



Modelled Multivariate Overlap: A method for measuring vowel merger

Irene Smith¹, Morgan Sonderegger¹, The Spade Consortium

¹McGill University, Canada

irene.smith@mail.mcgill.ca, morgan.sonderegger@mcgill.ca

Abstract

This paper introduces a novel method for quantifying vowel overlap. There is a tension in previous work between using multivariate measures, such as those derived from empirical distributions, and the ability to control for unbalanced data and extraneous factors, as is possible when using fitted model parameters. The method presented here resolves this tension by jointly modelling all acoustic dimensions of interest and by simulating distributions from the model to compute a measure of vowel overlap. An additional benefit of this method is that computation of uncertainty becomes straightforward. We evaluate this method on corpus speech data targeting the PIN-PEN merger in four dialects of English and find that using modelled distributions to calculate Bhattacharyya affinity substantially improves results compared to empirical distributions, while the difference between multivariate and univariate modelling is subtle.

Index Terms: vowel merger, distributional similarity, Bayesian modelling, corpus phonetics, sociolinguistics

1. Introduction

Vowels that are close together in articulation can have overlapping distributions in acoustic space (e.g., F1 and F2). It is often useful to have a measure of how much their distributions overlap, for research in areas including phonetics, sociolinguistics, and L2 acquisition [1]. We focus on its application for the study of *vowel merger*, where two distinct vowel categories come to be pronounced identically, a topic that is well-studied in the sociolinguistics and phonetics literature. Merger can be *unconditioned*, where the two categories lose their contrast entirely (such as the COT-CAUGHT merger), or *allophonic*, where the contrast is lost only in a specific phonological context, like the PIN-PEN merger, where the contrast is lost only prenasally [2].

2. Measures of vowel overlap

Sociolinguistic studies have employed different overlap measures to quantify the degree of vowel merger. Such gradient, acoustically-derived measures are especially useful when studying partial merger and merger- (or un-merger) in-progress.

Nycz and Hall-Lew [3] survey approaches to quantifying vowel merger in F1/F2 space: of interest here are Euclidean distance [4, 5, 6], adjusted Euclidean distance from mixed-effects regression modelling [7, 8], and Pillai scores [9]. A commonly-used alternative to Pillai scores is Bhattacharyya affinity [10], discussed further below. Kelly and Tucker [1] assess Pillai scores alongside three other measures, including the *a posteriori* probability-based metric (APP: [11]). Other methods exist (e.g. [12, 13]) but are not considered here.

Nycz and Hall-Lew define several desiderata for any measure of vowel merger, which we adopt for our proposed measure of overlap. (1) *Multivariate* measures are preferred over *univariate* measures. (2) Measures should control for unbalanced data, which is common in naturalistic data, to avoid undue influence from frequent lexical items and to control for phonological context. (3) There should be a way to assess uncertainty associated with the measure.

Euclidean distance between two categories is defined as the distance between the centroids of each category in F1/F2 space, and has the advantage of interpretability. It only satisfies desideratum (1). Adjusted Euclidean distance uses predictions from mixed-effects regression modelling to compute Euclidean distance between two classes from two univariate models (predicting F1 and F2). This allows (2) and (3) to be met, but not (1).

Pillai scores, a test statistic from a Multivariate ANOVA, indicate the variance in F1 and F2 accounted for by vowel class [1]. While satisfying all three desiderata, Pillai scores have limitations: they allow some integration of control predictors, but random effects cannot be included in the R implementation [3], a potential barrier for many researchers. Furthermore, Pillai scores and their associated *p*-values are sensitive to sample size [14] and assume normally distributed categories with equal variance, which is often not the case [10]. We do not consider Pillai scores further.

Bhattacharyya affinity (BA) is a promising alternative to Pillai scores which avoids these assumptions [10]. BA gives an interpretable number for any two arbitrary distributions, and only indicates “perfect overlap” when the two distributions are identical. Like Euclidean distance, BA calculated from raw distributions satisfies (1) but not (2) or (3), but it has a theoretical advantage over Euclidean distance in that it accounts for the *shape* of category distributions.

APP is based on the rate of misclassification between two categories. It is compatible with our proposed method, but is not considered further due to space considerations.

The measures discussed above show a pattern: either the measure is calculated directly from the distribution, yielding multivariate but potentially skewed results due to unbalanced data, or the measure is modelled with appropriate controls but is univariate. Moreover, none produce satisfying measures of uncertainty [3, 10, 1]. The method presented below satisfies all these desiderata, while maintaining the flexibility to calculate any measure.

[3, 10] distinguish between measures of *distance* (difference in a measure of central tendency between categories) and *overlap* (amount of shared acoustic space). In Section 4, we use BA, a measure of overlap; equivalent Euclidean distance plots are available in the Github repository (linked below).

3. MMO: Modelled Multivariate Overlap

This section describes the Modelled Multivariate Overlap method in general terms. The method, as described, is designed for vowel overlap in formant space (here, F1 and F2), but it could in principle be applied to other types of phones using any set of acoustic dimensions.

We use Bayesian linear mixed-effects models [15] to implement the proposed method. Bayesian modelling is particularly convenient because it directly outputs modelled distributions for both the data and for parameter values. In principle, this approach could also be taken using frequentist models (e.g. linear mixed-effects models) by simulating from fitted models [16]. The code for fitting the models in R using brms [17] and Stan [18], as well as to calculate MMO, is available as a Github repository.¹




Figure 1: *Illustration of the steps involved in MMO. Vowel category (V1 or V2) is indicated by linetype (solid or dashed); context (con. 1 or con. 2) is indicated with color; and (normalized) formant (F1 or F2) is indicated with shape (circle or triangle).*

Calculating MMO involves four steps, illustrated in Figure 1:

1. Normalize vowel tokens so that F1 and F2 are comparable across speakers to control for anatomical differences but preserve linguistically relevant variation [19]. Restrict the dataset to tokens of the two vowels of interest. Tokens can be filtered by other criteria, e.g., syllable stress. Normalization may be done before or after filtering tokens, as appropriate.
2. Jointly model F1 and F2 as a function of the predictors of interest, as well as any controls.
3. Simulate the joint distribution of F1 and F2 for each combination of predictors of interest, marginalizing over or adjusting for each control predictor as appropriate. For example, to calculate vowel overlap per speaker, speaking rate could be held constant, while speaker gender would be factored into predictions.
4. Compute the desired vowel merger measure from the simulated distributions. The measure should be computed separately for each salient context (e.g., for the conditioning environment and the “elsewhere” context for conditioned mergers). Do this for repeated simulations to generate a *distribution* of possible merger values for estimating uncertainty.

While MMO is agnostic to the overlap measure calculated from it, it is important to consider the assumptions of the statistical model relative to the assumptions of this measure. For example, Pillai scores should only be used if the model is structured such that the simulated distributions are (roughly) normally distributed with equal variance.

Finally, the model structure must be suited to answering the research questions. For example, if vowel merger is of interest,

¹https://github.com/MontrealCorpusTools/MMO_vowel_overlap

there must be a predictor of *vowel* category. If the merger of interest is conditioned, an additional *context* predictor and its interaction with *vowel* must be included so that the distance between vowel categories is allowed to vary by context. In what follows, we refer to the model that has only the predictors necessary to answer research questions as the *minimal model*, and any model that goes beyond these minimal parameters by adding controls, modelling variance, etc. as an *expanded model*.

4. Example application: PIN-PEN merger

This section demonstrates an application of MMO. We use data targeting the PIN-PEN merger, which is a conditioned merger, that has been described as /ɪ/ and /ɛ/ becoming merged before nasal consonants. Southern US English and African American English have been described as having this merger [2].

We examine the overlap between /ɪ/ and /ɛ/ both prenasally and preorally in four dialects of English. The data used here were assembled as part of the SPADE project [20], and the four corpora used here were chosen to represent different expectations for the behavior of /ɪ/ and /ɛ/, both prenasally and preorally. The four dialects chosen for this study were the US South, non-Southern North American, Southern England, and Scottish. The expectations for each dialect with respect to PIN-PEN are outlined below. All of the data was extracted from spontaneous speech produced in sociolinguistic interviews, and F1 and F2 were measured at 1/3 of the duration of each vowel token using the algorithm from [21] implemented in PolyglotDB [22].

US South is the dialect that is expected to have the PIN-PEN merger, and therefore is predicted to have full overlap in prenasal contexts and less overlap in preoral contexts. North American English has been described as having substantial extent and magnitude of prenasal coarticulation, that varies substantially across speakers [23]. Vowel nasality has the effect of compressing the F1 space compared to the oral vowel space, meaning that /ɪ/ and /ɛ/ are expected to be closer together in nasalized vowels than in oral vowels. The predicted consequence is that /ɪ/ and /ɛ/ should be slightly more similar prenasally than preorally for North American speakers, with possible variation due to the variation in prenasal coarticulation. Southern England is expected to have relatively little prenasal coarticulation [24], and therefore relatively little change in vowel overlap prenasally compared to preorally, especially given that the formants were measured 1/3 into each vowel’s articulation. Finally, Scottish English has been described as having only a marginal contrast between /ɪ/ and /ɛ/ [25], with no height (F1) difference. Thus, Scottish is expected to have high overlap both prenasally and preorally, and context is expected to have little-to-no effect on overlap.

4.1. Methods

4.1.1. Preprocessing

F1 and F2 were measured at 1/3 of the vowel duration and Lobanov normalized [26] (i.e., z-scored by speaker using tokens of every vowel). The data was then subsetted to include only stressed tokens of /ɪ/ and /ɛ/. The set of words chosen was constant across all four dialects, and was designed to include only stressed vowels in the KIT and DRESS lexical sets [27], using the UNISYN dictionary [28].

4.1.2. Modelling

MMO was implemented separately for each dialect, with several variations in implementation for comparison. We compare two different model structures: the “minimal” model structure described above, and an extended model incorporating more controls.

The minimal model models F1 and F2 with fixed effects of *vowel* ($/i/$ or $/\varepsilon/$), *context* (nasal or oral), and their interaction, random intercepts of *speaker* and *word*, and by-*speaker* random slopes of *vowel*, *context*, and their interaction. This model structure allows the degree of overlap calculated from the model to differ by context for each speaker.

The expanded model includes all of the terms in the minimal model. Additionally, it includes a fixed effect of $\log(\text{duration})$ and its interaction with *vowel* and *context* (to capture any differing effects of duration on the *vowel*-by-*context* interaction), random intercepts of *following consonant* nested above the by-*word* random intercepts (one of several ways to control for phonological context). The variance of both the fixed and by-*speaker* random effects of *vowel*, *context*, and their interaction was modelled [17, 29], meaning that each *vowel*-by-*context*-by-*speaker* pair was allowed to have a differently-shaped distribution.

To assess the importance of multivariate modelling, both a multivariate and two univariate models (one for F1 and one for F2) were fit for each of the two model structures described above. The two univariate models were then deployed together to simulate an (uncorrelated) two-dimensional distribution.

All models were fit using brms [17] in R. Brms’s default (flat or weakly informative) priors were used in all cases except for correlations, which had LKJ(1.5) priors, to discourage (near)-perfect random effect correlations [30].

4.1.3. Simulation and merger metrics

The distribution for each vowel-by-context pair was simulated from each of the models defined above. Both by-*speaker* and “average speaker” distributions were simulated: “average speaker” predictions come from ignoring the entire random effects structure (including by-*word* random effects).

Two measures of vowel overlap were calculated. BA was calculated from 1000 draws from the posterior predicted distribution, 100 times (to estimate uncertainty). Euclidean distance was estimated 100 times from a single draw from the posterior of the estimate for the mean. The results for BA are shown below, and the results for Euclidean distance are available in supplementary materials (on Github).

Additionally, measures of vowel merger were calculated for each speaker from the raw empirical distributions to compare MMO to current standard methods. In order to make a passing attempt at balancing the data without modelling, the raw distributions were also averaged by word (to minimize the influence of highly frequent words) before calculating the vowel merger metrics for each speaker. Errorbars are not shown for the two empirical approaches, since it is non-trivial to define uncertainty given by-*speaker* and word variability.

Table 1 summarizes all of the different types of distributions used to calculate our two metrics of vowel merger.

4.2. Results

Most plots shown in this section focus on the empirical and multivariate results. The univariate results are almost identical, and thus were excluded for space. Only Figure 4 shows results for

Table 1: Types of distributions used to calculate vowel overlap.

	empirical	modelled univariate	modelled multivariate
raw	•		
averaged	•		
minimal		•	•
expanded		•	•

all models and empirical distributions to demonstrate the similarity between multivariate and univariate results. Plots showing all the results are available in supplementary materials (on Github).

Figure 2 shows the raw (unaveraged) empirical 10% quantile ellipses (top row) and simulated distributions from the two multivariate models for the average speaker of each dialect (second and third rows). 10% intervals are used in order to better illustrate the relative location of categories. The predicted distributions are qualitatively similar across the different methods for each dialect, but also show differences, which will propagate into all measures computed from the distributions.




Figure 2: 10% quantile ellipses of empirical distributions and multivariate modelled distributions.

Figure 3 compares overlap across dialects, shown for each method. The two empirically-derived measures (top row) are computed by averaging the BA of all the speakers in each dialect. The four methods shown give qualitatively the same pattern. US South has high prenasal overlap and lower preoral overlap, and Scotland has high overlap both prenasally and preorally. For these two dialects, the modelled results have much higher overlap estimates compared to the empirically-derived estimates, making the modelled overlap more in line with the predictions from Section 4. North America and Southern England have relatively low overlap both prenasally and preorally, with possibly higher prenasal than preoral overlap, as expected.

We evaluate each method (i.e., row in Figure 4) on the basis of its compatibility with what we already know about each dialect. For this type of evaluation heuristic, it is particularly use-




Figure 3: Predicted Bhattacharyya affinity calculated from empirical distributions (top row) and modelled multivariate distributions (bottom row).

ful to look at patterns *across* speakers. Figure 4 shows BA estimates for each speaker from each method. Each black dot represents the prenasal (y-axis) and preoral (x-axis) overlap for a single speaker. A speaker whose dot falls directly on the dashed $y = x$ line has the same prenasal and preoral overlap, and speakers falling above the line have higher overlap prenasally than preorally (a BA of 1 means that the two phones have the same distribution).

In Figure 4, there is almost no difference between the multivariate and univariate modelled results, hence the choice to exclude the univariate models from other figures. Additionally, the difference between the minimal model and expanded model is subtle. The most striking result from Figure 4 is the contrast between the modelled estimates and the empirically-derived ones. All models show speakers falling into the predicted patterns to a much larger degree than the empirical distributions do. Note that there is no structural reason why the models would be expected to predict speakers falling into a particular pattern. Additionally, the mismatch between the modelled and empirical results should not be, *a priori*, alarming, since the models factor in multiple control variables that surface as “noise” in the empirical distributions.

Compared to the “average speaker” predictions (Figure 3), the by-speaker plot shows a clear difference in behavior between North America and Southern England: both dialects show an effect of prenasal coarticulation in the expected direction, but North America has much more interspeaker variation in the amount of prenasal overlap relative to preoral overlap than Southern England, in line with high interspeaker variability in prenasal coarticulation in American English [23].

5. Discussion

We found that computing measures from modelled data was the single most important choice affecting the quality of overlap measures. Word-averaging the empirical distributions was implemented as a simple data-balancing measure, but it did not produce substantially better results than calculating directly from the raw empirical distributions. We thus strongly recommend some modelling steps before calculating overlap.

While the plots in Figure 4 show clear qualitative differ-

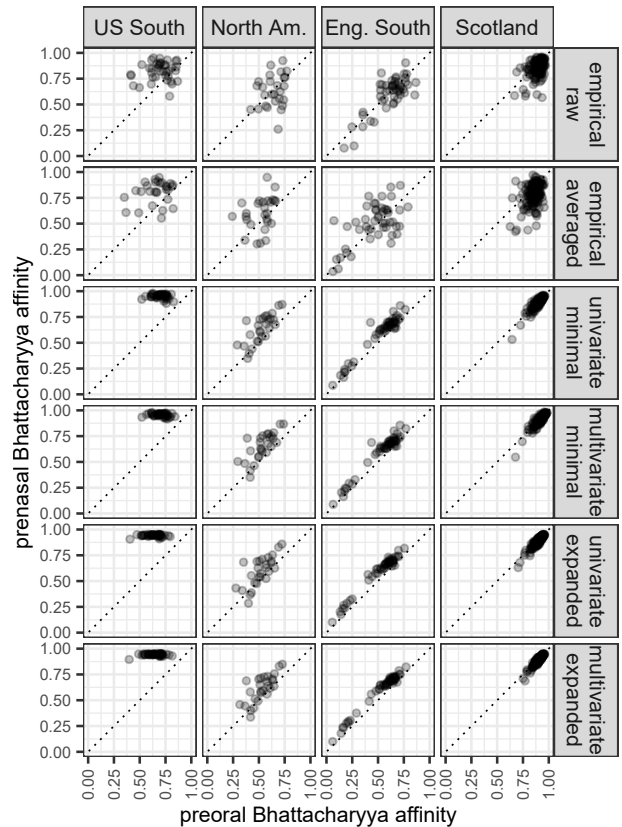


Figure 4: Predicted Bhattacharyya affinity by speaker, for each dialect and modelling approach.

ences between the modelled and empirical results, we did not have a quantitative metric evaluating different methods. One possibility for future work would be simulating data where the underlying distributional properties are known (similar to [14]).

In the case of the PIN-PEN merger, multivariate modelling shows no discernible improvement over univariate modelling. However, multivariate modelling is still theoretically preferable, and might turn out to be more important in other mergers with substantial between-category differences in both F1 and F2 (e.g., the COT-CAUGHT merger), so further testing on other cases is necessary. With the example model syntax in the Github repository as a template, multivariate modelling is no more challenging than univariate modelling, but it does require more computational resources.

6. Conclusion

The method presented here, Modelled Multivariate Overlap, is a flexible framework for calculating vowel overlap. It allows for the computation of any vowel overlap measure, and the modelling step of its implementation can be adapted based on personal preference. MMO offers a substantial improvement over the same measures calculated from empirical distributions while still being multivariate, and with the added benefit of easy-to-calculate measures of uncertainty.

7. Acknowledgements

We acknowledge an FRQSC PhD fellowship to IS (B2Z-333884) and grants from NSERC (RGPIN-2023-04873) and the Canada Research Chair program (CRC-2023-00009) to MS. We thank Robin Dodsworth, Meghan Clayards, Jane Stuart-Smith, and Massimo Lipari for feedback, and Michael Smith for proof-reading.

8. References

- [1] M. C. Kelley and B. V. Tucker, “A comparison of four vowel overlap measures,” *The Journal of the Acoustical Society of America*, vol. 147, no. 1, pp. 137–145, 2020.
- [2] W. Labov, S. Ash, and C. Boberg, *The Atlas of North American English: Phonetics, Phonology and Sound Change*. De Gruyter Mouton, 2006.
- [3] J. Nycz and L. Hall-Lew, “Best practices in measuring vowel merger,” *Proceedings of Meetings on Acoustics*, vol. 20, no. 1, 2014.
- [4] M. A. Baranowski, “Phonological variation and change in the dialect of Charleston, South Carolina,” Ph.D., University of Pennsylvania, 2006.
- [5] T. L. Irons, “On the status of low back vowels in Kentucky English: More evidence of merger,” *Language Variation and Change*, vol. 19, no. 2, 2007.
- [6] A. J. Dinkin, “Dialect boundaries and phonological change in upstate New York,” Ph.D., University of Pennsylvania, 2009.
- [7] J. Nycz, “Second Dialect Acquisition: Implications for Theories of Phonological Representation,” Ph.D., New York University, 2011.
- [8] —, “New contrast acquisition: Methodological issues and theoretical implications,” *English Language and Linguistics*, vol. 17, no. 2, pp. 325–357, 2013.
- [9] J. Hay, P. Warren, and K. Drager, “Factors influencing speech perception in the context of a merger-in-progress,” *Journal of Phonetics*, vol. 34, no. 4, pp. 458–484, 2006.
- [10] D. E. Johnson, “Quantifying overlap with Bhattacharyya’s affinity,” *New Ways of Analyzing Variation 44, Toronto*, 2015. [Online]. Available: https://danielezrajohnson.shinyapps.io/nwav_44/
- [11] G. S. Morrison, “Comment on ‘A geometric representation of spectral and temporal vowel features: Quantification of vowel overlap in three linguistic varieties,’” *The Journal of the Acoustical Society of America*, vol. 123, no. 1, pp. 37–40, 2008.
- [12] A. B. Wassink, “A geometric representation of spectral and temporal vowel features: Quantification of vowel overlap in three linguistic varieties,” *The Journal of the Acoustical Society of America*, vol. 119, no. 4, pp. 2334–2350, 2006.
- [13] E. F. Haynes and M. Taylor, “An assessment of acoustic contrast between long and short vowels using convex hulls,” *The Journal of the Acoustical Society of America*, vol. 136, no. 2, pp. 883–891, 2014.
- [14] J. A. Stanley and B. Sneller, “Sample size matters in calculating Pillai scores,” *The Journal of the Acoustical Society of America*, vol. 153, no. 1, pp. 54–67, 2023.
- [15] S. Vasishth, B. Nicenboim, M. E. Beckman, F. Li, and E. J. Kong, “Bayesian data analysis in the phonetic sciences: A tutorial introduction,” *Journal of Phonetics*, vol. 71, pp. 147–161, 2018.
- [16] M. Sonderegger, *Regression Modeling for Linguistic Data*. MIT Press, 2023.
- [17] P.-C. Bürkner, “Advanced Bayesian multilevel modeling with the R package brms,” *The R Journal*, vol. 10, no. 1, pp. 395–411, 2018.
- [18] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, and A. Riddell, “Stan: A probabilistic programming language,” *Journal of Statistical Software*, vol. 76, 2017.
- [19] C. C. Voeten, W. Heeringa, and H. Van de Velde, “Normalization of nonlinearly time-dynamic vowels,” *Journal of the Acoustical Society of America*, vol. 152, no. 5, pp. 2692–2710, 2022.
- [20] M. Sonderegger, J. Stuart-Smith, M. McAuliffe, R. Macdonald, and T. Kendall, “Managing data for integrated speech corpus analysis in *SPeech Across Dialects of English* (SPADE);” in *The Open Handbook of Linguistic Data Management*, A. L. Berez-Kroeker, B. McDonnell, E. Koller, and L. B. Collister, Eds., 2022, pp. 195–207.
- [21] J. Mielke, E. R. Thomas, J. Fruehwald, M. McAuliffe, M. Sonderegger, J. Stuart-Smith, and R. Dodsworth, “Age vectors vs. axes of intraspeaker variation in vowel formants measured automatically from several English speech corpora,” in *Proceedings of the 19th International Congress of Phonetic Sciences*, 2019, pp. 1258–1262.
- [22] M. McAuliffe, E. Stengel-Eskin, M. Socolof, and M. Sonderegger, “Polyglot and speech corpus tools: A system for representing, integrating, and querying speech corpora,” in *Interspeech 2017*. ISCA, 2017, pp. 3887–3891.
- [23] P. S. Beddor, “A coarticulatory path to sound change,” *Language*, vol. 85, no. 4, pp. 785–821, 2009.
- [24] J. Gwizdzinski, S. Barreda, C. Carignan, and G. Zellou, “Perceptual identification of oral and nasalized vowels across American English and British English listeners and TTS voices,” *Frontiers in Communication*, vol. 8, 2023.
- [25] A. J. Aitken, “The Scottish vowel length rule,” *So Many People, Longages, and Tongues, Edinburgh: Middle English Dialect Project*, pp. 131–157, 1981.
- [26] B. Lobanov, “Classification of Russian vowels spoken by different speakers,” *Journal of the Acoustical Society of America*, vol. 49, no. 2B, pp. 606–608, 1971.
- [27] J. C. Wells, *Accents of English: Volume 1*. Cambridge University Press, 1982.
- [28] S. Fitt, “Documentation and user guide to UNISYN lexicon and post-lexical rules,” Centre for Speech Technology Research, Edinburgh, Tech. Rep., 2000.
- [29] L. A. Ciaccio and J. Veríssimo, “Investigating variability in morphological processing with Bayesian distributional models,” *Psychonomic Bulletin & Review*, vol. 29, no. 6, pp. 2264–2274, 2022.
- [30] D. Lewandowski, D. Kurowicka, and H. Joe, “Generating random correlation matrices based on vines and extended onion method,” *Journal of Multivariate Analysis*, vol. 100, no. 9, pp. 1989–2001, 2009.