Evaluating generalised additive mixed modelling strategies for dynamic speech analysis

Márton Sóskuthy^a

^a*University of British Columbia, BC, Canada*

Abstract

Generalised additive mixed models (GAMMs) are increasingly popular in dynamic speech analysis, where the focus is on measurements with temporal or spatial structure such as formant, pitch or tongue contours. GAMMs provide a range of tools for dealing with the non-linear contour shapes and complex hierarchical organisation characteristic of such data sets. This, however, means that analysts are faced with non-trivial choices, many of which have a serious impact on the statistical validity of their analyses. This paper presents type I and type II error simulations to help researchers make informed decisions about modelling strategies when using GAMMs to analyse phonetic data. The simulations are based on two real data sets containing F2 and pitch contours, and a simulated data set modelled after the F2 data. They reflect typical scenarios in dynamic speech analysis. The main emphasis is on (i) dealing with dependencies within contours and higher-level units using random structures and other tools, and (ii) strategies for significance testing using GAMMs. The paper concludes with a small set of recommendations for fitting GAMMs, and provides advice on diagnosing issues and tailoring GAMMs to specific data sets. It is also accompanied by a GitHub repository including a tutorial on running type I error simulations for existing data sets: https://github.com/soskuthy/gamm_strategies.

Keywords: generalised additive mixed models, type I/II error simulations, dynamic speech analysis, random effects, significance testing

1. Introduction

There is increasing interest in dynamic properties of speech in phonetics and related fields. Dynamic analyses of speech focus on *contours* (e.g. formant, pitch or tongue contours), which are series of measurements with temporal or spatial structure. The direct study of contours facilitates the discovery and analysis of complex phonetic patterns without relying on overly simplistic assumptions about the data – for instance, that all vowels can be reduced to single-point F1/F2 measurements, or that variation in tongue shapes is simply a function of frontness and height. However, dynamic analyses present a number of challenges that are not easily overcome using traditional statistical tools. First, phonetic contours are often inherently wiggly, and this wiggleness is difficult to capture using parametric methods such as linear regression. Second, the focus in dynamic analyses is rarely on a single contour (as opposed to, e.g. studies of changes in global temperature;

Cahill et al. 2015): the aim is typically to make comparisons across contours as a function of various categorical or continuous predictors such as segmental context or speaker year of birth. As a result, dynamic phonetic data sets have not only spatial/temporal structure, but also hierarchical structure, with individual measurements (such as formant estimates at specific time points) grouped into contours. These contours are in turn often grouped into higher-order units such as words or speakers, creating complex dependencies among individual measurements. Statistical analyses that ignore the wiggleness of phonetic contours or dependencies among measurements will almost certainly lead to inaccurate conclusions.

Generalised additive mixed models (GAMMs; Hastie and Tibshirani 1990; Wood 2017) provide an elegant and flexible solution to these challenges. GAMMs are an extension of generalised linear mixed models that allow so-called *smooth terms* alongside conventional parametric terms. Smooth terms can approximate a wide range of non-linear shapes, estimating their degree

Email address: marton.soskuthy@ubc.ca (Márton Sóskuthy)

of wiggleness directly from the data. Moreover, as the name suggests, GAMMs are a type of ‘mixed’ model, which means that they include random effects that bring dependencies among grouped data points directly into the model, and allow for well calibrated inferences even in the presence of hierarchical structure. These properties make GAMMs an ideal tool for dynamic phonetic analysis.

GAMMs have been gaining in popularity in phonetics (Kösling et al., 2013; Wieling et al., 2016; Strycharczuk and Scobbie, 2017; Sóskuthy et al., 2018; Noiray et al., 2019; Carignan et al., 2020; Renwick and Stanley, 2020) and there is a growing literature that provides guidance on using them for linguistic analysis (Baayen et al. 2016, 2017; Sóskuthy 2017; Winter and Wieling 2016; Wieling 2018; Chuang et al. 2020; Baayen and Linke in press). The work presented in this paper serves as a complement to these introductions. It evaluates emerging practices in the modelling of phonetic data using GAMMs via type I and II error simulations.

Simulations often constitute an important supplement to theoretical arguments in statistics (see e.g. Marra and Wood 2012; Wood et al. 2016). They have also played an essential role in applied contexts by highlighting the severe consequences of using inadequate models (see e.g. Barr et al. 2013 and Matuschek et al. 2017). Following this line of work, the simulations in this paper will be used to (i) illustrate how different choices about model specification may affect the validity of GAMM analyses of phonetic data and (ii) provide comparisons where multiple alternative strategies are available. While some general recommendations are made at the end, the goal is not to provide a one-size-fits-all approach to GAMM analysis, but to empower analysts to make informed decisions about their models.

All simulations are carried out using the `mgcv` (Wood, 2017) and `itsadug` (van Rij et al., 2017) packages in R (R Core Team, 2013). Several different data sets are used: a set of formant measurements for the diphthong /aɪ/ in New Zealand English (Sóskuthy et al., 2019); a simulated data set modelled after these formant measurements; and pitch contours representing contrastive focus in German (Roessig et al., 2019). While real data are not commonly used in simulation studies, they play a crucial role in the current paper. Due to their temporal/spatial structure, phonetic contours show complex patterns of variation that are essential to capture in the simulations. While the simulated data sets are also modelled after real contours, the use of real data substantially boosts the ecological validity of the work presented here, and also shows how the methods under investigation generalise across different types of phonetic

contours.

The simulations focus on type I error rates (false positives) and power. In these simulations, the true underlying patterns are known in advance, making it possible to check whether the rate of false positives is nominal (close to $\alpha = 0.05$) and whether the model can reliably detect true differences between sets of contours. The data sets under examination are hierarchically structured: all of them contain measurements grouped into contours, and some of them further group these contours into higher-order units (speakers). The simulations look at a range of different problems: dealing with dependencies within contours; capturing dependencies in higher-order items such as sets of contours grouped into speakers; and choosing among different methods for significance testing and model fitting.

The paper is structured as follows. The rest of the introduction provides a brief overview of previous work in dynamic phonetic analysis (1.1) and some background to GAMMs (1.2). Section 2 outlines the general methodological approach. The results section (3) is divided into four parts. Each part addresses a specific question about fitting GAMMs to dynamic speech data, and presents further methodological details alongside the results. Section 4 concludes the paper with a discussion of the main findings and a small set of recommendations.

The aim of this paper is to make GAMMs accessible to a wider audience in phonetics by providing advice on modelling strategies. However, it is not a tutorial. Therefore, only essential pieces of R code are included in the main text. The paper is accompanied by a GitHub repository that includes the code for all of the simulated data sets and models presented below: https://github.com/soskuthy/gamm_strategies. This repository is an integral part of the paper, and it is recommended that readers consult the files in the `example_code` subfolder. To make this easier, the paper contains frequent references to relevant files in this folder. These references are of the form `example_code/filename.md` and refer to md documents that can be viewed in any browser via GitHub. All in-text references to md documents are clickable hyperlinks that take the reader directly to the relevant page. Note that Rmarkdown documents with executable code are also included in the folder. The repository includes all the code for running and analysing the simulations in the paper. In addition to these materials, the repository also includes a brief tutorial on running type I error simulations for real data sets.

1.1. Dynamic speech analysis

For the purposes of this paper, dynamic speech analysis is defined as the analysis of phonetic contours. Contours include (i) temporally ordered measurements such as formant tracks, pitch contours and articulatory trajectories (measured through techniques such as electromagnetic articulography); and (ii) spatially ordered measurements such as ultrasound tongue contours. There is no single gold-standard in dynamic speech analysis, and practices differ widely across disciplines. This is understandable given the considerable variation in types of data and research questions across different fields. In this section, we review a few examples that illustrate the questions and methods that are characteristic of dynamic speech analysis. These will provide a useful background to the rest of the paper, which focuses more specifically on GAMMs and problems that are amenable to analysis through generalised additive modelling.

The notion of dynamic speech analysis is best approached by considering a continuum of methods that vary in the amount of dynamic information that they make use of. At one end of this continuum are methods that reduce contours to a single measurement. This is still the prevailing approach in acoustic analysis. For instance, vowels are typically represented by formant values taken near their temporal midpoint, ‘steady-state’ section or formant maxima / minima (Peterson and Barney, 1952; Harrington et al., 2000; Labov et al., 2005). Such single-point measurements correspond well with traditional views of the vowel space (Jones, 1918), make it easy to visualise the vowel system of a language variety as a whole, and do an excellent job at capturing patterns of variation and change for many vowels despite their simplicity (see e.g. Labov 1994). Single-point measurements are less wide-spread in work on spatial contours (such as tongue contours), partly because of the difficulties involved in choosing a suitable measurement point and partly because the temporal dynamics of articulator movements are often of key importance (Gick et al., 2006; Scobbie et al., 2012; Carignan et al., 2020). Thus, even when spatial contours are reduced to a single point (e.g. a single location along the tongue surface identified through EMA or ultrasound), the focus is often on tracking the movements of this point over time (Gick et al., 2006; Strycharczuk and Scobbie, 2017).

The bulk of the continuum mentioned above is made up by methods that go beyond single-point analyses by utilising some dynamic information, but still rely on abstract generalisations over contour shapes. Such approaches typically capture dynamic information in a

small number of numeric parameters that reflect specific properties of contours. For instance, Hillenbrand et al.’s (1995) acoustic analysis of vowels in American English considers vowel duration and changes in F1 and F2 (based on measurements at 20%, 50% and 80% of vowel duration) alongside static measures of fundamental frequency and formants. Their analysis shows that coarse-grained information about formant movements contributes substantially to the discriminability of vowels. Changes in formant values are especially important in the analysis of diphthongs, and are frequently captured by calculating the *Euclidean distance* between the nucleus and the offglide (or onglide) of the diphthong in F1/F2 space (Haddican et al., 2013; Cardoso, 2015). Euclidean distance can then be used as a measure of diphthongisation and compared across different environments or speakers. Another technique that is often used in the analysis of vowel formants is the so-called *discrete cosine transform* (DCT; Watson and Harrington 1999; Harrington et al. 2008). The DCT represents contours using a set of coefficients obtained through a simple closed-form mathematical transformation. The first three of these coefficients correspond to the mean, the slope and the curvature of a contour. Watson and Harrington (1999) use DCT to show that while average formant values and slope play an important role in distinguishing vowels in Australian English, curvature does not. Studies focusing on tongue contours often rely on yet another method: *principal components analysis* (PCA; Slud et al. 2002; Turton 2015). PCA is a flexible method that allows researchers to detect relevant dimensions of variation in complex data, and reparameterise the observations in terms of these dimensions. Turton (2015) uses PCA to establish the degree of darkening in realisations of /l/ in different varieties of English based on ultrasound contours, which allows her to make observations about the degree to which dark /l/ and clear /l/ constitute separate categories in these varieties.

At the other end of the continuum are fully non-parametric approaches that do not restrict comparisons to specific characteristics of contours such as slope or curvature. Such methods are best suited to situations where the exact nature of the predicted variation in contour shapes is not clear in advance. They are often used when plotting contours for the purposes of visual illustration. Most work that considers dynamic aspects of speech relies on plots of contours to give readers a visual impression of the phonetic objects under investigation. These contours are sometimes presented in raw format, and sometimes aggregated and/or smoothed using moving averages, locally weighted scatterplot smoothing, smoothing splines and other methods.

Non-parametric methods can also be used for statistical inference. Davidson (2006) introduces *smoothing spline ANOVA* (SS-ANOVA) for comparing tongue contours obtained through ultrasound imaging. SS-ANOVAs model sets of contours and their differences as smooth functions, and allow researchers to detect regions where the contours diverge by looking at confidence intervals around the estimated difference curve. GAMMs also belong to this general group of methods, and are in many ways similar to SS-ANOVA, although current state-of-the-art implementations of GAMMs are far more flexible (Wood, 2017). Previous applications of GAMMs in phonetics include analyses of EMA trajectories from different dialects of Dutch (Wieling et al., 2016); a comparison of native English *versus* Dutch and German productions of English obstruents from an articulatory perspective (Wieling et al., 2017); an ultrasound-based dynamic study of articulatory fronting in Southern British English high back vowels in different segmental contexts (Strycharczuk and Scobbie, 2017); an investigation of /u/-fronting across three generations in Derby English based on F2 contours (Sóskuthy et al., 2018); a study of changing vowel dynamics in the front vowel system of Southern American English (Renwick and Stanley, 2020); and a GAMM analysis of real-time MRI data focusing on dynamic changes in vocal tract shapes (Carignan et al., 2020).

Another non-parametric method that has recently been used in the analysis of phonetic contours is functional data analysis (FDA; Gubian et al. 2015). FDA combines smoothing via splines, landmark registration using dynamic time warping, and functional principal component analysis to automatically discover salient dimensions of variation in contour shapes, and represent curves in terms of a small number of scalar parameters along these dimensions. These parameters can then be analysed using conventional linear mixed effects regression. FDA has been used to analyse formant and pitch contour data for diphthongs and vowels in hiatus in Spanish (Gubian et al., 2015) and L1 and L2 prosody as reflected by f0 contours in Japanese and German (Asano and Gubian, 2018).

1.2. Generalised additive modelling

This section provides an introduction to generalised additive modelling, focusing on essential concepts underlying the simulations in later sections. This introduction is necessarily superficial. Interested readers are referred to more detailed tutorials such as Sóskuthy (2017), Wieling (2018), Winter and Wieling (2016), Chuang et al. (2020), Baayen and Linke (in press) and Wood’s (2017) comprehensive book-length treatment. All of the

Figure 1: The faint filled circles connected by lines indicate five simulated F2 contours measured at 5% intervals. The darker line indicates a simple generalised additive model fitted to the contours. The insets show the two model terms: a parametric term (the intercept) and a smooth term centred around 0.

discussion below refers to generalised additive models as implemented in the *mgcv* package (Wood, 2017).

1.2.1. GAMs

Let us start with GAMs, or generalised additive models, which differ from GAMMs (generalised additive *mixed* models) in that they do not include random effects. The fundamentals of GAMs will be illustrated through a toy data set, shown in figure 1. The data set consists of five F2 contours (modelled after the diphthong /aɪ/) represented by measurements taken at 5% intervals throughout their duration. The contours are time-normalised, that is, they are all shown on the same scale regardless of their absolute duration.

The dark line in the main panel of figure 1 illustrates a simple GAM fitted to this data set. This model was fitted using the following code:

```
bam(f2 ~ s(measurement.no, k=10),
    data=f2.dat.1)
```

The syntax for GAMs is similar to conventional linear regression models in R, except for the use of the function `bam()` (an efficient implementation of GAMs in the package *mgcv*) and the smooth term specification via the function `s()`. The smooth term specification tells *mgcv* to fit a smooth over `measurement.no` with 9 basis functions (this is set using the `k` parameter, which equals the number of basis functions plus one for most fixed effect smooths). The notions of smooths and basis functions are explained in the next few paragraphs.

Similar to other regression models, the model predictions in the plot are obtained by summing separate model terms. In this case, one of these model terms is the intercept. As usual, the intercept does not need to be explicitly specified in the model formula. The intercept is a so-called *parametric* term, which means that it behaves in the same way as it would in a conventional linear model. The intercept is shown separately in the inset on the left in figure 1. The other model term is a *smooth* term, which allows the model to capture the non-linear pattern in the contours using a smooth function. This term is shown in the second inset in figure 1. Smooth terms are extremely flexible and can take on a wide range of different shapes. Note that most fixed-effect smooth terms in `mgcv` are centred around zero with respect to the outcome variable. As a consequence, parametric terms represent height-related patterns that hold across the entire length of the contour (in this case, the average F2), while smooth terms represent patterns related to the *shape* of the contour (in this case, the non-linear rise in F2 over time).

Smooth terms are created by representing the data using a set of simple non-linear functions, multiplying each of these functions by a separate coefficient estimated from the data, and then summing them. This process is conceptually very similar to polynomial regression, where both a variable x and some of its powers (x^2 , x^3 , ...) are included in a model as predictors. These simple component functions are called *basis functions*. GAM smooths are usually *splines* such as cubic or thin plate regression splines (henceforth *cr* and *tp* splines), and are composed of special basis functions that ensure reasonable smoothing performance. Just as with polynomials, the number of basis functions has a strong influence on the flexibility of a spline. The smooth term in figure 1 has 9 basis functions. Figure 2 shows two additional smooth terms fitted to the same data with 3 (left) and 14 (right) basis functions. The faint dashed curves are the basis functions multiplied by their coefficients, while the dark curves show the smooths obtained by summing them pointwise along the horizontal axis. This process of summing basis functions, represented visually in figure 2, is how complex shapes are captured *via* smooths. The smooth with 3 basis functions is less flexible than the one with 14 basis functions, and (in this case) provides an inferior fit to the data.

The smooth with 9 basis functions (figure 1) looks nearly identical to the one with 14 basis functions (figure 2, right), even though we might expect it to be wigglier given the higher number of basis functions. This is because the number of basis functions simply imposes an upper limit on the degree of wiggleness; the actual

degree of wiggleness is constrained by a *smoothing parameter*. Crucially, this smoothing parameter is estimated from the data, and does not need to be changed by the user. Model summaries of GAMs generated using the `summary()` function display the *estimated degrees of freedom* (EDF) for each smooth term, which decreases monotonically as a function of the smoothing parameter. The EDF is mapped onto the range between 1 and the number of basis functions, providing an intuitive measure of the actual amount of wiggleness used by the smooth (with 1 corresponding to a straight line and higher numbers corresponding to more wiggly curves). The smoothing parameter is part of a penalty term used during the estimation of the model parameters, which ensures that the model strikes a sensible balance between goodness of fit and smoothness.

Being able to estimate the smoothing parameter from the data reduces the involvement of the user in determining the degree of complexity necessary for a given model. As long as the number of basis functions is sufficiently high, the model will be able to settle on a reasonable smoothing parameter. However, the user still has to make a decision about the number of basis functions. This decision is often based partly on intuition: for instance, 3 basis functions can only capture relatively simple patterns, while 100 basis functions can capture a great deal of complexity in contour shapes. The package `mgcv` also offers a function called `gam.check()` that can help the user determine whether a smooth needs a higher number of basis functions or if it is already sufficiently complex.

Figure 3 illustrates three smooths with the same number of basis functions but different smoothing parameters. The smooth with the lowest smoothing parameter has the highest degree of wiggleness, which results in overfitting. As the smoothing parameter is increased, the fitted curves become less wiggly. Note that conventional fixed effect smooths converge to a straight line estimate at high values of the smoothing parameter (cf. the dark line in figure 3).

1.2.2. Residual autocorrelation

One problem with the simple GAMs above is that they fail to recognise hierarchical and spatial/temporal dependencies within the data. Measurements from the same contour are not independent, and measurements that are closer together within the same contour tend to be more similar than ones that are further apart. This issue can be expressed in terms of *residual autocorrelation*: individual contours deviate from group trends in systematic ways, leaving patterns in the model residuals such that residuals for neighbouring data points

Figure 2: Two GAMs fitted to the simulated F2 contours. The faint dashed curves show the basis functions multiplied by coefficients estimated from the data. The dark curves show the model predictions. Left: a *cr* smooth with 3 basis functions. Right: a *cr* smooth with 14 basis functions.

Figure 3: Smooths with different smoothing parameters. Darker lines indicate a higher smoothing parameter. The faint circles represent the measurements to which the models are fitted.

Figure 4: Left: Residuals for three contours plotted against time (residuals from the same contour have the same shade and are connected by lines). Right: Residual autocorrelation values (vertical axis) at different lags (horizontal axis).

are strongly correlated. The left-hand panel of figure 4 shows the residuals for three of the five contours plotted against time. There is clear temporal patterning in these residual series. The right-hand panel illustrates this by plotting the correlation between pairs of residuals at fixed distances (or *lags*) along the time dimension, that is, the within-contour autocorrelation (for the current plot, these sets of correlations are calculated separately for each contour and then averaged across contours). For instance, the autocorrelation at lag 1 (the second vertical line) is calculated by taking the residuals at measurement points 0, 1, 2, ... and correlating them with the residuals at measurement points 1, 2, 3, ... The autocorrelation at lag 0 is always 1, as it is simply the correlation between every point and itself. In an ideal scenario, the remaining autocorrelation values would all be close to zero. This is certainly not the case for the current model: for example, the autocorrelation at lag 1 is around 0.8, which indicates a high degree of dependence between adjacent data points that is not captured by the model. Statistical models tend to be overconfident in the presence of such dependencies, which makes confidence intervals overly narrow and biases *p*-values downward.

GAMs offer two options for removing autocorrelations in the residuals, which can be used either on their own or in combination (Baayen et al. 2016). The first of these is to allow the model to adjust its predictions separately for each contour by using random effects, making it a GAMM. GAMMs can include three different types of random effects: random intercepts, random slopes and random smooths (these options are illustrated in more detail in Sóskuthy 2017, Wieling 2018 and in `example_code/set_1-real_f2.md`). Random intercepts adjust the model intercept separately for each contour, shifting the predicted contours upward or downward without changing their shape. Random slopes over time essentially tilt the predicted contours to yield a closer fit to the data. Since the set of contours in figure 1 shows relatively complex and fine-grained variation in trajectory shapes, random intercepts and slopes are unlikely to be sufficient for removing the patterns in the residuals. Random smooths over time are much more flexible: they allow the model to fit separate smooths to each group in the data (in this case, each contour), which all share the same smoothing parameter. Similar to other smooth terms, the number of basis functions determines the maximum wiggleness that can be accommodated by a random smooth, but the actual amount of wiggleness is estimated from the data. Random smooths behave as proper random effects in that a model with random smooths can also include a fixed effect smooth

representing the overall pattern in the data set.

The code below illustrates a version of our model that includes random smooths:

```
bam(f2 ~
    s(measurement.no, k=10) +
    s(measurement.no, traj,
      bs="fs", m=1, k=10),
  data=f2.dat.1)
```

Random smooths are specified by setting the value of the parameter `bs="fs"`, adding a grouping variable (in this case, trajectory IDs coded by the variable `traj`) and setting `m=1` (which induces a small amount of shrinkage in the estimated random smooths).

While random smooths are essential in some contexts, they are computationally very costly. Each level of the grouping factor adds *k* coefficients to the model, where *k* is the number of basis functions. This creates a significant computational bottleneck in terms of processing time, and, more importantly, memory.

The second option is to use an autoregressive error model (AR model). There are many different kinds of autoregressive models, but the current efficient implementation of GAM in `mgcv` (the `bam()` function) only supports the simplest of these: the so-called AR1 error model. An AR1 error model estimates the model parameters under the assumption that the errors for neighbouring observations in the contour are correlated: for instance, the error at measurement number 4 is partly determined by the error at measurement number 3. The AR1 model assumes that such correlations only exist between immediately adjacent points in the contour. In `mgcv`, the correlation parameter for an AR1 model has to be set manually, but an estimate from an identical GAM without an AR1 component is usually adequate (cf. Baayen et al. 2016; Wieling et al. 2016). The use of AR1 error models is illustrated in more detail in `example_code/set_1-real_f2.md`.

1.2.3. Hierarchical dependencies

Random effects in GAMMs can also be used to capture higher-level hierarchical dependencies. For example, each speaker may contribute multiple formant contours to the same data set. Contours from the same speaker are likely to be more similar to each other than to contours from different speakers. This hierarchical patterning should be brought into the model. If the main predictor of interest varies *across* speakers, but not within speakers, this can be done by including random intercepts by speaker, random slopes over time by speaker or random smooths over time by speaker. These work

very similarly to the random effects described above, the only difference being that they now capture by-speaker dependencies instead of by-contour dependencies.

If the main variable of interest varies *within* speakers, a different solution may be necessary. Let us assume that the variable of interest is phonological context. One option, then, is to include both a random intercept by speaker and a random slope over context by speaker, which adjusts the size of the effect of context on average contour height separately for each speaker:

```
s(speaker, bs="re") +  
s(speaker, context, bs="re")
```

This solution is not particularly flexible, in that it cannot deal with situations where context also varies across speakers in terms of its effect on contour shape. There are several potential remedies for this type of situation. Wieling et al. (2016) and Sóskuthy (2017) use random smooths over time with a special grouping variable that has separate levels for each combination of speaker and context (speaker₁+voiced, speaker₁+voiceless, speaker₂+voiced, speaker₂+voiceless, etc.):

```
s(measurement.no, speakerContext,  
  bs="fs", m=1)
```

This type of random smooth will be referred to as an ‘item × effect’ random smooth. Wieling (2018) relies on a similar method, but instead of setting up a new grouping variable, he simply uses `speaker` as the grouping variable and adds `context` using the `by` parameter for smooth specifications:

```
s(measurement.no, speaker,  
  by=context, bs="fs", m=1)
```

We will refer to this as an ‘item-by-effect’ type random smooth (cf. `example_code/set_3-sim.f2.md`).

‘Item × effect’ and ‘item-by-effect’ random smooths are conceptually analogous to a conventional mixed effects model with a random structure of the form $(1 \mid \text{speakerContext})$, that is, random intercepts for each combination of speaker and context. This is not how one would typically deal with within-speaker fixed effects in mixed effects models: the standard solution would be to use $(1 + \text{context} \mid \text{speaker})$, that is, a random intercept by speaker and a random slope over context by speaker. The advantage of this second method is that both the by-speaker intercepts and slopes are explicitly penalised, leading to shrinkage in both (Baayen, 2008, pp. 299–303). As we will see, this shrinkage is essential in that it ensures that random effects do not ‘overreach’ by erroneously accounting for

variance that should, in fact, be captured by a fixed effect in the model (which leads to overly conservative estimates and *p*-values). Building on this observation, we explore a further random smooth structure for dealing with within-speaker effects, one that is closely analogous to the combination of random intercepts and slopes discussed above. This method relies on two separate sets of random smooths: a set of random reference smooths (corresponding to random intercepts in conventional mixed models) and a separate set of random difference smooths (corresponding to random slopes):

```
s(measurement.no, speaker, bs="fs", m=1) +  
s(measurement.no, speaker,  
  by=context.ord, bs="fs", m=1)
```

The distinction between reference *versus* difference smooths is explained in the next section (including the variable `context.ord`, which is an ordered factor); the method proposed here is simply an extension of this distinction to random smooths. This type of random smooth structure will be referred to as ‘random reference/difference smooths’ (cf. `example_code/set_3-sim.f2.md`).

1.2.4. Significance testing with GAMMs

Let us now turn to significance testing using GAMMs, focusing on a specific question: are two sets of contours significantly different? GAMMs offer at least three broader sets of methods for significance testing (van Rij, 2015): (i) tests based on summary statistics about the parametric and smooth terms that are easy to extract from the model; (ii) model comparison between a full and a nested model; and (iii) visual methods. The first set of methods requires GAMMs with so-called *difference terms*. Figure 5 illustrates these using an extended toy data set that now has an additional five contours generated using a different underlying curve. The panel on the left shows the raw contours (faint lines and filled circles/triangles) and the predictions from a GAM fit with difference terms (dark lines). The panel on the right shows the four model terms required to represent this model. The first set of contours (solid blue lines on the left) can be modelled in the same way as before by summing the intercept and a smooth term, which are both shown in the top row of the right-hand panel. The terms in the bottom row capture the estimated parametric and smooth differences between the contours. These indicate that (i) on average, the second set of contours is about 75 Hz lower than the first one and (ii) the second set of contours has a different shape from the first set. The model predictions for the second group (dark

Figure 5: Left: A GAM fitted to an extended data set with two groups of contours (indicated by different shapes, colours and line types). The faint lines and filled circles/triangles show the contours, while the dark lines show the model predictions. Right: The model terms from the GAM. Summing the parametric and smooth terms in the top row yields the solid blue curve in the left-hand panel. The sum of the parametric and smooth terms in the bottom row represents the difference between dashed red and solid blue curves in the left-hand panel.

red dashed line in the left-hand panel) are obtained by summing all four of the terms on the right.

The code below illustrates a GAMM with a difference smooth fitted to the extended data set:¹

```
bam(f2 ~ group +
    s(measurement.no) +
    s(measurement.no, by=group.ord) +
    s(measurement.no, traj,
        bs="fs", m=1),
    data=f2.dat.2)
```

As before, the intercept is not explicitly specified in the model. The first predictor represents the parametric difference term, the second one the reference smooth and the third one (with the `by` argument) the difference smooth. The last predictor is a random smooth by trajectory.

Using difference terms in the model allows us to test whether two sets of contours are different. The `summary()` command in `mgcv` performs separate statistical tests to check whether (i) the parametric term is different from zero (i.e. whether they have different heights) and whether (ii) the smooth term is different from zero (i.e. whether they have different shapes), returning separate p -values for each of them (these tests

are based on methods described in Marra and Wood 2012; Wood 2013). Using two separate tests to check for differences between contours might be problematic when our hypothesis does not specify whether the expected differences are in height or shape, leading to inflated false positive rates due to multiple comparisons (cf. section 3.4).

Another option is to use visual methods based on confidence intervals. Figure 6 illustrates two possible options. The first of these is to check whether there is overlap in the confidence intervals around the fitted contours (including the intercept and parametric effects, but excluding random effects; see e.g. Davidson 2006, 411). This method is problematic: overlapping confidence intervals do not necessarily imply non-significant differences. An intuitively more reliable approach is to plot the estimated difference between the curves and check whether the confidence interval around it includes zero (Davidson, 2006; Sóskuthy, 2017). If it does, the difference at that specific point is non-significant, otherwise it is significant. There are at least two scenarios where this procedure is useful and generally accepted. First, visual methods are key to interpreting the output of GAMMs and offer an ideal tool for exploratory investigations. When used in this way, they are an essential complement to the methods of significance testing above. Second, the researcher may have a prior hypothesis about where a difference is expected between two curves, in which case using the second visual method described above may be entirely adequate.²

¹The variable that codes the difference between the two sets of contours (`group.ord`) needs to be converted to an ordered factor with contrast coding before fitting this model. This is simply how difference smooths are coded in GAM(M)s. Sóskuthy (2017) and Wieling (2018) provide more guidance on this (see also the code in the example files).

²Note that some plotting functions can also control for the addi-

Figure 6: Left: Model predictions for the two groups of contours with 95% pointwise confidence intervals. Right: The estimated difference between the two groups with the associated 95% pointwise confidence interval. The highlighted area indicates where the confidence interval excludes zero.

However, visual techniques are sometimes used in a confirmatory setting without such hypotheses. In such cases, the researcher may declare a significant result if there is any region of any length where the confidence interval around the difference curve excludes zero. Intuitively, this approach may lead to erroneous inferences, since it uses pointwise differences as a proxy for an ‘overall’ difference between the two curves. This is especially problematic when the region where the two contours differ is narrow: for instance, should one declare a significant result if the confidence interval around the difference curve excludes zero along a mere 2% of the entire contour?

Given the above considerations, a promising alternative strategy might be to perform a model comparison between a full model and a nested model that simultaneously excludes both difference terms (but retains the intercept and the reference smooth). The nested model corresponding to the model shown above and plotted in figure 5 is as follows:

```
bam(f2 ~ s(measurement.no) +
    s(measurement.no, traj,
      bs="fs", m=1),
  data=f2.dat.2)
```

Likelihood ratio-based comparisons between models with different fixed effects are only reliable when they

are fitted using maximum likelihood estimation (ML). The `bam()` function in `mgcv` offers several different estimation methods, including ML, restricted ML (REML) and fast restricted ML (fREML).³ ML is computationally much less efficient than REML and especially fREML. In addition, more recent versions of `mgcv` also allow ‘marginal discretization of model covariates’ and parallel computation (Wood et al., 2017, 1), which can yield substantial performance gains, but are only available when the model is estimated via fREML. As a result, ML often ends up being several orders of magnitude slower (and more memory-intensive) than alternative methods, which can make model comparisons impossible.

There are three possible ways to deal with this issue. First, model comparisons based on conditional AIC (as computed by the `AIC()` function) do not suffer from the same restriction, and can be used with models estimated via fREML (Wood et al., 2016). This simply involves obtaining AIC scores for the full and nested model, and comparing them: if the full model has a lower AIC score, the additional model complexity is justified. However, it is difficult to control type I error rates using AIC as this method is not based on null-hypothesis significance testing. It may therefore be useful to interpret the output of AIC comparisons in a more graded way, using the actual size of the difference in AIC scores to evaluate the strength of the evidence for

tional uncertainty that results from the estimation of the smoothing parameter, yielding more accurate confidence intervals. The functions in the `itsadug` package used for plotting in this paper do not yet have this functionality, but the `plot.gam()` function in `mgcv` (the package that we use for fitting GAMs) does: it can be enabled by setting the parameter `unconditional` to the value `TRUE`.

³A discussion of the differences among these three different types of estimation methods is beyond the scope of the present paper. Oehlert (2014) provides an accessible exposition of the difference between ML and REML. Wood et al. (2015) presents the technical details of fREML.

a given model over another (cf. Burnham and Anderson 2004). The second option is to use a so-called *binary smooth*, which integrates parametric and smooth difference terms into a single smooth (Wieling, 2018). Performing significance testing on the binary smooth is equivalent to the model comparison method described above, but it does not require an actual likelihood ratio test, and is therefore valid even when the model is not estimated using ML. A third alternative relies on the fact that REML and fREML can be used for model comparison when the models only differ in their *random* effects. It is possible to replace the parametric difference term with a random intercept and to place an additional null space penalty on the smooth difference term (by setting the `select` parameter of `bam()` to the value `TRUE`), effectively turning both of them into random effects. Model comparison can then be performed using models estimated via fREML.

All of the above options for significance testing are illustrated in `example_code/set_4-real_f2.md`.

2. General methods

The simulations are organised into four sets, most of which are run for all three data types (i.e. the real F2 contours, the real pitch contours and the simulated F2 contours). The first set looks at how well different random structures and error models can deal with dependencies within contours (cf. 1.2.2). The second set examines the ability of different random structures to capture dependencies in higher-order items (e.g. contours grouped into speakers) in cases where the main treatment variable only varies *across* items. The third set looks at a similar situation but with a treatment variable that varies *within* items (cf. 1.2.3). The final set of simulations compares different methods for significance testing (cf. 1.2.4).

2.1. Data

The simulations are run using three different data types: (i) a set of real F2 contours representing the diphthong /aɪ/ in New Zealand English (Sóskuthy et al., 2019); (ii) a set of pitch (F0) contours representing nuclear accented syllables produced with contrastive focus in German (Roessig et al., 2019); and (iii) a set of simulated F2 contours. In type I error simulations, each contour is randomly assigned to one of two different groups (referred to as *A* and *B*), resulting in a situation where the true curves underlying the data do not differ across the groups. In type II error simulations, the contours in groups *A* and *B* come from different underlying curves.

Both type I and type II error simulations are run for the simulated F2 contours. For the real contours, only type I error simulations are run, as type II error simulations would require arbitrary decisions about the exact shape of the difference between the underlying curves for groups *A* and *B*.

The following sections provide a broad description of each data type. Further simulation-specific details are given in section 3.

2.1.1. Real F2 contours

The real F2 contours come from a study of diachronic changes to diphthong dynamics in New Zealand English (Sóskuthy et al., 2019), based on the Origins of New Zealand English corpus (Gordon et al., 2007). The contours represent the vowel /aɪ/ (e.g. *pride*, *sigh*). In order to avoid significant sources of systematic variation in our sample, the contours are all from 30 female speakers born between 1963 and 1982, and are limited to /aɪ/ before voiced consonants and in final position.

For each vowel, 11 equally spaced F2 measurements were collected automatically using LaBB-CAT (Fromont and Hay, 2008) and *Praat* (Boersma and Weenink, 2009). These measurements thus represent samples at 0%, 10%, 20%, ..., 100% into the vowel. The 30 female speakers chosen for inclusion in our analyses all had reasonably high-quality automatic measurements, and all produced a minimum of 30 contours. The actual number of contours per speaker varies between 31 and 123, with a median of 39. There are some missing measurements for some of the contours due to formant tracking errors. A few representative speakers are shown in the top left panel of figure 7.

2.1.2. Real pitch contours

The real F0 contours are from a study of the prosody of different focus types in German (Roessig et al., 2019).⁴ In the current study, we only use F0 contours representing contrastive focus, which are fairly consistently realised with a sharp rising contour. The data come from 27 speakers of German (17 female and 10 male). All the contours are from sentences using the same syntactic frame: *Er hat X auf Y gelegt*, where *X* is the name of a tool (e.g. *den Hammer* ‘the hammer’) and *Y* is a disyllabic nonce form with initial stress (e.g. *die Wohse* ‘the wohse’). These sentences (and the corresponding question prompts) were designed to elicit a nuclear accent on the stressed syllable of *Y*. Each speaker is represented by 18 to 20 contours in the data.

⁴We thank Simon Roessig, Doris Mücke and Martine Grice for making their raw data available to us.

Figure 7: *top left*: example contours from the real F2 data set; *top right*: example contours from the pitch data set; *bottom left*: parameterisation of the simulated F2 contours, with the red labels indicating additional parameters used to define the second set of contours in type II simulations; *bottom right*: example contours from the simulated F2 data set.

We used REAPER (Talkin, 2015) to extract F0 measurements at 10 ms intervals from the beginning of the nuclear accented syllable to 100 ms after its end. This yielded 23 to 59 measurements per contour, with a median of 39. Note that this data set differs substantially from the F2 data in that the number of measurements per contour are not fixed, and there are also more measurements per contour than for F2. As a consequence, there is substantial variation in the duration of the measurement intervals relative to the entire contour duration (ranging from 1.7% to 4.3%). Some measurements are missing due to pitch tracking errors. We use log F0 in our analyses. The systematic effect of speaker sex on average contour height has been removed from the data via residualisation. This does not affect random across-speaker variation in contour height. The top right panel of figure 7 shows the full set of F0 contours for a few representative speakers.

2.1.3. Simulated F2 contours

The bottom left panel of figure 7 illustrates the underlying curves used to generate the simulated F2 contours. The blue solid curve is used for all contours in type I simulations and for group A contours in type II simulations. Group B contours in type II simulations are based on the red dashed curve. The labels and arrows in the graph indicate parameters that vary across contours. Each contour is therefore defined by a set of parameters sampled from independent probability distributions (described in more detail in the appendix).

The contours are generated using logistic curves parameterised using four values: the starting and end points (START, END), the horizontal location of the midpoint of the transition (TRANSITION_{horiz}) and the slope of the transition (TRANSITION_{slope}). The formula used for generating these curves is shown in Appendix A. In type II simulations, START, TRANSITION_{horiz} and TRANSITION_{slope}

are sampled from the same distribution for all contours, but END_2 comes from a separate distribution for group B . This distribution has a lower mean, which yields curves that are identical near the beginning but end at a lower value. To simulate measurement noise and natural variation in formant contours, random noise is added to each measurement point along the contours at the end of the data generation procedure. All parameters are sampled from distributions that approximate those of the real F2 contours described in section 2.1.1 (e.g. the mean and the standard deviation of the starting point and end point of the contours is roughly the same, the amount of random noise is comparable, etc.). Similar to the real F2 data, each contour consists of 11 equally spaced measurements.

The bottom right panel of figure 7 shows contours from a few representative simulated ‘speakers.’ Unlike the real data sets, where the contours always come from the same limited set of speakers (and where randomness is achieved mainly via permutation of the group labels), the simulated contours are created anew for each iteration of the simulation. Comparing the top left and bottom right panels in figure 7 also shows that the real and simulated F2 contours are broadly similar, though there are differences both in the trajectory shapes and in the exact nature of the random variation across contours. It is therefore important not to equate these two data sets with each other in drawing conclusions from the simulations.

2.2. Procedure

For each set of simulations, 10,000 data sets are generated for all possible combinations of simulation type (type I vs. II) and data type (real F2, real pitch, simulated F2). Each data set is then analysed using a range of different modelling set-ups, described in more detail in the following sections. Each cell in the summary tables in section 3 presents a proportion or average based on an overall 10,000 simulation runs.

All models share a few general properties. They are all fit using the `bam()` function from version 1.8-31 of the `mgcv` package with `fREML` estimation (using the `method="fREML"` argument) and discretised covariates (`discrete=TRUE`) unless otherwise stated. Fixed smooth terms are constructed using *tp* smooths, while random smooth terms are constructed using either *tp* or *cr* smooths. With a few exceptions, all full models are of the following general structure (where *group* is an ordered factor):

$$f2 \sim 1 + \text{group} + s(\text{time}) +$$

$$s(\text{time}, \text{by}=\text{group}) + \dots$$

That is, the models contain an intercept (1), a parametric difference term (*group*), a reference smooth (`s(time)`) and a difference smooth (`s(time, by=group)`). The dots indicate further terms such as random effects, which vary across models. For the real and simulated F2 contours, the number of basis functions is set to 9 for fixed smooths (i.e. setting $k=10$); for the pitch contours, it is set to 19 ($k=20$). Random intercepts, slopes and smooths are implemented as recommended in Sósuthy (2017) and Wieling (2018). When the models include an AR1 autoregressive error model, the autocorrelation at lag 1 is estimated by extracting the residual autocorrelation from an identical model without AR1 fit to the same data (AR1 is never used with the simulated data, as the data generation process does not include autocorrelated errors).

For all models, we extract *p*-values for the parametric and smooth difference terms based on `summary.gam()`. In addition, for the simulated formant data we also extract the total processing time in seconds (including any nested models that need to be fit), and an estimate of peak memory usage in megabytes, measured by capturing the output of automatic garbage collection performed by `gc()` during the execution of the `bam()` command. Some additional information is recorded for certain models, such as *p*-values for model comparisons, or the amount of overlap between estimated confidence intervals when visual comparisons are performed.

The tables in the following section report type I error rate, power (one minus type II error rate) where available, median processing time in seconds and median peak memory usage in megabytes. The type II data sets are designed with relatively small differences between the two groups, typically yielding power estimates between 0.3–0.6 when the type I error rate is nominal. Type I error rates are reported for both the parametric and smooth terms. Power is only reported for the smooth term, since the simulated differences in the type II data sets reside mainly in the shapes of the contours. Note that power cannot be compared across different types of data in a meaningful way.

3. Simulations

3.1. Set 1: Dealing with autocorrelation

3.1.1. Methods

The first set of simulations compares different modelling strategies with respect to their ability to deal with dependencies within contours. Data sets are generated

as follows. For the real formant contours, one of the 30 speakers is selected randomly. Focusing on a single speaker eliminates the main source of higher-level hierarchical dependencies in the data (these dependencies are tackled in the next set of simulations). 50 or fewer contours are sampled from this speaker (each contour is only used once). The same process is used for the pitch contours, though the number of contours per simulation is only 18–20. For the simulated formant data, 50 different contours are generated randomly. This is analogous to using data from a single speaker, as the distributions from which the contour parameters are drawn are the same for each contour. For all data types, half of the contours are assigned to group *A* and half to group *B*. For type I error simulations, this assignment is random, ensuring that there are no systematic underlying differences between the two groups. For type II simulations (only run for the simulated data), the group assignment is non-random in that the underlying curves are slightly different (see appendix Appendix A.1 for details). These data sets do not have any further hierarchical structure, which makes them somewhat artificial in the context of phonetic research, where contours are usually grouped by speakers and other higher-order units. However, it provides an ideal test for methods of controlling autocorrelation within contours.

The following model components are evaluated in terms of their ability to capture within-contour dependencies:

- Random intercepts by contour.

```
s(trajectory, bs="re")
```

- Random intercepts plus slopes by contour.

```
s(trajectory, bs="re") +  
s(trajectory, measurement.no, bs="re")
```

- Random smooths by contour with different types/numbers of basis functions (types: *tp* and *cr*; numbers: $k = 3, 5$ and 10 for the F2 data; $k = 3, 5, 10, 15$ and 20 for the pitch data). An example of a random smooth of type *tp* with $k = 5$:

```
s(measurement.no, trajectory, bs="fs",  
m = 1, xt = "tp", k = 5)
```

- An AR1 autoregressive error model (not used for the simulated data, as the function used to create random noise around the measurements did not include correlations between neighbouring measurements). These are set by specifying the *AR.start*

and *rho* arguments of the function *bam()* (see the example code files for more detail).

In addition, the simulations also test random smooths with a flexible number of basis functions where the optimal value for k is determined using the function *gam.check()*. This is done by testing random smooths with increasing levels of complexity, and settling on a final model when (i) *gam.check()* suggests that the number of basis functions does not need to be increased any further (i.e. the p -value reported by *gam.check()* for the random smooths is higher than 0.05) or (ii) an upper limit is reached. The figures reported in the next section are calculated using this final model. For the F2 data, the value of k was increased in steps of 2 starting at 4 and reaching the upper limit at 10. For the pitch data, the value of k was increased in steps of 3 starting at 4 and reaching the upper limit at 16.

Type I error rate is calculated twice: once for parametric and once for smooth difference tests based on the model summary. Power is only calculated for the smooth difference term (see above). Processing times and memory usage are similar across the different data sets, and are therefore only shown for the simulated F2 data.

The data and models used in this section are illustrated in *example_code/set_1-real_f2.md*, *example_code/set_1-pitch.md* and *example_code/set_1-sim_f2.md*.

3.1.2. Results

Table 1 summarises the simulation results for all three data sets. The table is admittedly rather heavy on detail, so a brief walkthrough is provided below. Each row presents a summary for a different modelling set-up, with the rows ordered into blocks according to major modelling decisions. The first block shows the results for models without random smooths; the second block shows the results for random smooths with a thin plate (*tp*) basis; the third block for random smooths with a cubic regression spline (*cr*) basis; and the fourth block for AR1 error models. Error rates are calculated separately for the parametric and smooth tests. For instance, the very first numeric cell of the table in the upper left corner (with a value of 0.261) represents the type I error rate calculated using the parametric difference term from a model with no random components fitted to the real formant data. The cell to its right (0.323) shows the type I error rate for the same model and same data but based on the smooth difference term. The cell underneath (0.049) shows the parametric type I error rate for a different model (this time with random intercepts) fit-

MODEL COMPONENT	REAL FORMANTS		REAL PITCH		SIMULATED FORMANTS			MEM.	TIME
	PARAM.	SMOOTH	PARAM.	SMOOTH	PARAM.	SMOOTH			
	TYPE I	TYPE I	TYPE I	TYPE I	TYPE I	TYPE I	POWER		
NO SMOOTHS									
no components	0.261	0.323	0.619	0.609	0.439	0.303	0.682	1.3	3.3
r intcpt	0.049	0.404	0.067	0.766	0.057	0.470	0.786	0.2	3.3
r intcpt + slope	0.001	0.261	0.009	0.659	0.001	0.145	0.494	0.2	3.3
TP SMOOTHS									
r smooth, tp, 3 bs	0.017	0.169	0.037	0.538	0.049	0.082	0.378	0.8	3.4
r smooth, tp, 5 bs	0.013	0.127	0.033	0.341	0.053	0.092	0.402	9.4	3.6
r smooth, tp, 10 bs	0.007	0.116	0.042	0.100	0.053	0.114	0.461	38.8	4.9
r smooth, tp, 15 bs	—	—	0.037	0.069	—	—	—	—	—
r smooth, tp, 20 bs	—	—	0.037	0.074	—	—	—	—	—
r smooth, tp, gam.check	0.015	0.128	0.034	0.405	0.053	0.094	0.403	3.2	3.6
CR SMOOTHS									
r smooth, cr, 3 bs	0.029	0.134	0.047	0.545	0.058	0.119	0.459	0.0	3.4
r smooth, cr, 5 bs	0.041	0.084	0.047	0.359	0.058	0.079	0.377	5.6	3.6
r smooth, cr, 10 bs	0.047	0.077	0.056	0.129	0.058	0.082	0.381	50.6	4.3
r smooth, cr, 15 bs	—	—	0.060	0.086	—	—	—	—	—
r smooth, cr, 20 bs	—	—	0.062	0.083	—	—	—	—	—
r smooth, cr, gam.check	0.036	0.093	0.050	0.440	0.058	0.089	0.398	0.4	3.6
AR1	0.102	0.252	0.187	0.336	—	—	—	0.0	5.7

Table 1: Simulation set 1 – type I error rates for the real F2 data and the real pitch data, and type I error rates, power, memory usage (MB) and processing time (s) for the simulated F2 contours. The type I error rates and power figures are presented separately for parametric / smooth difference terms. Power is only shown for the simulated formants. Statistics for random smooths with more than 10 basis functions are only given for the pitch data, as these smooths cannot be fitted to the formant data. When a model component yields an unacceptable type I error rate (higher than 0.1), both the type I error rate and power are shown in grey. The methods represented by the different rows are illustrated in the same order in `example_code/set.1-real.f2.md`, `example_code/set.1-pitch.md` and `example_code/set.1-sim.f2.md`.

ted to the same data set. Type I error rates that are higher than 0.1 and corresponding power estimates are shown in grey. This allows us to quickly dismiss methods that yield catastrophically high rates of false positives.⁵ It is useful to first concentrate on the type I error rates and ignore power, as power estimates are not particularly informative in cases where the type I error rate is not well calibrated. Note also that type II error simulations from which the power figures are calculated were only run for the simulated formant data.

Looking at the distribution of figures in grey quickly reveals a few general patterns. First of all, the top row

⁵We do not mean to suggest that type I error rates lower than 0.1 are necessarily acceptable. Since the p -values reported by GAM(M)s are approximate (Wood, 2013) small deviations from the nominal type I error rate of 0.05 are expected. The (somewhat arbitrary) cut-off of 0.1 serves simply to rule out modelling strategies that are clearly inappropriate, and allow for more nuanced discussion of methods with error rates that are closer to nominal.

shows that models with no random effects or AR1 components all produce heavily inflated type I error rates. Second, the parametric terms produce low type I error rates for any model with random effects. This is easily seen by scanning the columns labelled ‘param. type I’, which are almost all in black (with the exception of the top and the bottom rows, which do not include random effects). This is because even simple random intercepts are capable of adjusting for between-contour differences in height. Third, focusing on the top block, random intercepts and slopes do not really improve on models without random effects in terms of type I error rates for the smooths. Again, this is by no means surprising, as neither random intercepts nor random slopes can capture the non-linear patterns in the formant and pitch contours investigated here. Finally, models with AR1 error components but without random effects (in the last block) do not perform particularly well: while their error rates are consistently lower than those of the

vanilla models in the top row, they are still far too high. A likely explanation is that AR1 error models are too simple to fully capture the complex patterns of residual autocorrelation in these contours. Higher-level autocorrelation functions (e.g. the `corARMA()` autocorrelation-moving average correlation structure) would likely do better, but – at the moment – these cannot be fitted efficiently using the `bam()` function.

Let us now look at how random smooths fare with respect to dependencies within contours. As noted above, the error rates for parametric terms are all sufficiently low for models with random effects, so our focus will be on error rates for smooth terms. For now, we will also set aside the results in the `gam.check()` rows. The smooth error rates exhibit some interesting variation across the data sets. For the real formants, `tp` smooths do not perform particularly well, with smooth error rates remaining above 0.10 regardless of the number of basis functions. The `cr` smooths in the third block yield better calibrated error rates, especially when 10 basis functions are used. For the pitch data, `tp` smooths do especially well, although the error rates for some `cr` smooths are also under 0.10 – but only when 15 or more basis functions are used. Finally, the simulated formant contours behave similarly to the real ones with respect to `cr` smooths, but the results for `tp` smooths are a little different in that some of these (with a lower number of basis functions) also produce acceptable results. These differences across data sets highlight the fact that there is no one-size-fits-all strategy when it comes to random smooths: different contour types seem to require different types and numbers of basis functions. This, of course, raises the question of how the analyst can decide on the best option for a given data set.

One potential (partial) solution that we examined in the simulations is to use the `gam.check()` function to find the optimal number of basis functions for random smooths. The results of this investigation – summarised in the `gam.check()` row of the table – are mixed. For the real and simulated formant data, `gam.check()` produces error rates that are close to those of the more complex random smooths. However, it performs rather poorly on the pitch data, producing smooth error rates that are far outside the acceptable range. Looking at the actual `k` values chosen by `gam.check` (not shown here) reveals that it only very rarely allowed the models to move beyond the starting value of 4. It therefore appears that `gam.check` sometimes erroneously suggests that a given level of complexity in the random smooths is sufficient even when it is not. Note that model diagnostics based on the residuals also do not reveal obvious issues with these overly simple models. The discussion sec-

tion will suggest a few potential ways of dealing with this issue and finding the right random smooths for a given data set.

Looking at the power figures adds little to our observations so far: models of the simulated formant data that produce acceptable type I error rates all seem to have roughly the same power. Processing times are also fairly constant across the different models, with slightly longer times required for random smooths with higher numbers of basis functions. However, the memory usage statistics clearly show a non-linear increase in resource requirements as a function of model complexity: for instance, random `cr` smooths with 10 basis functions require nearly 10 times as much memory as random smooths with 5 basis functions. This underlines the fact that the success of complex random smooths at dealing with within-contour dependencies comes at a high computational cost.

A reviewer suggests that combining an AR1 component with simpler random smooths (e.g. `k = 3` or `5`) may allow computationally less costly models to achieve the same results as models with complex random smooths. An additional set of simulations showed this not to be the case: while models without random smooths showed some modest improvements in terms of error rates, the error rates for models with random smooths remained essentially the same.

3.2. Set 2: Across-item effects

3.2.1. Methods

The simulations in the previous section are somewhat contrived in that contours in real-life dynamic phonetic data sets are almost always grouped into higher-order units such as words or speakers. This section therefore attempts to answer the following question: what random structures should be used to capture dependencies within such higher-order units? We use dependencies within speakers to explore this issue.

For the real formant data, the data set is generated by randomly sampling 20 contours from each of the 30 speakers (i.e. 600 contours altogether). For the pitch data, we use all contours from all speakers (but each simulated data set is different in that the category labels are assigned randomly). For the simulated formant data, 50 speakers are generated with 40 contours per speaker. The parameters of the contours are sampled from speaker-specific distributions, which in turn are controlled by a set of hyperparameters (shown in detail in Appendix A.2). As shown by the example contours in figure 7, each data type is characterised by substantial dependencies within speakers, such that there is less

variation in contour shapes within individual speakers than there is in the data set as a whole. For all data sets, half of the speakers were assigned to group *A* and half to group *B*. All contours representing a given speaker belong to the same group. In other words, this set of simulations focuses on across-item (specifically, across-speaker) effects.

Since there can be no variation in the size or shape of the effect within speakers, the random structures tested here are conceptually equivalent to by-speaker random intercepts in conventional linear mixed effects models. However, since the curves underlying the speakers vary not only in overall height, but also in shape, random intercepts are unlikely to be sufficient to capture within-speaker dependencies. The simulations therefore compare a range of model structures: no random structures; random intercepts by speaker; random intercepts plus slopes by speaker; by-speaker random smooths of varying complexity (3, 5 and 10 basis functions for the formant data and 3, 5, 10 and 15 basis functions for the pitch data); and by-speaker random smooths selected by the `gam.check()` function. The example model formulae provided for the previous set of simulations apply for these random effects as well, although the grouping factor in this case is `speaker`, not `trajectory`. In addition, all models for real data include an AR1 component to control for within-contour dependencies, as has been done in previous work (e.g. Wieling et al. 2016; Sóskuthy et al. 2018; Carignan et al. 2020). Although by-contour random smooths may be more adequate for this purpose, fitting a random smooth with many hundreds of levels would consume a significant amount of memory and take a long time, making it unsuitable for simulations that need to be run many thousands of times. Since the key dependencies in this data set are within speakers, missing patterns at the level of individual contours may not impact statistical inferences too heavily (see section 4 for more discussion). Models for the simulated data did not include an AR1 component as this data was simulated without autocorrelation in the noise-generating process.

The data and models used in this section are illustrated in `example_code/set.2-real.f2.md`, `example_code/set.2-pitch.md` and `example_code/set.2-sim.f2.md`.

3.2.2. Results

Table 2 presents the results of the second set of simulations. The overall pattern is similar to the findings of the previous section. Parametric type I error rates are well controlled by all random structures for all three data types, while tests based on the smooth term pro-

duce much more variable error rates. The latter set of figures also shows that models without random smooths (in the top block) are simply not capable of capturing by-speaker trends.

Turning now to the different random smooth specifications and focusing on the type I error rates for the smooth terms, we also see results that are broadly in line with those from the first set of simulations, though with some key differences. By-speaker dependencies in both the real formant data and the simulated formant data are best captured by `cr` random smooths, though – unlike in the previous set – a smooth with 5 basis functions seems sufficient. Similar to the first set of simulations, `tp` smooths also have some limited success in controlling type I error rates, but are slightly edged out by `cr` smooths. The results for the pitch data diverge from those in the previous set: only complex `cr` smooths produce acceptable type I error rates. It bears mentioning that the best-performing models do a great job at bringing within-speaker dependencies under control, with type I error rates that are very close to nominal (around 0.04-0.06), even more so than in the previous section. This is especially encouraging given that the data under investigation in this section are much more similar to real data sets.

Model selection based on `gam.check()` runs into the same issues seen in the previous section. While it is more or less successful at controlling error rates for the real and simulated formant data, it fails catastrophically for the pitch data. As before, inspection of the `k` values chosen by `gam.check()` reveals a general reluctance to move beyond low values even when they are insufficient to account for dependencies within groups.

The power figures for the simulated formant data are all roughly the same for models that produce acceptable type I error rates. Memory usage and processing times are also broadly in line with what we observed in the previous section, with complex models requiring substantially more resources.

3.3. Set 3: Within-item effects

3.3.1. Methods

The third set of simulations is similar to the previous one in that the contours are organised into speakers, but this time each speaker is represented by 20 contours from group *A* and 20 contours from group *B*. In other words, group membership varies within speakers. The data set is generated in such a way that the size and the exact realisation of the differences also varies systematically across speakers. These simulations are only run for the simulated data.

MODEL COMPONENT	REAL FORMANTS		REAL PITCH		SIMULATED FORMANTS			MEM.	TIME
	PARAM.	SMOOTH	PARAM.	SMOOTH	PARAM.	SMOOTH			
	TYPE I	TYPE I	TYPE I	TYPE I	TYPE I	TYPE I	POWER		
NO SMOOTHS									
no components	0.601	0.622	0.718	0.941	0.813	0.760	0.949	0.0	3.2
r intcpt	0.060	0.660	0.060	0.960	0.058	0.780	0.953	0.2	3.2
r intcpt + slope	0.004	0.365	0.037	0.924	0.000	0.286	0.734	0.2	3.3
TP SMOOTHS									
r smooth, tp, 3 bs	0.042	0.137	0.039	0.735	0.053	0.089	0.501	0.5	3.5
r smooth, tp, 5 bs	0.037	0.087	0.036	0.436	0.059	0.096	0.536	10.0	3.6
r smooth, tp, 10 bs	0.037	0.179	0.038	0.137	0.060	0.191	0.696	25.7	3.9
r smooth, tp, 15 bs	—	—	0.037	0.211	—	—	—	—	—
r smooth, tp, gam.check	0.042	0.091	0.039	0.397	0.058	0.094	0.538	0.0	4.2
CR SMOOTHS									
r smooth, cr, 3 bs	0.049	0.105	0.049	0.812	0.062	0.202	0.662	0.0	3.5
r smooth, cr, 5 bs	0.057	0.067	0.057	0.451	0.059	0.063	0.481	13.6	3.6
r smooth, cr, 10 bs	0.061	0.063	0.060	0.064	0.059	0.085	0.509	30.9	4.0
r smooth, cr, 15 bs	—	—	0.060	0.042	—	—	—	—	—
r smooth, cr, gam.check	0.054	0.070	0.056	0.463	0.060	0.093	0.527	0.6	4.2

Table 2: Simulation set 2 – type I error rates for the real F2 data and the real pitch data; and type I error rates, power, memory usage (MB) and processing time (s) for the simulated F2 contours. The methods represented by the different rows are illustrated in the same order in `example_code/set_2-real_f2.md`, `example_code/set_2-pitch.md` and `example_code/set_2-sim_f2.md`.

This type of within-speaker dependency is implemented by setting up separate distributions for groups *A* and *B* within each speaker and sampling contour-level parameter values from these distributions. The means of the group *A* and *B* speaker-level distributions are sampled in pairs from bivariate normal distributions with positive correlation parameters. In more intuitive terms, group *A* and group *B* contours representing the same speaker are generated by separate underlying curves, but each speaker still has ‘their own identity’ that is visible across both groups (see the example plot in `example_code/set_3-sim_f2.md`). As before, group *B* contours have a slightly different underlying shape for type II simulations (see Appendix A.3 for more detail).

Speaker-specific patterns in the size and shape of the group *A* vs. *B* contrast can be captured by using random structures that are conceptually equivalent to by-speaker random intercepts and slopes in conventional linear mixed effects models (cf. section 1.2.3). All random smooths mentioned below are of type *cr* and they are always fitted at three different levels of complexity (3, 5 and 10 basis functions). Five different types of random structures are tested:

1. Random smooths by speaker as in the previous set. These random smooths cannot capture variation

in the implementation of the group effect across speakers.

2. A combination of two random effects: random smooths by speaker plus a random slope by speaker over group, that is `s(speaker, group, bs="re")` using `mgcv` notation. These random effects can capture distinct contour shapes across speakers as well as by-speaker variation in the vertical size of the group *A* vs. *B* contrast.
3. Item $\times$ effect random smooths, with separate smooths fitted to each combination of speaker and group:

```
s(measurement.no, speakerGroup,
  bs="fs", m=1)
```

This type of random effect can capture all the by-speaker patterns described above as well as variation in the exact realisation of the group *A* vs. *B* contrast. Note, however, that it is still suboptimal in that it treats group *A* and *B* tokens from the same speaker as independent of each other.

4. Item-by-effect random smooths:

```
s(measurement.no, speaker,
  bs="fs", m=1, by=group.factor)
```

We expect this type of random effect to behave very similarly to item $\times$ effect random smooths.

5. Random reference/difference smooths:

```
s(measurement.no, speaker,
  bs="fs", m=1) +
s(measurement.no, speaker,
  bs="fs", m=1, by=group.ordered)
```

Given that this random structure separates each level of the hierarchical variable into a reference smooth and a difference smooth, it should be able to produce shrinkage for both and thereby produce type I error rates and power that are well-calibrated.

The data and models used in this section are illustrated in `example_code/set_3-sim_f2.md`.

3.3.2. Results

The results are shown in table 3. As expected, models with only by-speaker random smooths (first block) produce unacceptably high type I error rates for both parametric and smooth terms, regardless of the complexity of the smooth. This is because they are unable to represent differences between the groups within speakers. The models with random smooths by speaker and random slopes by speaker over group (second block) do slightly better. They manage to reduce the parametric type I error rate to nominal levels, but the error rates for the smooth difference terms remain largely unaffected.

The last three blocks yield more interesting results. The item $\times$ effect and the item-by-effect smooths produce essentially the same error rates and power. The type I error rates are anomalously low for both methods, which is likely a reflection of low power: the tests do not return significant results even in cases where a reasonably large difference arises by chance. This is confirmed by the power figures, which are low for item $\times$ effect and item-by-effect smooths when compared to the random reference/difference smooths in the last block.

The reason for this low power is simply that these models do not recognise the connection between contours representing different levels of the fixed effect variable within the same speaker: they treat them as if they were completely independent. In other words, there is no shrinkage on the random within-speaker differences between groups. This results in a situation where some of the systematic variation that should be captured by the fixed effect terms in the model (i.e. the parametric term and the smooth difference term) is, instead, hijacked by the random effect term. A similar effect is observed when a conventional mixed effect model with the structure

```
outcome ~ group + (1 + group | speaker)
```

is, instead, run as

```
outcome ~ group + (1 | speakerGroup)
```

In a simple simulation (included at the end of `example_code/set_3-sim_f2.md`), the first model structure showed a close-to-nominal type I error rate of 0.06 and a power of 0.54, while the second model structure produced a type I error rate of 0 and a power of 0.006.

The random reference/difference smooths shown in the last block produce nominal type I error rates and substantially higher power than the previous methods. These results clearly favour this last method, and suggest that it is the right tool for dealing with within-speaker dependencies in the presence of effects that vary within speakers. Note also that even the simplest random smooths produce perfectly acceptable results within this block. Importantly, reference/difference random smooths are computationally no more costly than item $\times$ effect smooths or item-by-effect smooths, as shown by their memory usage and processing time.

3.4. Set 4: Significance testing

3.4.1. Methods

The previous simulations all relied on pairs of significance tests that test whether the parametric and smooth difference terms are different from zero. However, hypotheses about contours tend not to separate these effects: the main research question is often whether two sets of contours differ in any way, not whether they differ specifically in height or in shape. Since the method above requires two tests of significance for a single question, it may lead to increased rates of false positives. The current section therefore looks at a broader range of methods for significance testing.

Most of the methods have already been described in detail in section 1.2.4, so they are only reviewed briefly here. The first of these is the approach described above: performing separate significance tests at $\alpha = 0.05$ for both the parametric and smooth terms, and declaring a significant overall difference if either of them is significant (henceforth the ‘parametric/smooth method’). Since this method is almost certainly anti-conservative, the same tests are also performed using Bonferroni correction, lowering the α -level to 0.025. Another solution is to use model comparison based on a likelihood ratio test, simultaneously excluding the parametric and smooth difference terms from the nested model (see section 1.2.4 for example model formulae). The p -value for the model comparison is generated using `compareML()`

R SMOOTH TYPE	PARAM.		SMOOTH		MEMORY	TIME
	TYPE I	POWER	TYPE I	POWER		
SIMPLE RANDOM SMOOTHS						
cr, 3 bs	0.422	0.821	0.376	0.907	10.1	2.9
cr, 5 bs	0.425	0.822	0.380	0.909	24.6	3.1
cr, 10 bs	0.426	0.823	0.380	0.909	67.0	3.8
RANDOM SMOOTHS + SLOPES						
cr, 3 bs	0.054	0.440	0.379	0.908	11.7	3.1
cr, 5 bs	0.054	0.440	0.383	0.909	31.5	3.3
cr, 10 bs	0.054	0.440	0.383	0.910	83.4	4.2
ITEM $\times$ EFFECT						
cr, 3 bs	0.000	0.019	0.008	0.228	20.2	3.9
cr, 5 bs	0.000	0.018	0.000	0.057	73.8	4.6
cr, 10 bs	0.000	0.015	0.000	0.060	238.0	8.6
ITEM-BY-EFFECT						
cr, 3 bs	0.000	0.025	0.008	0.228	13.0	4.8
cr, 5 bs	0.000	0.024	0.000	0.055	77.5	5.0
cr, 10 bs	0.000	0.022	0.000	0.059	176.1	9.2
REFERENCE/DIFFERENCE						
cr, 3 bs	0.060	0.456	0.049	0.570	14.8	4.4
cr, 5 bs	0.060	0.455	0.051	0.574	86.4	4.9
cr, 10 bs	0.059	0.454	0.050	0.573	185.4	9.4

Table 3: Simulation set 3 – type I error rate, power, memory usage (MB) and processing time (s) for the simulated F2 data. The five different random smooth types are separated into labelled blocks. The methods represented by the different rows are illustrated in the same order in `example_code/set_3-sim_f2.md`

from version 2.3 of `itsadug`. This method requires fitting via the much slower ML method. Results from inadequate fREML-based comparisons are also presented below to demonstrate that the p -values from such tests cannot even be considered approximate.

The three alternative solutions from section 1.2.4 are also included. The first is to use the conditional AIC score to compare models fitted with fREML. This method does not follow the principles of null-hypothesis significance testing, and therefore cannot control type I error rates in the same way as other tests do. The figures reported for AIC in the next section are based on a simple criterion: if the AIC of the full model is lower than that of the nested model, we declare a positive result; if there is no difference or the nested model has the lower value, we declare a negative result. The type I error rates and power figures based on this rule are marked typographically to set them apart from the error rates for significance tests. The second workaround is to use a single binary smooth that incorporates both parametric and smooth differences, and test whether this term is different from zero (this information can be extracted from

the model summary without performing model comparison). The third solution is to use an fREML-based model comparison where the parametric and smooth difference terms both receive extra penalties, effectively turning them into random effects. Since fREML-based comparisons between nested models with different random effects are valid, this method may yield reasonable error rates, although using random terms to estimate a replicable effect is technically incorrect.

The simulations also test two visual methods. One of these is based on checking for overlap between the 95% confidence intervals for the predicted contours for A and B . The issue of translating pointwise differences into overall significance is dealt with by using a cut-off criterion: if the confidence intervals are separated along a region of the curve whose length (in percentage) is equal to or greater than this criterion, the difference is deemed significant. The second visual method is to check whether the 95% confidence interval around the estimated difference curve excludes 0. Again, a cut-off criterion is used to estimate overall significance. Both of the visual methods rely on predictions and confidence

intervals from models estimated via fREML.

As the focus is now on significance testing, not random structures, we simply use the data sets from the second set of simulations in section 3.2. The models we fitted use the best-performing random structures from section 3.2.

The data and models used in this section are illustrated in `example_code/set_4-real_f2.md`, `example_code/set_4-pitch.md` and `example_code/set_4-sim_f2.md`.

3.4.2. Results

The results for the non-visual methods are presented in table 4. The first block shows results from significance tests whose results are included in standard model summaries generated via the `summary()` command, and which therefore do not require costly model comparisons. The overall pattern of error rates is similar across the three data sets. As predicted, the parametric/smooth method leads to inflated type I error rates when the tests are performed using the usual $\alpha = 0.05$ criterion. These rates are roughly twice as high as those for better calibrated methods. The Bonferroni correction helps to lower this rate to a more acceptable level. Binary smooths produce essentially the same well-calibrated type I error rates, but with higher power than the Bonferroni-corrected tests.

Turning now to the likelihood ratio tests, model comparison with ML produces excellent type I error rates, though slightly lower power than the binary smooth test. Note, however, that models fitted via ML are by far the least efficient, using substantially more memory and taking longer to converge than any other model. Since the relationship between the size of the data set and memory usage / processing time is non-linear, this difference is expected to be even greater for models fitted to larger data sets.

Unsurprisingly, the output of likelihood ratio tests based on models estimated via fREML is unusable when the full and nested models differ in their fixed effects: the type I error rates in the second row are sometimes anomalously low, and sometimes unacceptably high, suggesting that they are essentially randomly distributed. Model comparisons relying on penalised difference terms do not seem to fare much better.

As noted above, the AIC values in the table cannot be taken at face value, as AIC is not designed to control type I/II error rates. Nonetheless, the figures point in the right direction, with more or less acceptable type I error rates between 0.05–0.15 for all data sets, and a power of 0.286 for the simulated data. The simple cut-off rule used to evaluate AIC also obscures some

finer patterns in the data: although both the real formant and pitch data show a relatively high proportion of false positives, the size of the AIC differences for these models is typically very close to 0. If the cut-off for AIC differences is moved to -0.3 instead of 0 (i.e. the AIC value for the full model needs to be at least 0.3 lower than the AIC value for the nested model), the error rates for the formant and pitch data decrease to 0.0483 and 0.0361, respectively. Since AIC provides a graded scale for comparing models, AIC differences that are so close to 0 should not be interpreted as strong evidence for the full model over the nested one. While an AIC-based approach may be appropriate in terms of type I error rates, it seems less convincing in terms of its power. For the simulated data with a true underlying effect, the full model was selected in a mere 29% of all iterations, which is much lower than the corresponding power figures for model summary and LRT-based tests (and this figure would be even lower for a more stringent cut-off criterion). In any case, AIC is difficult if not impossible to compare to the other tests examined here, and one should not draw general conclusions based on power figures from a single set of simulations. On the whole, AIC paired with fREML estimation appears to be a promising strategy, but one that (i) is fundamentally different from the null-hypothesis significance testing methods presented above and (ii) may need some further investigation to better understand its performance in terms of power.

Figure 8 presents type I error rates and power for the visual methods. As explained above, visual testing requires a decision about the cut-off point for overall significance defined in terms of the relative width of the region where the 95% confidence intervals show no overlap / exclude zero. The graphs in figure 8 illustrate type I error rates and power as a function of cut-off values between 1% and 100% (shown along the horizontal axis). Let us first look at the graph on the right, which summarises the results for the simulated F2 contours. The graph shows error rates and power for both visual tests. For both tests, power and the type I error rate decrease as the cut-off criterion is increased. This is because a higher cut-off criterion corresponds to a more stringent test: the difference needs to be realised across a wider section of the curves. As predicted, the overlap-based test has substantially lower power than the difference-based test, and is therefore not discussed further. Type I error rates for the difference-based test are mostly under 15%, and there is a fairly broad range of values between 30–50% where it produces reasonably high power and close-to-nominal type I error rates. Note, however, that the success of this method hinges on the somewhat ar-

SIGNIFICANCE TEST	R. FORMANTS	R. PITCH	S. FORMANTS		MEMORY	TIME
	TYPE I	TYPE I	TYPE I	POWER		
MODEL SUMMARY						
par/smooth	0.114	0.112	0.130	0.773	6.5	2.0
par/smooth bonf.	0.064	0.062	0.085	0.685	6.5	2.0
binary smooth	0.066	0.063	0.090	0.753	0.1	2.0
LIKELIHOOD RATIO TESTS						
ML	0.029	0.042	0.062	0.675	87.5	7.0
fREML	1.000	0.004	1.000	1.000	6.9	3.5
fREML trick	0.008	0.020	0.983	0.999	0.0	3.4
AIC, fREML	*0.138	*0.146	*0.052	*0.286	6.9	3.5

Table 4: Simulation set 4 – type I error rate for the real formant data and the pitch data, and type I error rate, power, memory usage (MB) and processing time (s) for the simulated formant data. The label ‘trick’ refers to the extra penalty workaround, i.e. turning fixed effects into random effects. The error rates for the AIC-based comparisons are marked by asterisks and italicised to highlight the fact that AIC is not designed to control type I/II error rates. The methods represented by the different rows are illustrated in the same order in `example_code/set_4-real_f2.md`, `example_code/set_4-pitch.md` and `example_code/set_4-sim_f2.md`

bitrary choice of the cut-off value: with cut-off values below 30%, type I error rates become rather high, and with cut-off values over 50%, power falls quickly. Essentially the same results are observed for the real F2 and pitch contours in the panel on the left, although there are no power figures for these simulations. While a cut-off value of 40% seems to strike a good balance between type I and II errors for all the data types examined here, there is no guarantee whatsoever that the error rates would vary in the same way for other types of data. Therefore, it is not advisable to rely solely on visual information for significance testing without using some of the other, more consistent methods examined in the current section.

4. Discussion and recommendations

Let us first review the main findings focusing on their implications for the use of GAMMs in phonetics. This is followed by a more general discussion that covers the limitations of the work presented in this paper, and makes a few more general recommendations as well as suggestions for finding the optimal random structure for a given data set.

The first set of simulations (section 3.1) demonstrates that dependencies within contours are best captured using random smooths. The type of random smooth to be used (tp *versus* cr) and the optimal number of basis functions vary across data sets. The simulations also suggest that relying solely on `gam.check()` to determine the necessary number of basis functions does not work for all data types. While there are no hard-and-fast

rules for making decisions about the type and number of basis functions, we do suggest some general strategies for comparing models with different random smooths at the end of this section.

Models without random components, models with random intercepts only and models with random slopes in lieu of random smooths all produce high type I error rates. This is because they are not able to adequately capture non-linear patterns in contour shapes. Models with AR1 autoregressive error components produce lower error rates than models without smooths, but the type I error rates are still rather high. It is more than likely that more complex models of residual autocorrelation (autoregressive models with a higher order such as AR2 or AR3, or autoregressive moving average models) would do a better job at capturing within-contour dependencies. Unfortunately, at the moment these can only be fitted using the `gamm()` function in `mgcv`, not the significantly more efficient and more fully implemented `bam()` function that we have used in this paper.

The results from the second set of simulations suggest that appropriately specified random smooths are adequate for handling higher-level dependencies when testing for across-item effects. As with the first set of simulations, the precise details of the best-performing smooths vary across data sets, and the `gam.check()` function is not always sufficient for making decisions about these details.

The third set of simulations shows that random reference/difference smooths perform well in scenarios with higher-level dependencies and within-item effects. The same set of simulations suggests that using simple ran-

Figure 8: Type I error rates (red lines) and power (blue lines) for visual significance tests plotted against the significance cut-off criterion. The dashed lines represent tests based on overlap between separate confidence intervals, while the solid lines represent tests based on the confidence interval around the difference curve. All error rate curves are labelled for convenience. The figures for the real F2 and pitch data are plotted on the left, while the simulated F2 data on the right. The dotted grey line marks $\alpha = 0.05$. The visual methods evaluated in these plots are illustrated in section 7 of `example_code/set_4-f2.md` and `example_code/set_4-pitch.md`.

dom smooths or combining random smooths with random slopes may lead to anti-conservative estimates for this type of data. The reason for this is simply that these random effects cannot account for cases where a given predictor affects contour shapes differently across speakers (or words, items, etc.). Conversely, item $\times$ effect and item-by-effect random smooths (which behave essentially identically) are overly conservative: while they do not produce inflated type I error rates, they lead to reduced power. This is because different levels of a fixed effect variable within the same speaker are represented as completely separate. As a result, the model estimates across-speaker differences in the realisation of the fixed effect without imposing a penalty term on them, and the corresponding random smooths end up soaking up some of the variation that should, in fact, be assigned to the fixed effect terms. This issue does not arise when using random reference/difference smooths, which directly penalise the estimates for random individual deviations from a fixed effect.

The fourth set of simulations focuses on significance testing when the research question is not specifically about height or shape differences between groups *A* and *B*, but simply whether the two groups are different at all. Such situations are common in dynamic speech analysis.⁶ The simulations show that there are a number of reliable methods for significance testing in such

⁶This is not to say that differentiating between height and shape effects is not important. However, it is essential to be clear about what aspect of the contours one's hypotheses relate to, and to use significance tests that focus on that aspect directly. If there is no prior reason

cases: parametric/smooth tests with Bonferroni correction, model comparisons using ML estimation and binary difference smooths. Using model comparison based on AIC also produced promising results, though these comparisons must be interpreted differently from tests performed within the null hypothesis significance testing framework, and need further investigation to rule out potential issues with power. Model comparison with ML produced marginally lower type I error rates than the other methods, but at the cost of a massive increase in computational cost. Binary smooths, on the other hand, produced the highest power and did not prove to be any more resource-intensive than the rest of the tests.

None of the other methods for significance testing produced acceptable type I error rates. Parametric/smooth tests without correction for multiple comparisons roughly double the type I error rate, while likelihood-ratio tests combined with fREML lead to catastrophic results.

Visual methods are absolutely essential for communicating results from GAMMs, and they may also be useful for evaluating hypotheses that relate to differences between contours at a specific location. However, the simulation results confirm our initial suspicion that they are not appropriate as a standalone method for significance testing in the absence of specific hypotheses of the type mentioned above. Though they may produce acceptable type I error rates under certain conditions,

to expect that a difference will reside primarily in contour height or shape, it is better to use a single confirmatory test.

they require arbitrary decisions about the cut-off criterion that may not generalise across all data sets.

Before making any recommendations based on the findings above, it is important to make their potential limitations explicit. As with any simulation-based work, the results in this paper need to be viewed in the context of the specific data sets that were used, as these can have a significant impact on the findings. For instance, all the contours investigated here represent relatively short intervals of time and, although they show substantial variation, they are realised more or less consistently across speakers (cf. figure 7). It is possible that longer or more variable trajectories may behave differently with respect to some of the modelling strategies investigated here. Therefore, such data sets need further consideration before they are subjected to analysis via GAMMs. Nonetheless, the formant and pitch data examined here are representative of many contour types that feature in phonetic work. In addition, using resampling-based simulations with real data substantially increases the ecological validity of the work presented here.

Another limitation relates to model complexity. Typical models in dynamic speech analysis often include several predictors of interest and various interactions. Moreover, the simulations in this paper only test categorical contrasts, even though continuous variables (e.g. duration, speech rate, etc.) are very common in phonetic data. There is no reason to assume that the results do not generalise to more complex models and continuous predictors, but a more systematic investigation focusing on these scenarios may uncover further complications (e.g. related to concavity; see Tomaschek et al. 2018a for a discussion). It should also be noted that this paper addresses questions about different aspects of GAMMs in a modular fashion: for instance, within-contour dependencies are investigated separately from within-speaker dependencies. While this is necessary to make the key arguments of the paper easier to follow, it would also be interesting to see how different components of GAMMs interact. To give an example, would the performance of the simulations in set 2 change if we used by-contour random smooth instead of an AR1 error model to control for dependencies within contours?

In keeping with the point about the generalisability of the results above, we only make general recommendations about GAMMs when those recommendations are straightforwardly grounded in broader statistical arguments. These arguments are made explicit for each of the recommendations below. The list of recommendations will be followed by a few further suggestions, which, however, cannot be used in an entirely general

fashion.

1. Use appropriately specified by-item (i.e. by-speaker, by-word, etc.) random smooths to capture higher-level dependencies when testing for across-item effects. When the contours contain non-linear patterns, random intercepts and random slopes are not sufficient to capture these.
(Missing such non-linear patterns leaves within-item dependencies in the residuals, impairing the model's ability to appropriately moderate its inferences. See below for some suggestions on how to establish the optimal random smooth specification.)
2. Use appropriately specified random reference and difference smooths to capture higher-level dependencies when testing for within-item effects. Item $\times$ effect and item-by-effect random smooths may lead to overly conservative inferences, while simpler random structures may be anti-conservative.
(Using a combination of random reference and difference smooths penalises by-item estimates of deviations from a given fixed effect, making sure that the random smooths do not soak up too much of the variance.)
3. Model comparison with ML estimation, binary smooths and separate parametric/smooth tests with Bonferroni correction all provide well-calibrated tests of significance for differences across sets of contours. AIC-based model comparison for models fitted with fREML is a possible alternative, though the results cannot be interpreted in the same way as those of null hypothesis significance tests.
(These methods are all supported by theoretical and simulation-based work on GAMMs, e.g. Wood 2013, 2017; Wood et al. 2016. Likelihood ratio tests on fixed effects for models fitted with REML or fREML do not yield meaningful results.)
4. Use visual methods as a complement to other methods of significance testing. Do not use them as standalone tests of significance when there are no hypotheses about the location of the expected differences.
(Visual methods require a decision about whether two sets of curves differ along a sufficiently wide interval, and this decision is difficult to make in a principled manner.)
5. As with other mixed effects models, treat p -values with a certain amount of caution: values that are close to 0.05 (or whatever value α is set to) should be explicitly flagged when reporting the results of statistical tests based on GAMMs.

(Even for the best-performing models, type I error rates are often slightly higher than nominal. The approximate nature of p-values for smooth terms and likelihood ratio tests in the context of GAMMs is discussed in detail in Wood 2017, 305–315, who also shows that these approximations perform well.)

It is possible to make a few additional suggestions based on the simulations above. Though these suggestions are rooted in consistent patterns in the simulations, it is possible that they reflect unique aspects of the data we analysed. Therefore, they should not be applied uncritically to data that are substantially different from the F2 and pitch contours analysed here. A potentially useful strategy for ascertaining their validity for specific data sets is to use resampling-based type I error simulations as described at the end of this section.

The first set of simulations indicates that within-contour dependencies are best handled using appropriately specified by-contour random smooths. However, as noted above, these simulations are not particularly realistic in that they assume no higher-level dependencies beyond those within contours. In practical terms, this is much like fitting models to data from a single speaker (or a single word), which does not often happen in phonetics. When higher-level dependencies are also present, an AR1 error model seems to be sufficient for controlling within-contour dependencies without severely impacting error rates. This is reassuring given that this is also the strategy adopted in recent work in phonetics using GAMMs (Wieling et al., 2016; Wieling, 2018; Sóskuthy et al., 2018; Carignan et al., 2020). There is no indication in the current set of simulations that this strategy needs to be revised.

A number of earlier simulations (not included in this paper) investigated artificial pitch contours that were generated without random noise to emulate the effects of smoothing that is sometimes applied to such contours before statistical analysis. These simulations produced unacceptably high error rates regardless of what modelling strategy was followed. The likely reason for the failure of GAMMs to cope with such data is that their underlying assumptions are severely violated. Since the smooths estimated by GAMMs can fit noiseless contours very closely, there are hardly any residuals left in the data, and what is left is not normally distributed. The confidence intervals and p -values calculated by GAMMs assume the presence of random noise and also that this noise is normally distributed. They become unreliable when these assumptions are not met. In any case, the simulations in the current pa-

per also demonstrate that smoothing trajectories before GAMM analysis is not necessary: all the contours that we analysed contain substantial noise, and yet GAMM modelling strategies with acceptable error rates could be found for all of them. If the analyst has no choice but to work with completely smooth trajectories, it is not possible to use standard GAMMs. In such cases, a feasible alternative might be to use QGAMs (Fasiolo et al., 2020), which do not make the same distributional assumptions as conventional GAMs. QGAMs have recently been used in the analysis of tongue movement data collected via electromagnetic articulography (Tomaschek et al., 2018b).

The recommendations above mention ‘appropriately specified’ random smooths, which refers mainly to the number and type of basis functions (or k and bs for short, based on the notation used by `mgcv`). So far, little guidance has been given on what may count as appropriate and how this can be established from the data. The simulations in sets 1 and 2 suggest that using `gam.check()` to choose the value of k in a stepwise fashion (starting with a lower number and moving towards higher values) is not always reliable. This procedure often settles on random smooths with a smaller number of basis functions than necessary to achieve well-calibrated type I error rates. A reviewer suggests that a better strategy would be to avoid looking at the p -values generated by `gam.check()`, and focus on how close the EDF for the smooth is to k' (the upper limit of the EDF for that term; both of these figures are also reported by `gam.check()`). If these two figures are very close, this might indicate that the number of basis functions needs to be increased. Another option may be to use a similar stepwise process, but rely on AIC for comparing models with different numbers or types of basis functions. Some preliminary experimentation with this method shows promise: for the real formant data in set 2, AIC seems to correctly select cubic regression splines, and also rules out smooths with too few basis functions. However, it seems to have a tendency to select more complex random smooths even in cases where the error rates do not justify this complexity. It seems that further work is needed to evaluate the performance of AIC in this context.

Another possibility is to use resampling-based type I error simulations structured along the same lines as those presented in this paper. The online supplementary materials contain a detailed tutorial on how such simulations can be set up for specific data sets (`tutorial.md`). The key is to (i) choose a homogeneous, coherent and representative subset of the data, and (ii) fit the same models to resampled versions of these subsets with a

randomly assigned mock predictor many times in order to get an estimate of their type I error rates. The advantage of this method is that it provides direct access to type I error rates and that it is completely flexible in terms of the structure of the models to be compared. Thus, it does not limit comparisons to models with different types/numbers of random smooths, but also allows the analyst to look at completely different model structures (e.g. using AR1, random intercepts / smooths, or even a different function for fitting GAMMs such as `gamm()`). The drawback of such simulations is that they are computationally intensive. This can be mitigated by being careful in choosing a reasonably small subset of the data for simulations, and taking advantage of performance-oriented features of GAMMs (such as using multiple processor cores, discretising the data, etc.). In our experience, a reasonably powerful laptop computer can run 1000 iterations of such simulations overnight.

Before concluding the paper, it is important to highlight the importance of model checking in the context of GAMMs. The strategies proposed here by no means obviate the need for the careful checking of models along the lines described in Wieling (2018) and Chuang et al. (2020). It is essential to check that specific data sets and model specifications do not violate the underlying assumptions of GAMMs. The first step should be to use standard diagnostics such as autocorrelation plots or QQ-plots of the residuals to test for obvious issues. The implementation of GAMMs in `mgcv` offers many solutions for addressing distributional violations, such as using Gaussian location scale additive models for dealing with heteroskedascity (enabled by using the `family = gaulss()` option for `bam()`) or GAMMs with a scaled t distribution for residuals when they show a heavy tail (`family = scat()`). More severe violations of modelling assumptions might be possible to address by using non-parametric approaches such as QGAMs (Fasiolo et al., 2020).

The simulations in this paper demonstrate that generalised additive mixed modelling is a powerful and reliable tool for analysing dynamic speech data. We have evaluated a wide range of different modelling strategies for a varied set of different data types, and found a number of well-calibrated solutions in each case. Relying partly on the results of these simulations and partly on previous work and theoretical arguments, the paper has made a number of general recommendations for fitting GAMMs to phonetic contours. It has also suggested ways in which different GAMM modelling strategies can be compared and evaluated, and made some more specific suggestions for the types of data analysed here.

References

- Asano, Y., Gubian, M., 2018. “Excuse meeee!”: (Mis)coordination of lexical and paralinguistic prosody in L2 hyperarticulation. *Speech Communication* 99, 183–200.
- Baayen, R. H., 2008. *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge University Press, Cambridge.
- Baayen, R. H., Linke, M., in press. An introduction to the generalized additive model. In: Paquot, M., Gries, S. T. (Eds.), *A practical handbook of corpus linguistics*. Springer, Berlin.
- Baayen, R. H., van Rij, J., de Cat, C., Wood, S. N., 2016. Autocorrelated errors in experimental data in the language sciences: Some solutions offered by generalized additive mixed models. *arXiv preprint arXiv:1601.02043*.
- Baayen, R. H., Vasisht, S., Bates, D., Kliegl, R., 2017. The cave of shadows. addressing the human factor with generalized additive mixed models. *Journal of Memory and Language* 94, 206–234.
- Barr, D. J., Levy, R., Scheepers, C., Tily, H. J., 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language* 68 (3), 255–278.
- Boersma, P., Weenink, D., 2009. Praat: doing phonetics by computer (version 5.0.38) [computer program]. Version 5.1.17; Retrieved on 29/09/2009 from <http://www.praat.org/>.
- Burnham, K. P., Anderson, D. R., 2004. Multimodel inference: understanding AIC and BIC in model selection. *Sociological methods & research* 33 (2), 261–304.
- Cahill, N., Rahmstorf, S., Parnell, A. C., 2015. Change points of global temperature. *Environmental Research Letters* 10 (8), 084002.
- Cardoso, A., 2015. Variation in nasal-obstruent clusters and its influence on PRICE and MOUTH in Scouse. *English Language and Linguistics* 19 (3), 505–532.
- Carignan, C., Hoole, P., Kunay, E., Pouplier, M., Joseph, A., Voit, D., Frahm, J., Harrington, J., 2020. Analyzing speech in both time and space: Generalized additive mixed models can uncover systematic patterns of variation in vocal tract shape in real-time MRI. *Laboratory Phonology: Journal of the Association for Laboratory Phonology* 11 (1).
- Chuang, Y.-Y., Fon, J., Baayen, R. H., 2020. Analyzing phonetic data with generalized additive mixed models. *PsyArXiv*. URL <https://doi.org/10.31234/osf.io/bd3r4>
- Davidson, L., 2006. Comparing tongue shapes from ultrasound imaging using smoothing spline analysis of variance. *The Journal of the Acoustical Society of America* 120 (1), 407–415.
- Fasiolo, M., Wood, S. N., Zaffran, M., Nedellec, R., Goude, Y., 2020. Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 1–11.
- Fromont, R., Hay, J., 2008. ONZE Miner: the development of a browser-based research tool. *Corpora* 3 (2), 173–193.
- Gick, B., Campbell, F., Oh, S., Tamburri-Watt, L., 2006. Toward universals in the gestural organization of syllables: A cross-linguistic study of liquids. *Journal of Phonetics* 34 (1), 49–72.
- Gordon, E., Maclagan, M., Hay, J., 2007. The ONZE Corpus. In: Beal, J. C., Corrigan, K. P., Moisl, H. L. (Eds.), *Models and methods in the handling of unconventional digital corpora: Volume 2, Diachronic Corpora. Vol. 2*. Palgrave Macmillan, Basingstoke, Hampshire, pp. 82–104.
- Gubian, M., Torreira, F., Boves, L., 2015. Using functional data analysis for investigating multidimensional dynamic phonetic contrasts. *Journal of Phonetics* 49, 16–40.
- Haddican, B., Foulkes, P., Hughes, V., Richards, H., 2013. Interaction of social and linguistic constraints on two vowel changes in northern England. *Language Variation and Change* 25 (03), 371–403.
- Harrington, J., Kleber, F., Reubold, U., 2008. Compensation for coarticulation, /u/-fronting, and sound change in Standard Southern

- British: an acoustic and perceptual study. *Journal of the Acoustical Society of America* 123, 2825–2835.
- Harrington, J., Palethorpe, S., Watson, C., 2000. Monophthongal vowel changes in Received Pronunciation: An acoustic analysis of the Queen's Christmas broadcasts. *Journal of the International Phonetic Association* 30 (1-2), 63–78.
- Hastie, T., Tibshirani, R., 1990. Generalized additive models. Chapman & Hall/CRC, Boca Raton.
- Hillenbrand, J., Getty, L. A., Clark, M. J., Wheeler, K., 1995. Acoustic characteristics of American English vowels. *Journal of the Acoustical Society of America* 97, 3099–3111.
- Jones, D., 1918. An outline of English phonetics. Teubner, Leipzig.
- Kösling, K., Kunter, G., Baayen, H., Plag, I., 2013. Prominence in tri-constituent compounds: Pitch contours and linguistic theory. *Language and Speech* 56 (4), 529–554.
- Labov, W., 1994. Principles of Linguistic Change. Vol. 1: Internal Factors. Blackwell, Oxford.
- Labov, W., Ash, S., Boberg, C., 2005. The Atlas of North American English: Phonetics, phonology and sound change. Mouton de Gruyter, Berlin.
- Marra, G., Wood, S., 2012. Coverage properties of confidence intervals for generalized additive model components. *Scandinavian Journal of Statistics* 39 (1), 53–74.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., Bates, D., 2017. Balancing type I error and power in linear mixed models. *Journal of Memory and Language* 94, 305–315.
- Noiray, A., Wieling, M., Abakarova, D., Rubertus, E., Tiede, M., 2019. Back from the future: Nonlinear anticipation in adults' and children's speech. *Journal of Speech, Language, and Hearing Research* 62 (8S), 3033–3054.
- Oehlert, G. W., 2014. A few words about REML. Unpublished manuscript, accessed at <http://users.stat.umn.edu/~corbett/classes/5303/REML.pdf> on 30 June 2018.
- Peterson, G. E., Barney, H. L., 1952. Control methods used in the study of vowels. *Journal of the Acoustical Society of America* 24, 175–184.
- R Core Team, 2013. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org>
- Renwick, M. E., Stanley, J. A., 2020. Modeling dynamic trajectories of front vowels in the American South. *The Journal of the Acoustical Society of America* 147 (1), 579–595.
- Roessig, S., Mücke, D., Grice, M., 2019. The dynamics of intonation: Categorical and continuous variation in an attractor-based model. *PloS one* 14 (5).
- Scobbie, J. M., Stuart-Smith, J., Lawson, E., 2012. Back to front: a socially-stratified ultrasound tongue imaging study of Scottish English /u/. *Italian Journal of Linguistics* 24, 103–148.
- Slud, E., Stone, M., Smith, P. J., Goldstein Jr, M., 2002. Principal components representation of the two-dimensional coronal tongue surface. *Phonetica* 59 (2-3), 108–133.
- Sóskuthy, M., 2017. Generalised Additive Mixed Models for dynamic analysis in linguistics: a practical introduction. arXiv:1703.05339 [stat:AP].
- Sóskuthy, M., Foulkes, P., Hughes, V., Haddican, B., 2018. Changing words and sounds: the roles of different cognitive units in sound change. *Topics in Cognitive Science* 10, 787–802.
- Sóskuthy, M., Hay, J., Brand, J., 2019. Horizontal Diphthong Shift in New Zealand English. *Proceedings of the 19th International Congress of Phonetic Sciences (ICPhS)*, Melbourne, Australia.
- Strycharczuk, P., Scobbie, J. M., 2017. Fronting of Southern British English high-back vowels in articulation and acoustics. *The Journal of the Acoustical Society of America* 142 (1), 322–331.
- Talkin, D., 2015. REAPER: Robust epoch and pitch estimator. Available from: <https://github.com/google/REAPER>.
- Tomaschek, F., Hendrix, P., Baayen, R. H., 2018a. Strategies for addressing collinearity in multivariate linguistic data. *Journal of Phonetics* 71, 249–267.
- Tomaschek, F., Tucker, B. V., Fasiolo, M., Baayen, R. H., 2018b. Practice makes perfect: The consequences of lexical proficiency for articulation. *Linguistics Vanguard* 4 (s2).
- Turton, D., 2015. Determining categoricity in English /l/-darkening: A principal component analysis of ultrasound spline data. In: *Proceedings of the XVII International Congress of Phonetic Sciences*.
- van Rij, J., 2015. Overview GAMM analysis of time series data. <http://www.sfs.uni-tuebingen.de/~jvanrij/Tutorial/GAMM.html>. Accessed on 23/01/2017.
- van Rij, J., Wieling, M., Baayen, R. H., van Rijn, H., 2017. itsadug: Interpreting time series and autocorrelated data using gamms. R package version 2.3.
- Watson, C. I., Harrington, J., 1999. Acoustic evidence for dynamic formant trajectories in Australian English vowels. *The Journal of the Acoustical Society of America* 106, 458–468.
- Wieling, M., 2018. Analyzing dynamic phonetic data using generalized additive mixed modeling: A tutorial focusing on articulatory differences between L1 and L2 speakers of English. *Journal of Phonetics* 70, 86–116.
- Wieling, M., Tomaschek, F., Arnold, D., Tiede, M., Bröker, F., Thiele, S., Wood, S. N., Baayen, R. H., 2016. Investigating dialectal differences using articulatory data. *Journal of Phonetics* 59, 122–143.
- Wieling, M., Veenstra, P., Adank, P., Tiede, M., 2017. Articulatory differences between L1 and L2 speakers of English. *Proceedings of The 11th International Seminar on Speech Production*, Tianjin, China, October 16–19.
- Winter, B., Wieling, M., 2016. How to analyze linguistic change using mixed models, Growth Curve Analysis and Generalized Additive Modeling. *Journal of Language Evolution* 1 (1), 7–18.
- Wood, S., 2013. On p-values for smooth components of an extended generalized additive model. *Biometrika* 100 (1), 221–228.
- Wood, S., 2017. Generalized Additive Models: An Introduction with R, 2nd Edition. CRC Press, Boca Raton.
- Wood, S. N., Goude, Y., Shaw, S., 2015. Generalized additive models for large data sets. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 64 (1), 139–155.
- Wood, S. N., Li, Z., Shaddick, G., Augustin, N. H., 2017. Generalized Additive Models for gigadata: Modeling the UK Black Smoke Network daily data. *Journal of the American Statistical Association*, 1–12.
- Wood, S. N., Pya, N., Säfken, B., 2016. Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association* 111 (516), 1548–1563.

Appendix A. Simulation details

This section first presents the formulae for generating simulated formant contours, and then provides details about how the parameters of the contours were sampled for each type of simulation.

The following formula was used for generating individual simulated formant contours (where x corresponds to the time dimension, scaled to the range [0,1]; the rest of the parameters have been described in section 2):

$$f2 = \frac{\text{END} - \text{START}}{1 + \exp[-\text{TRANSITION}_{\text{slope}}(x - \text{TRANSITION}_{\text{horiz}})]} + \text{START} + \epsilon$$

ϵ is a Gaussian error term with $\mu = 0$ and $\sigma = \text{NOISE}$.

Appendix A.1. Set 1: Dealing with autocorrelation

The parameter values for the contours in the first set of simulations are sampled from the following distributions (normal distributions are parameterised using the mean and standard deviation). Note that the time dimension ranges from 0 to 1 in the actual simulations, not 0 to 100 as in the figures in the main text. This is also reflected in the parameters below, which generate the wrong contours if the time dimension is taken to be 0 to 100.

Set 1, F2 contours, type I

$$\begin{aligned} \text{START} &\sim \mathcal{N}(1300, 150) \\ \text{END} &\sim \mathcal{N}(1650, 150) \\ \text{TRANSITION}_{\text{horiz}} &\sim \mathcal{N}(0.6, 0.025) \\ \text{TRANSITION}_{\text{slope}} &\sim \mathcal{N}(15, 4) \\ \text{NOISE} &\sim \mathcal{N}(0, 40) \end{aligned}$$

Set 1, F2 contours, type II

$$\begin{aligned} \text{START} &\sim \mathcal{N}(1300, 150) \\ \text{END} &\sim \mathcal{N}(1600, 150) \\ \text{END}_2 &\sim \mathcal{N}(1700, 150) \\ \text{TRANSITION}_{\text{horiz}} &\sim \mathcal{N}(0.6, 0.025) \\ \text{TRANSITION}_{\text{slope}} &\sim \mathcal{N}(15, 4) \\ \text{NOISE} &\sim \mathcal{N}(0, 40) \end{aligned}$$

Appendix A.2. Set 2: Across-word effects

These simulations differ from the first set in that the parameter values for the contours come from separate distributions for each word, whose means are drawn from higher-level distributions. The parameter distributions are shown below. Word-level means are marked by the subscript w (a superscript is used when a subscript is already present). Contour-level parameters have the same names as before.

Set 2, F2 contours, type I

$$\begin{aligned} \text{START}_w &\sim \mathcal{N}(1300, 90) \\ \text{START} &\sim \mathcal{N}(\text{START}_w, 150) \\ \text{END}_w &\sim \mathcal{N}(1650, 90) \\ \text{END} &\sim \mathcal{N}(\text{END}_w, 150) \\ \text{TRANSITION}_{\text{horiz}}^w &\sim \mathcal{N}(0.6, 0.02) \\ \text{TRANSITION}_{\text{horiz}} &\sim \mathcal{N}(\text{TRANSITION}_{\text{horiz}}^w, 0.015) \\ \text{TRANSITION}_{\text{slope}}^w &\sim \mathcal{N}(15, 4) \\ \text{TRANSITION}_{\text{slope}} &\sim \mathcal{N}(\text{TRANSITION}_{\text{slope}}^w, 3) \\ \text{NOISE} &\sim \mathcal{N}(0, 40) \end{aligned}$$

Set 2, F2 contours, type II

$$\begin{aligned} \text{START}_w &\sim \mathcal{N}(1300, 90) \\ \text{START} &\sim \mathcal{N}(\text{START}_w, 150) \\ \text{END}_w &\sim \mathcal{N}(1600, 90) \\ \text{END} &\sim \mathcal{N}(\text{END}_w, 150) \\ \text{END}_2^w &\sim \mathcal{N}(1690, 90) \\ \text{END}_2 &\sim \mathcal{N}(\text{END}_2^w, 150) \\ \text{TRANSITION}_{\text{horiz}}^w &\sim \mathcal{N}(0.6, 0.02) \\ \text{TRANSITION}_{\text{horiz}} &\sim \mathcal{N}(\text{TRANSITION}_{\text{horiz}}^w, 0.015) \\ \text{TRANSITION}_{\text{slope}}^w &\sim \mathcal{N}(15, 4) \\ \text{TRANSITION}_{\text{slope}} &\sim \mathcal{N}(\text{TRANSITION}_{\text{slope}}^w, 3) \\ \text{NOISE} &\sim \mathcal{N}(0, 40) \end{aligned}$$

Appendix A.3. Set 3: Within-word effects

The contour-level parameter values in these simulations are sampled from separate word-level distributions for groups *A* and *B*. The means of the word-level distributions are sampled in pairs from bivariate normal distributions with positive correlation parameters. The parameter distributions are shown below. The notation is the same as before, but the word-level subscripts also include an additional letter indicating group membership. Note that the bivariate distributions are parameterised using covariance matrices.

Set 3, F2 contours, type I

$$\begin{aligned}
\begin{bmatrix} \text{START}_{A,w} \\ \text{START}_{B,w} \end{bmatrix} &\sim \mathcal{N}\left(\begin{bmatrix} 1350 \\ 1350 \end{bmatrix}, \begin{bmatrix} 90^2 & 90^2 \times 7/8 \\ 90^2 \times 7/8 & 90^2 \end{bmatrix}\right) \\
\text{START}_A &\sim \mathcal{N}(\text{START}_{A,w}, 150) \\
\text{START}_B &\sim \mathcal{N}(\text{START}_{B,w}, 150) \\
\\
\begin{bmatrix} \text{END}_{A,w} \\ \text{END}_{B,w} \end{bmatrix} &\sim \mathcal{N}\left(\begin{bmatrix} 1650 \\ 1650 \end{bmatrix}, \begin{bmatrix} 90^2 & 90^2 \times 7/8 \\ 90^2 \times 7/8 & 90^2 \end{bmatrix}\right) \\
\text{END}_A &\sim \mathcal{N}(\text{END}_{A,w}, 150) \\
\text{END}_B &\sim \mathcal{N}(\text{END}_{B,w}, 150) \\
\\
\begin{bmatrix} \text{TRANSITION}_{\text{horiz}}^{A,w} \\ \text{TRANSITION}_{\text{horiz}}^{B,w} \end{bmatrix} &\sim \mathcal{N}\left(\begin{bmatrix} 0.6 \\ 0.6 \end{bmatrix}, \begin{bmatrix} .02^2 & .02^2 \times 7/8 \\ .02^2 \times 7/8 & .02^2 \end{bmatrix}\right) \\
\text{TRANSITION}_{\text{horiz}}^A &\sim \mathcal{N}(\text{TRANSITION}_{\text{horiz}}^{A,w}, 0.015) \\
\text{TRANSITION}_{\text{horiz}}^B &\sim \mathcal{N}(\text{TRANSITION}_{\text{horiz}}^{B,w}, 0.015) \\
\\
\begin{bmatrix} \text{TRANSITION}_{\text{slope}}^{A,w} \\ \text{TRANSITION}_{\text{slope}}^{B,w} \end{bmatrix} &\sim \mathcal{N}\left(\begin{bmatrix} 15 \\ 15 \end{bmatrix}, \begin{bmatrix} 4^2 & 4^2 \times 7/8 \\ 4^2 \times 7/8 & 4^2 \end{bmatrix}\right) \\
\text{TRANSITION}_{\text{slope}}^A &\sim \mathcal{N}(\text{TRANSITION}_{\text{slope}}^{A,w}, 3) \\
\text{TRANSITION}_{\text{slope}}^B &\sim \mathcal{N}(\text{TRANSITION}_{\text{slope}}^{B,w}, 3) \\
\\
\text{NOISE} &\sim \mathcal{N}(0, 40)
\end{aligned}$$

The type II simulations only differ from the type I simulations in that their $\text{END}_{A,w}$ and $\text{END}_{B,w}$ values are sampled from a different distribution:

Set 3, F2 contours, type II

$$\begin{aligned}
\begin{bmatrix} \text{END}_{A,w} \\ \text{END}_{B,w} \end{bmatrix} &\sim \mathcal{N}\left(\begin{bmatrix} 1630 \\ 1660 \end{bmatrix}, \begin{bmatrix} 90^2 & 90^2 \times 7/8 \\ 90^2 \times 7/8 & 90^2 \end{bmatrix}\right) \\
\text{END}_A &\sim \mathcal{N}(\text{END}_{A,w}, 150) \\
\text{END}_B &\sim \mathcal{N}(\text{END}_{B,w}, 150)
\end{aligned}$$