relatively constant in examining how word length interacts with other social or linguistic variables.

3. Dataset

Refer to Table 1 for a summary of the datasets we use in our work. Below, we briefly profile phono-

tactic constraints within Dutch and Min dialects and dialectal variation of such constraints—which affects phonotactic complexity. We also describe the corresponding datasets we examine in this work.

3.1. Dutch

3.1.1. Phonotactic Properties of Dutch Dialects

The Dutch dialect landscape in the Netherlands is typically considered to be a continuum, where abrupt breaks rarely occur (Tældeman and Hinskens, 2014; Heeringa and Nerbonne, 2001). Such a continuum yields a particularly interesting situation for the study of phonotactic complexity, where a given pair of sites may share nearly identical phoneme inventories but display significant differences in phonological phenomena (Tældeman and Hinskens, 2014). For instance, /v/ varies across Dutch dialects in the onset and in the coda (Van der Torre, 2003). In particular, in the onset, the /v/ in /vr-/ sometimes becomes a fricative — [vr-] or [fr-] — due to a constraint that onsets with two sonorants are illicit. /v/ is sometimes replaced with [m] and [b] due to constraints on place of articulation. Such phonological differences form the basis for differences in phonotactic complexity between Dutch dialects.

3.1.2. Dutch Dialect Data

Our data comes from the Goeman-Taeldeman-Van Reenen-Project (Taeldeman, Johan and Goeman, A, 1996), which is a Dutch dialect dataset collected between 1980 and 1995 in the Netherlands and Belgium. It contains phonetic transcriptions of 1,876 lexical items across 613 dialect sites elicited in a concept-aligned manner that enables fair comparison between sites.² This dataset has enjoyed widespread usage in dialectometry (Wieling et al., 2007), and the version we utilize in this study comes in a preprocessed form that is typical for dialectometric usage, where the diacritics are removed to reduce sparsity. Furthermore, a subset of 562 lexical items are selected to exclude morphological inflections of the same item, along with multiword expressions. Due to differences in transcription between the data in Belgium and in the Netherlands, we base our analysis on data from the latter country, which contains 424 sites (Wieling et al., 2007). After removing Frisian, a separate language that belongs to a different branch of West Germanic, we have 366 Dutch sites.

3.2. Min

3.2.1. Phonotactic Properties of Min Dialects

Min is a subgroup of internally divergent Sinitic languages spoken primarily in the region of Fujian, China, along with several migration-induced speech communities in Taiwan and Southeast Asia. As Sinitic languages, Min languages have tone superimposed upon syllables, where the syllable structure in Min is $C_1V(C_2)$ (Lien, 2015), but variation exists across dialects as to what is allowed in each position. A notable example is the phenomenon of denasalization in Southern Min, where for certain dialects (e.g. Xiamen, Zhangzhou and Quanzhou), if the coda of a closed syllable is a nasal, the initial must not be nasal (Lien, 2000). However, the exact rule does not equally apply to all dialect sites, where denasalization may occur in different forms. Compounded with the differences between dialects in licit phonemes for each position in the syllable and differences in tonal systems, considerable variation in phonotactic complexity occur even between closely-related varieties of Min.

3.2.2. Min Dialect Data

Our Min data comes from Centre for the Protection of Language Resources of China (2023), which contains phonetic transcriptions for 1,000

²The original dataset contains 1876 items, but we use Wieling et al. (2007)'s 562-item subset.

characters and 1,200 concepts across 1,289 Sino-Tibetan varieties in mainland China and Taiwan.³ We only include the Min dialects spoken in Fujian Province (where most Min dialects are concentrated).

We choose concept-aligned over character-aligned data for Min since the number of phones per character does not vary much due to the relatively constant syllable structure across Sinitic languages⁴, and character-aligned data would artificially restrict the word length since concepts are not always expressed as single characters in Chinese. As with any concept-aligned data, this may mean that different Min dialect entries may not be cognate with each other due to lexical innovation or borrowing. Furthermore, for concept-aligned data, the elicitation process often yields pronunciation variants that may or may not be related to one another within a site. For the purposes of this study, we pick the first variant, as our goal is to model the phonotactic sequence, and would arguably not lose essential information in keeping only one variant.

4. Methods

4.1. Phonotactic Language Modeling

We use Pimentel et al. (2020)'s unidirectional LSTM-based (Hochreiter and Schmidhuber, 1997) phoneme-level language model, which assigns probabilities to the phonetic transcriptions of words and learns phonotactically valid sequences of phones. A separate language model is learned for each dialect in our dataset. When the correlation between phonotactic complexity and word length is calculated, word length is measured in phonemes. To measure the phonotactic complexity of a language (or dialect in our case), Pimentel et al. (2020) estimate the Shannon entropy of the variety, which quantifies the average information encoded in each phoneme of a word (bits per phoneme). Entropy is an appropriate measure of phonotactic complexity, given the branching interpretation we provide in section 1. Because it is intractable to iterate over all possible strings to calculate the exact Shannon entropy, the authors use the average negative log-probability of a word to approximate the cross-entropy between the true, unknown phonotactic probability distribution of the language p_{lex} and the one learned by the model q_{lex} . Cross-entropy in turn provides an upper bound on the true entropy:

³Unfortunately, they do not allow the data to be publicly released. We obtained our version on June 12, 2023. The script we used to obtain the data is in the code.

⁴Each character is one syllable.

Language	Source	# dialects (we use)	# words
Dutch	Taldeman, Johan and Goeman, A (1996)	366	562
Min	Centre for the Protection of Language Resources of China (2023)	60	1200

Table 1: Statistics on the datasets in our experiments

$$H(p_{lex}) \leq H(p_{lex}, q_{lex}) \approx -\frac{1}{N} \sum_{i=1}^N \log q_{lex}(\tilde{\mathbf{x}}^{(i)}) \quad (1)$$

We refer the reader to their work for the full derivation.

4.2. Phonotactics and Syllable Structure

Pimentel et al. (2020)’s phonotactic language model does not explicitly consider how the location of a phoneme or set of phonemes within a syllable can determine the behavior of the phoneme. Onsets, for example, are often the locus of phonemic contrast, while codas are often absent or are the location of positional neutralization (Hayes, 2008). Phonological rules often apply at such syllable edges (Blevins, 1995). Syllable constituents—the onset, nucleus, and coda—also differ in the types of complex (multi-segment) sequences allowed (Blevins, 1995). Thus an accurate account of phonotactic constraints and thus phonotactic complexity must account for syllable structure. If we measure phonotactic complexity by the number of options at some position in a phoneme then how many options are next depends on what syllable constituent the next phone occupies.

4.3. Syllabification

To obtain syllable constituents for Dutch, we use *syllabiphon*⁵. This Python package syllabifies phoneme sequences according to the Sonority Sequencing Principle (Selkirk, 1984), which states that—across different languages—syllables rise in sonority until a peak, after which the sonority falls. We first identify syllable boundaries and then split each syllable into onset (which is level or rising in sonority), nucleus (the sonority peak), and coda (the remainder of the syllable, which remains level or falls in sonority). For example, [ʔɔrdɛ] (*aarde* in the Aalsmeer NH dialect) is syllabified as:

[Syl(ons='ʔ', nuc='ɔ', cod='rdɛ'),
Syl(ons='d', nuc='e', cod='')]

Sonority scores for each phone are obtained from PanPhon and are on a scale from 1–9 (9 being the

⁵<https://github.com/dmort27/syllabiphon>

most sonorous) (Mortensen et al., 2016). Vowels have minimum sonority 8.

We postprocess the result of syllabiphon to merge diphthongs and triphthongs into the nuclei of their syllables. On a set of one randomly chosen word from 100 random Dutch dialects, syllabiphon, with the aforementioned diphthong/triphthong postprocessing, correctly syllabified 85% of the items upon manual inspection. Almost all of the errors concern consonant clusters involving [s] (e.g. [sp], [st], [sx]) being split into different constituents.⁶ ML-based syllabification models achieve higher performance on individual languages. Syllabiphon has the advantage, though, of performing reasonably well on IPA transcriptions of any language without supervision.

As for Min, the data is already implicitly segmented by syllable since each Han character (which comes with its own tone) is monosyllabic in Sinitic, so we use the tone numbers as syllable boundaries. To identify syllable constituents, we assign all phones with a [-cons] feature from PanPhon (Mortensen et al., 2016) (except for ʔ) to the nucleus. Phones with [+cons] (and ʔ) are assigned to the onset when no nucleus has been identified for the syllable yet. Additionally, contour tones are treated as one token. To summarize, 'ŋi31tʰœ51' is syllabified as [Syl('ŋ', 'i', ' ', '31'), Syl('tʰ', 'œ', ' ', '51')].

4.4. The Syllable Constituency Prediction Task

To incorporate syllable constituency into our model, we add the objective of predicting the syllable constituents \mathbf{c} of a word given its phonetic transcription \mathbf{x} :

$$p(\mathbf{c}|\mathbf{x}) = \prod_{i=1}^{|\mathbf{x}|} p(c_i|\mathbf{x}_{\leq i}) \quad (2)$$

In doing so, the model learns about the distribution of phonemes in different syllable constituents—a key aspect of phonotactics as described earlier (subsection 4.2). We first feed

⁶Since [s] has a lower sonority when it is in a word-initial or word-final cluster with a stop (e.g. [sp], [ps]), we manually set the sonority of [s] in these positions to be 0 for Dutch to ensure the cluster is placed into the same constituent. This changed the accuracy to 93%.

the phoneme sequence in, apply an embedding layer to each phone, and feed the embedding sequence into the LSTM. The output of the LSTM is used for both the language modeling task and the auxiliary task of syllable constituency prediction. Specifically, at each step, we predict both the next phone and the next phone’s syllable constituent (Figure 3).

Note that we only include the phone-level language modeling loss in the calculation of the correlation, so we model the same distribution as [Pimentel et al. \(2020\)](#), which is the classic language modeling objective (reproduced from their paper):

$$p(\mathbf{x}) = \prod_{i=1}^{|\mathbf{x}|} p(x_i | \mathbf{x}_{<i}) \quad (3)$$

We cannot feed the syllable constituent as input to the language modeling prediction task because the syllable constituent is determined using global information, which autoregressive language modeling should not have access to. Doing so would artificially increase the predictability of the phoneme (and thus deflate the phonotactic complexity). A longer sequence may be more predictable (and thus have lower complexity) because it may have more syllables and thus more global information.

4.5. Multi-task Loss

To jointly optimize for both phone and constituency prediction, a simple strategy would be to simply sum up the losses, which are each weighted by hyperparameters $\lambda_{a,b}$:

$$\mathcal{L}(\mathbf{x}, \mathbf{c}) = \lambda_a \mathcal{L}_{phon}(\mathbf{x}) + \lambda_b \mathcal{L}_{syl}(\mathbf{x}, \mathbf{c}) \quad (4)$$

where $\mathcal{L}_{phon}(\mathbf{x})$ is the cross entropy version of the likelihood $p(\mathbf{x})$, where $\mathcal{L}_{syl}(\mathbf{x}, \mathbf{c})$ is the cross entropy version of the likelihood $p(\mathbf{c} | \mathbf{x})$, and where, again, only $\mathcal{L}_{phon}(x)$ is used for the correlation.

However, [Kendall et al. \(2018\)](#) show that the performance of multi-task models is highly dependent on an appropriate choice of weighting between each task’s loss, which may be difficult to decide on with manual selection. We therefore follow [Kendall et al. \(2018\)](#) in opting for dynamically weighing the losses by assigning less weight to more *uncertain* tasks, where uncertainty is computed here with the variance of the task-specific losses, σ_t^2 :

$$\tilde{L}_t = \frac{1}{2\sigma_t^2} L_t + \log \sigma_t \quad (5)$$

where $\log \sigma_t$ here serves to avoid division by zero. Such adjustments have improved multitask learning in various NLP tasks ([Hofmann et al., 2023](#); [Hung et al., 2023](#)). In practice, given the instability

in regressing the variance σ_t^2 , we train the network to predict the log variance $\eta_t := \log \sigma_t^2$:

$$\tilde{L}_t = \frac{1}{2} (e^{-\eta_t} L_t + \eta_t) \quad (6)$$

The final loss is then the sum of the uncertainty-adjusted losses over all tasks.

4.6. Generalized Additive Models

Having obtained approximations of phonotactic complexity, we use generalized additive models (GAM) ([Wood, 2011, 2017](#)) to model the scores as a function of their geographical coordinates. GAMs can be considered as an extension of generalized linear models (GLM) ([McCullagh and Nelder, 1989](#)). In a GLM, the mean μ of a random variable Y is related to a weighted sum of linear predictors X with coefficients β (where $X\beta$ is denoted η) through a link function g :

$$g(\mu(Y)) = \eta = X\beta \quad (7)$$

In contrast to linear regression, the link function in the GLM provides the possibility of relating a linear model to an η of a distribution in the exponential family. GAMs extend this by fitting functions on the predictors:

$$\eta = b_0 + f(x_1) + f(x_2) \dots + f(x_p) \quad (8)$$

The functions are smoothing splines that are learned to fit to the predictors, and can take any number of arguments. The smoothing splines are themselves linear combinations of basis functions, which are geometrically simpler functions that combine by way of a weighted sum to model the data. The non-parametric nature of the smoothing splines allows for much greater flexibility in fitting the distribution of the variable, while the additivity of the model allows for understanding the contribution of each predictor, as the fitted variables do not depend on one another. In this paper, we use the implementation of GAMs offered in the *mgcv* package,⁷ with thin plate regression splines as smooths ([Wood, 2003](#)). Our statistical results are shown in [Table 6](#) and [Table 7](#).

5. Results and Discussion

5.1. Phonotactic Complexity and Word Length

As shown in [Table 2](#), we observe a moderate negative correlation across both Dutch and Min dialects. For the Dutch dataset, we observe a

⁷<https://cran.r-project.org/web/packages/mgcv/index.html>

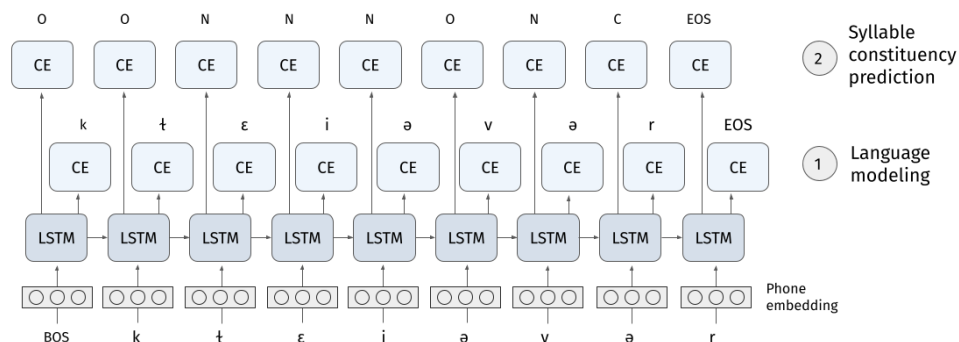


Figure 3: Diagram of our phonotactic language model with syllable constituency prediction as an auxiliary task, where O means onset; N, nucleus; C, coda; CE, cross-entropy loss. The example shown is the pronunciation of “klaver” in the Deurne NB dialect. The triphthong [εiə] was merged into the nucleus during postprocessing, but we treat diphthongs and triphthongs as sequences of multiple phonemes when calculating word length. The same LSTM is used for both language modeling and for syllable constituency prediction.

Pearson correlation of -0.68 both with and without syllabification, whereas for Min we observe a decrease from -0.72 to -0.7 after syllabification. The results show that incorporating the auxiliary task of syllable constituency does not strengthen the negative correlation for our datasets. This suggests that the LSTM-based phonotactic language model already implicitly learns syllable structure based on a phoneme’s context and position within a word. Despite modifying the setup with a syllable structure-based approach, we still reach a similar negative correlation, attesting to the robustness of the correlation. Furthermore, we also note that the tradeoff between linguistic complexities may in fact extend beyond such a binary model, where multiple factors may be at play in addition to phonotactic complexity and word length (Levshina, 2022). As such, the correlation should not be seen as an evaluation metric, since it is not known whether there is indeed a complete tradeoff between just the two dimensions.

5.2. GAMs

Table 3 and 4 summarize the statistical results of our GAMs on the Dutch dialect dataset, where for both phonotactic complexity and average word length, the coordinates as predictors are found to be statistically significant⁸. The regression surfaces are visualized in Figure 1 as contour plots mapped upon the coordinates of the Netherlands, where the color blue indicate that the values are low, yellow indicating that they are high, and green being a transitional value between the two. We also include statistics of the raw values for the Dutch dataset in Figure 4 for comparison.

⁸All GAM results including the plots are for non-syllabified runs.




Figure 4: Box plot of average word length (left) and phonotactic complexity (right) in Dutch dialects.

son. Values in both GAMs can be seen to show a high degree of similarity with nearby points, which corresponds with the fact that dialects which are spoken near one another tend to be more similar to one another than those that are spoken far away (Nerbonne, 2010). The larger blue region in the phonotactic complexity plot of Figure 1 is the region of Holland. For Min, high average word length and low phonotactic complexity is similarly concentrated around the capital of Fuzhou (Figure 2).

In the remaining sections, we focus our discussion on the Dutch dialect dataset due to a larger quantity of Dutch dialectology studies that examine this topic, although similar trends are observed also in the Min dataset.

5.3. The Compensation Hypothesis

Our results in Table 2 affirm that the compensatory relationship between phonotactic com-

Language	Syllable structure	Pearson r	Spearman ρ
Dutch	No	-0.678897	-0.63107
Dutch	Yes	-0.684041	-0.627137
Min	No	-0.720228	-0.692592
Min	Yes	-0.698437	-0.665027

Table 2: Correlation between phonotactic complexity and word length with and without syllable constituency prediction. The largest two-sided p-value among the Pearson correlations was $4.16e-10$.

Component	Term	Estimate	Std Error	t-value	p-value
parametric coefficients	Intercept	3.299	0.007	498.1	<0.001
Component	Term	edf	Ref. df	F-value	p-value
smooth terms	s(longitude, latitude)	25.4	28.22	18.12	<0.001
Adjusted R-squared					
0.581					

Table 3: Statistical results of GAMs on phonotactic complexity in Dutch dialects.

Component	Term	Estimate	Std Error	t-value	p-value
parametric coefficients	Intercept	4.805	0.009	524.4	<0.001
Component	Term	edf	Ref. df	F-value	p-value
smooth terms	s(longitude, latitude)	24.85	27.98	15.57	<0.001
Adjusted R-squared					
0.543					

Table 4: Statistical results of GAMs on average word length in Dutch dialects.

plexity and average word length observed in [Pimentel et al. \(2020\)](#) applies to the level of dialects as well. Dialects, however, provide us with a spatial dimension with social and historical correlates against which we can compare linguistic variables. We make use of this fact by modelling the phonotactic complexity and average word length of dialects as functions of their coordinates. In observing the inverse relationship between the distribution of high and low values of [Figure 1](#), we show how the negative correlation appears on a geographical level.

5.4. Koineization and the Linguistic Niche Hypothesis

In observing the center of low phonotactic complexity and high average word length to concentrate in the region of Holland, we interpret this as potential evidence of the linguistic niche hypothesis ([Lupyan and Dale, 2010](#); [Dale and Lupyan, 2012](#)), which states that linguistic complexity adapts to social constraints ([Trudgill, 2001](#); [McWhorter, 2007](#); [Bentz and Winter, 2014](#)). [Lupyan and Dale \(2010\)](#) conduct a study of over 2000 languages, and found morphological complexity to be negatively correlated with popula-

tion size, the degree of language contact, and the size of the region—the three of which are correlated—as languages spoken throughout regions with larger populations and broader geographical extents would also have a higher likelihood of having many non-native speakers. Extensive contact of this kind has been suggested to induce grammatical simplification, as non-native adult learners of the language would have difficulty in acquiring morphologically complex languages compared to a child.

As for dialects, Dutch dialectology has long utilized similar arguments, where the large influx of immigrants to early modern Holland is claimed to be the central cause of significant grammatical simplification in the urban vernacular of Holland ([Hendriks et al., 2018](#); [Howell, 2006](#)), the effect of which remains observable in its modern form ([Kloeke, 1927](#); [Hamans, 2011](#)). Holland in the 16th to 18th century saw a massive number of immigrants. Of 1.2 million immigrants, only over 200,000 were from within the province of Holland. The remainder came mostly from southern Netherlands and Germany ([Lucassen, 2002](#); [Hendriks et al., 2018](#)). This created a situation in which speakers of many different Dutch varieties intermixed (since natives of the province had become

a minority). Moreover, in an effort to integrate and achieve better socioeconomic outcomes, linguistic forms which carry strong regional characteristics were avoided; there was a bias towards phonologically and lexically simple forms due to the learning constraints of adult speakers. This led all parties to converge on simplified linguistic compromise. Children in such multi-dialectal contexts also have to make sense of the language used around them and, in the process, simplify the wide range of dialectal input in their environment. The combination of the aforementioned factors is claimed to have established a simpler variety, which became the native language of further generations (Howell, 2006; Kerswill and Williams, 2000).

In our results, Figure 1 is compatible with both the Linguistic Niche hypothesis and the Compensation Hypothesis—phonotactics was simplified through koineization, and word length increased concomitantly. As for Min, despite a relative lack of literature on linguistic complexity, the concentration of high average word length and low phonotactic complexity near the capital of Fuzhou may arguably also be due to similar processes (Figure 2), given the greater amount of socioeconomic opportunities offered.

Although Shcherbakova et al. (2023) find no inverse correlation between morphosyntactic complexity and sociodemographic factors such as the proportion of L2 speakers when examined against a database of 1314 languages (Skirgård et al., 2023), our study differs from theirs both in the dimensions of linguistic complexity measured, and also in the scale of linguistic variation, where our emphasis is placed on microvariation over geographically proximate dialects. As such, our results do not necessarily stand in conflict with theirs given the differences in measurement, and may point to the need of a sharper formulation of the Linguistic Niche hypothesis for more precise predictions of how linguistic complexity and sociodemographic factors affect one another.

6. Future Work

Although we have based this study on two languages of different language families to ensure the robustness of the results across typologically different languages, we seek to validate our results on additional dialect datasets, for which an interesting candidate would be the *Phonetischer Atlas von Deutschland* (Göschel, 2000),⁹ a German dataset containing 29,530 words transcribed in IPA across 183 cities in Germany.

We also seek to explore different methods of incorporating syllable structure beyond learning syl-

lable constituency as an auxiliary task. As Daland et al. (2011a) suggest, we can also explicitly supervise the model with sonority. Furthermore, despite Pimentel et al. (2020)'s finding that phonological features do not improve performance, we can explore more sophisticated methods of incorporating such features, as Torre (2003) finds that place of articulation is important in Dutch phonotactics. Finally, despite the lack of increase in negative correlation, we seek to perform a more thorough evaluation of our model enriched with syllable constituency knowledge on other downstream tasks.

Limitations

Our study is limited by the fact that the dataset we use comes in a form where diacritics are removed, which could potentially encode phonemic distinctions, thereby decreasing entropy. Additionally, our unidirectional LSTM limits our capability to capture regressive assimilation, which we hope to improve upon with more expressive architectures such as Transformers (Vaswani et al., 2017) in future work, though we acknowledge Mayer and Nelson (2020)'s suggestion that Elman RNNs may suffice to model most phonotactics. We also acknowledge that some phones may be inherently more complex than others due to differences in articulatory difficulty (Maddieson, 2009), which our model does not capture. Finally, we reduce the factors behind linguistic complexity to a binary tradeoff between two dimensions, but Levshina (2022) suggest the situation is actually multivariate. Even with a simplified model as ours, though, a moderate correlation of around 0.7 can be attained, lending credibility to the compensation hypothesis. Future work may examine other factors that play a role in the trade-off.

7. Bibliographical References

- Adam Albright. 2009. Feature-based generalisation as a source of gradient acceptability. *Phonology*, 26(1):9–41.
- R. Harald Baayen and Maja Linke. 2020. *Generalized Additive Mixed Models*, pages 563–591. Springer International Publishing, Cham.
- Christian Bentz, Ximena Gutierrez-Vasques, Olga Sozinova, and Tanja Samardžić. 2022. Complexity trade-offs and equi-complexity in natural languages: a meta-analysis. *Linguistics Vanguard*.

⁹<https://github.com/cysouw/PAD>

- Christian Bentz and Bodo Winter. 2014. *Languages with More Second Language Learners Tend to Lose Nominal Case*, pages 96 – 124. Brill, Leiden, The Netherlands.
- Juliette Blevins. 1995. The syllable in phonological theory. In John A. Goldsmith, editor, *The handbook of phonological theory*, chapter 6, pages 206–244. Blackwell, Cambridge, Mass., USA.
- Mary Bucholtz and Kira Hall. 2005. Identity and interaction: A sociocultural linguistic approach. *Discourse studies*, 7(4-5):585–614.
- Si-han Chen, Richard Futrell, and Kyle Mahowald. 2022. *Investigating information-theoretic properties of the typology of spatial demonstratives*. In *Proceedings of the 4th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 94–95, Seattle, Washington. Association for Computational Linguistics.
- Noam Chomsky and Morris Halle. 1968. The sound pattern of English.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. *What you can cram into a single \$&!#* vector: Probing sentence embeddings for linguistic properties*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Robert Daland, Bruce Hayes, James White, Marc Garellek, Andrea Davis, and Ingrid Norrmann. 2011a. Explaining sonority projection effects. *Phonology*, 28(2):197–234.
- Robert Daland, Bruce Hayes, James White, Marc Garellek, Andrea Davis, and Ingrid Norrmann. 2011b. Explaining sonority projection effects. *Phonology*, 28(2):197–234.
- Rick Dale and Gary Lupyan. 2012. *Understanding the origins of morphological diversity: The linguistic niche hypothesis*. *Advances in Complex Systems*, 15(03n04):1150017.
- David Gil. 1994. *The structure of Riau Indonesian*. *Nordic Journal of Linguistics*, 17(2):179–200.
- David Gil. 2008. *How complex are isolating languages? Language Complexity: Typology, contact, change*, 109-131 (2008).
- Camiel Hamans. 2011. *An early sociolinguistic approach towards standardization and dialect variation*, pages 369–387.
- Bruce Hayes. 2008. *Introductory phonology*, volume 7. John Wiley & Sons.
- Bruce Hayes and Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic inquiry*, 39(3):379–440.
- Wilbert Heeringa and John Nerbonne. 2001. *Dialect areas and dialect continua*. *Language Variation and Change*, 13(3):375–400.
- Jennifer Hendriks, Todd Ehresmann, Robert B. Howell, and Madison Mike Olson. 2018. *Migration and linguistic change in early modern holland: The case of leiden*. *Neuphilologische Mitteilungen*, 119(1):145–172.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Charles F. Hockett. 1958. *A course in modern linguistics*. Oxford & IBH, New Delhi.
- Valentin Hofmann, Goran Glavaš, Nikola Ljubešić, Janet B. Pierrehumbert, and Hinrich Schütze. 2023. *Geographic adaptation of pretrained language models*.
- Robert Howell. 2006. *Immigration and koineisation: The formation of early modern dutch urban vernaculars*. *Transactions of the Philological Society*, 104:207 – 227.
- Chia-Chien Hung, Anne Lauscher, Dirk Hovy, Simone Paolo Ponzetto, and Goran Glavaš. 2023. *Can demographic factors improve text classification? revisiting demographic adaptation in the age of transformers*.
- John Joseph and Frederick Newmeyer. 2012. ‘all languages are equally complex’: The rise and fall of a consensus. *Historiographia Linguistica*, 39.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. 2018. *Multi-task learning using uncertainty to weigh losses for scene geometry and semantics*.
- Paul Kerswill and Ann Williams. 2000. *Creating a new town koine: Children and language change in milton keynes*. *Language in Society*, 29(1):65–115.
- Diederik P. Kingma and Jimmy Ba. 2015. *Adam: A method for stochastic optimization*. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- James Kirby. 2021a. *Incorporating tone in the calculation of phonotactic probability*. In *Proceedings of the 18th SIGMORPHON Workshop on*

- Computational Research in Phonetics, Phonology, and Morphology, pages 32–38, Online. Association for Computational Linguistics.
- James Kirby. 2021b. [Incorporating tone in the calculation of phonotactic probability](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 32–38, Online. Association for Computational Linguistics.
- Gesinus Gerhardus Kloeke. 1927. *de Hollandsche expansie in de zestiende en zeventiende eeuw en haar weerspiegeling in de hedendaagsche Nederlandsche dialecten*. Martinus Nijhoff.
- Juliette Levin. 1985. *A metrical theory of syllabicity*. Ph.D. thesis, Massachusetts Institute of Technology.
- Natalia Levshina. 2022. *Communicative Efficiency: Language Structure and Use*. Cambridge University Press.
- Chinfa Lien. 2000. Denasalization, vocalic nasalization and related issues in southern min: A dialectal and comparative perspective. In *memory of Professor Li Fang-Kuei: Essays of linguistic change and the Chinese dialects*. Academic Sinica, Taipei.
- Chinfa Lien. 2015. [160Min languages](#). In *The Oxford Handbook of Chinese Linguistics*. Oxford University Press.
- Jan Lucassen. 2002. *Holland, een open gewest. Immigratie en bevolkingsontwikkeling*, pages 181–215. Verloren, Hilversum.
- Gary Lupyan and Rick Dale. 2010. [Language structure is partly determined by social structure](#). *PLOS ONE*, 5(1):1–10.
- Ian Maddieson. 2009. [Calculating phonological complexity](#), pages 83–110. De Gruyter Mouton, Berlin, New York.
- Connor Mayer and Max Nelson. 2020. [Phonotactic learning with neural language models](#). In *Proceedings of the Society for Computation in Linguistics 2020*, pages 291–301, New York, New York. Association for Computational Linguistics.
- P. McCullagh and J.A. Nelder. 1989. *Generalized Linear Models, Second Edition*. Chapman and Hall/CRC Monographs on Statistics and Applied Probability Series. Chapman & Hall.
- John McWhorter. 2001. [The world's simplest grammars are creole grammars](#). *Linguistic Typology*, 5:125–166.
- John McWhorter. 2007. *Language Interrupted: Signs of Non-Native Acquisition in Standard Language Grammars*. Oxford University Press.
- Steven Moran and Damián Blasi. 2014. [Cross-linguistic comparison of complexity measures in phonological systems](#). In *Measuring Grammatical Complexity*. Oxford University Press.
- David R. Mortensen, Patrick Littell, Akash Bharadwaj, Kartik Goyal, Chris Dyer, and Lori Levin. 2016. [PanPhon: A resource for mapping IPA segments to articulatory feature vectors](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3475–3484, Osaka, Japan. The COLING 2016 Organizing Committee.
- John Nerbonne. 2010. [Measuring the diffusion of linguistic change](#). *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 365(1559):3821–3828.
- Frederick J. Newmeyer. 2021. [Complexity and relative complexity in generative grammar](#). *Frontiers in Communication*, 6.
- Tiago Pimentel, Brian Roark, and Ryan Cotterell. 2020. [Phonotactic complexity and its trade-offs](#). *Transactions of the Association for Computational Linguistics*, 8:1–18.
- Michael Ramscar, Daniel Yarlett, Melody Dye, Katie Denny, and Kirsten Thorpe. 2010. [The effects of feature-label-order and their implications for symbolic learning](#). *Cognitive Science*, 34(6):909–957.
- Elisabeth Selkirk. 1984. On the major class features and syllable theory. *Language Sound Structure*, pages 107–136.
- Olena Shcherbakova, Susanne Maria Michaelis, Hannah J Haynie, Sam Passmore, Volker Gast, Russell D Gray, Simon J Greenhill, Damián E Blasi, and Hedvig Skirgård. 2023. Societies of strangers do not speak less complex languages. *Science Advances*, 9(33):eadf7704.
- Julius Steuer, Johann-Mattis List, Badr M. Abdullah, and Dietrich Klakow. 2023. [Information-theoretic characterization of vowel harmony: A cross-linguistic study on word lists](#). In *Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 96–109, Dubrovnik, Croatia. Association for Computational Linguistics.
- Johan Taeldeman and Frans Hinskens. 2014. *7. The classification of the dialects of Dutch*, pages 129–142. De Gruyter Mouton, Berlin, Boston.

- Erik Jan van der Torre. 2003. Dutch sonorants: The role of place of articulation in phonotactics. *Utrecht: LOT*.
- Peter Trudgill. 1989. *Contact and isolation in linguistic change*, pages 227–238. De Gruyter Mouton, Berlin, New York.
- Peter Trudgill. 2001. Contact and simplification: Historical baggage and directionality in linguistic change. *Linguistic Typology*, 5(2/3):371–374.
- Erik Jan Van der Torre. 2003. *Dutch sonorants: the role of place of articulation in phonotactics*, volume 81. LOT.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Soren Wichmann and Jeff Good. 2014. *Quantifying Language Dynamics: On the Cutting edge of Areal and Phylogenetic Linguistics*. Brill, Leiden, The Netherlands.
- Søren Wichmann and Eric Holman. 2023. *Cross-linguistic conditions on word length*. *PloS one*, 18:e0281041.
- Martijn Wieling, Wilbert Heeringa, and John Nerbonne. 2007. An aggregate analysis of pronunciation in the Goeman-Taeldeman-Van Reenen-Project data. *Taal en Tongval*, 59(1):84–116.
- Martijn Wieling, John Nerbonne, and R. Harald Baayen. 2011. *Quantitative social dialectology: Explaining linguistic variation geographically and socially*. *PLOS ONE*, 6(9):1–14.
- S. N. Wood. 2003. Thin-plate regression splines. *Journal of the Royal Statistical Society (B)*, 65(1):95–114.
- S. N. Wood. 2011. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)*, 73(1):3–36.
- S.N Wood. 2017. *Generalized Additive Models: An Introduction with R*, 2 edition. Chapman and Hall/CRC.
- Dellert, Johannes and Daneyko, Thora and Münch, Alla and Ladygina, Alina and Buch, Armin and Clarius, Natalie and Grigorjew, Ilja and Balabel, Mohamed and Boga, Hizniye Isabella and Baysarova, Zalina and others. 2020. *NorthEuraLex: A wide-coverage lexical database of Northern Eurasia*. Springer.
- Göschel, Joachim. 2000. *Der phonetische Atlas von Deutschland*. [\[link\]](#).
- Skirgård, Hedvig and Haynie, Hannah J. and Blasi, Damián E. and Hammarström, Harald and Collins, Jeremy and Latache, Jay J. and Lesage, Jakob and Weber, Tobias and Witzlack-Makarevich, Alena and Passmore, Sam and Chira, Angela and Maurits, Luke and Dinnage, Russell and Dunn, Michael and Reesink, Ger and Singer, Ruth and Bower, Claire and Epps, Patience and Hill, Jane and Vesakoski, Outi and Robbeets, Martine and Abbas, Noor Karolin and Auer, Daniel and Bakker, Nancy A. and Barbos, Giulia and Borges, Robert D. and Danielsen, Swintha and Dorenbusch, Luise and Dorn, Ella and Elliott, John and Falcone, Gida and Fischer, Jana and Ghanggo Ate, Yustinus and Gibson, Hannah and Göbel, Hans-Philipp and Goodall, Jemima A. and Gruner, Victoria and Harvey, Andrew and Hayes, Rebekah and Heer, Leonard and Herrera Miranda, Roberto E. and Hübler, Natalia and Huntington-Rainey, Biu and Ivani, Jessica K. and Johns, Marilen and Just, Erika and Kashima, Eri and Kipf, Carolina and Klingenberg, Janina V. and König, Nikita and Koti, Aikaterina and Kowalik, Richard G. A. and Krasnoukhova, Olga and Lindvall, Nora L.M. and Lorenzen, Mandy and Lutzenberger, Hannah and Martins, Tônia R.A. and Mata German, Celia and van der Meer, Suzanne and Montoya Samamé, Jaime and Müller, Michael and Muradoglu, Saliha and Neely, Kelsey and Nickel, Johanna and Norvik, Miina and Oluoch, Cheryl Akinyi and Peacock, Jesse and Pearey, India O.C. and Peck, Naomi and Petit, Stephanie and Pieper, Sören and Poblete, Mariana and Prestipino, Daniel and Raabe, Linda and Raja, Amna and Reimringer, Janis and Rey, Sydney C. and Rizaew, Julia and Ruppert, Eloisa and Salmon, Kim K. and Sammet, Jill and Schembri, Rhiannon and Schlabach, Lars and Schmidt, Frederick W.P. and Skilton, Amalia and Smith, Wikaliler Daniel and de Sousa, Hilário and Sverredal, Kristin and Valle, Daniel and Vera, Javier and Voß, Judith and Witte, Tim and Wu, Henry and Yam, Stephanie and Ye 田田, Jingting and Yong, Maisie and Yuditha, Tessa and Zariquiey, Roberto and Forkel, Robert and Evans, Nicholas and Levinson, Stephen C. and

8. Language Resource References

- Centre for the Protection of Language Resources of China. 2023. *The Chinese Language Resources Protection Project Collection and Display Platform*. [\[link\]](#).

Haspelmath, Martin and Greenhill, Simon J. and Atkinson, Quentin D. and Gray, Russell D. 2023. *Grambank Reveals Global Patterns in the Structural Diversity of the World’s Languages*. [\[link\]](#).

Taeldeman, Johan and Goeman, A. 1996. *Fonologie en morfologie van de Nederlandse dialecten: een nieuwe materiaalverzameling en twee nieuwe atlasprojecten*.

A. Syllabification

Algorithm 1 Algorithm for finding syllable boundaries given a list of sonority scores

```

1: function FindBound(s: list of sonority scores)
2:    $P \leftarrow []$ 
3:   for each  $i \in 0 \dots |s| - 1$  do
4:     if  $s[i] \geq \alpha$  then
5:       push( $P, i$ )
6:    $b \leftarrow [\text{True} \mid x \in s]$ 
7:   push( $b, \text{True}$ )
8:   for each  $i \in 0 \dots \text{length}(b[i : -1])$  do
9:     if  $i < P[0]$  or  $i > P[-1]$  then
10:       $b[i] \leftarrow \text{False}$ 
11:     if  $s[i - 1] < s[i]$  then
12:       $b[i] \leftarrow \text{False}$ 
13:     else if  $i = \text{length}(s) - 1$  then
14:       $b[i] \leftarrow \text{False}$ 
15:     if  $i > 2$  and  $s[i - 2] = s[i - 1] = s[i]$  then
16:       $b[i] \leftarrow \text{False}$ 
17:   return  $b$ 

```

Algorithm 2 Algorithm for parsing syllables into onset, nucleus, and coda

```

1: function ParseSyl(s: list of sonority scores,  $p$ : list of phonemes)
2:    $m \leftarrow \max(s)$ 
3:    $i \leftarrow 0$ 
4:    $o \leftarrow \varepsilon$ 
5:   while  $s < m$  do
6:      $o \leftarrow o \cup p[i]$ 
7:      $i \leftarrow i + 1$ 
8:    $n \leftarrow p[i]$ 
9:    $i \leftarrow i + 1$ 
10:   $c \leftarrow \text{join}([x \mid x \in p \dots p[i : -1]])$ 
11:  return  $\langle o, n, c \rangle$ 

```

B. Reproducibility

B.1. Hyperparameters

We use the same hyperparameters for each language and for the models with and without syllable structure. The hyperparameter values are

Hyperparameter	Value
Learning rate	1e-3
Adam betas	(0.9, 0.999)
Adam eps	1e-8
Dropout	0.2
Phoneme embedding size	256
Syllable constituent embedding size	256
Number of LSTM layers	1
Hidden size	256
Batch size	64

Table 5: Hyperparameter values for all languages and all models

the default ones provided by Pimentel et al. (2020). Refer to Table 5 for the values. Note that we use the Adam optimizer (Kingma and Ba, 2015). Following Pimentel et al. (2020), we use a 80-10-10 train-dev-test split.

B.2. Compute resources

Our model has 578,147 parameters and was trained on one GPU on a Ubuntu server. On GPU, it takes less than half an hour to train models for all 366 dialects.

B.3. GAM results for Min

Component	Term	Estimate	Std Error	t-value	p-value
parametric coefficients	Intercept	2.653	0.011	237.6	<0.001
Component	Term	edf	Ref. df	F-value	p-value
smooth terms	s(longitude, latitude)	12.34	16.75	6.123	<0.001
Adjusted R-squared					
0.638					

Table 6: Statistical results of GAMs on phonotactic complexity in Min dialects.

Component	Term	Estimate	Std Error	t-value	p-value
parametric coefficients	Intercept	6.888	0.026	261	<0.001
Component	Term	edf	Ref. df	F-value	p-value
smooth terms	s(longitude, latitude)	19.34	24.21	4.749	<0.001
Adjusted R-squared					
0.658					

Table 7: Statistical results of GAMs on average word length in Min dialects.




Figure 5: Contour plot of GAM results using a thin plate regression spline, where the response variable is non-syllabified phonotactic complexity of Min dialects, and the predictors, their latitude and longitude. Colder colors (blue) indicate low complexity, while warmer colors (yellow) indicate higher values.




Figure 6: Contour plot of GAM results using a thin plate regression spline, where the response variable is average word length of Min dialects, and the predictors, their latitude and longitude. Colder colors (blue) indicate shorter length, while warmer colors (yellow) indicate longer.

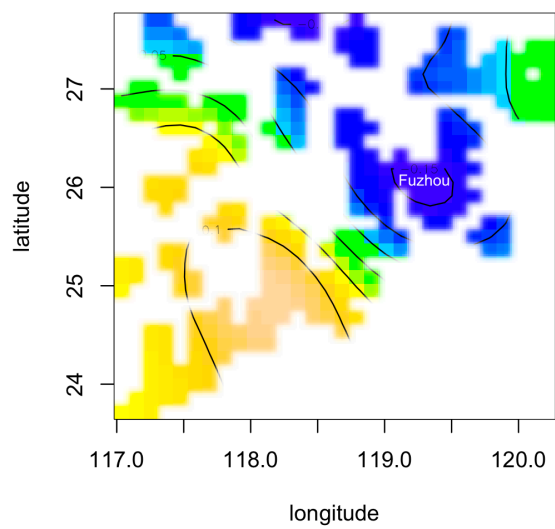


Figure 7: Contour plot of GAM results using a thin plate regression spline, where the response variable is syllabified phonotactic complexity of Min dialects, and the predictors, their latitude and longitude. Colder colors (blue) indicate low complexity, while warmer colors (yellow) indicate higher values.