

Modeling systematicity and individuality in nonlinear second language development

Murakami, Akira

DOI:

[10.1111/lang.12166](https://doi.org/10.1111/lang.12166)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Murakami, A 2016, 'Modeling systematicity and individuality in nonlinear second language development: The case of English grammatical morphemes', *Language Learning*, vol. 66, no. 4, pp. 834-871.
<https://doi.org/10.1111/lang.12166>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is the peer reviewed version of the following article: Murakami, A. (2016), Modeling Systematicity and Individuality in Nonlinear Second Language Development: The Case of English Grammatical Morphemes. *Language Learning*, 66: 834-871, which has been published in final form at: <https://doi.org/10.1111/lang.12166>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Modeling systematicity and individuality in nonlinear L2 development: The case of English
grammatical morphemes

Abstract

The present paper introduces two sophisticated statistical modeling techniques that allow researchers to analyze systematicity, individual variation, and nonlinearity in L2 development. Generalized linear mixed-effects models can quantify individual variation and examine systematic effects simultaneously, and generalized additive mixed models allow analysts to examine systematicity, individuality, and nonlinearity within a single model. Based on a longitudinal learner corpus, this article illustrates their utility in the context of the L2 accuracy development of English grammatical morphemes. This paper discusses the strengths of each technique and the ways in which these techniques can benefit second language acquisition research, further highlighting the importance of accounting for individual variation in modeling L2 development.

Modeling systematicity and individuality in nonlinear L2 development: The case of English grammatical morphemes

Methodological Challenges in Modeling Individual Variation and Nonlinearity in SLA Systematicity and Individuality in SLA Research

For many years, second language acquisition (SLA) research has focused on revealing systematicity in second language (L2) development. The problem with searching for systematicity alone is that the identification of systematic patterns often necessitates statistical averaging, and averaging conceals individual patterns (Dörnyei, 2009). Indeed, in psychology, it is well-known that the averaged pattern can differ from the individual patterns that constitute the data (e.g., Heathcote, Brown, & Mewhort, 2000). Therefore, there has recently been a growing interest in SLA in understanding the performance of individual learners (van Geert & van Dijk, 2002; Verspoor, Lowie, & van Dijk, 2008).

However, studying individual variation requires appropriate analytical tools. Conventional statistical techniques in SLA, such as ANOVA, cannot appropriately disentangle between- and within-learner variability. With recent developments in statistical modeling, however, we can now model and analyze group-level and individual-level features simultaneously. The technique, called mixed-effects modeling, is now widely used in (applied) linguistics, including SLA (e.g., Kozaki & Ross, 2011; Tremblay, Derwing, Libben, & Westbury, 2011; see also Cunnings, 2012 and Linck & Cunnings, 2015). In most studies employing mixed-effects models, however, the technique has been used to control for individual differences in testing the significance of predictors or to study the sources of these differences. While this is certainly useful, mixed-effects models can also provide information about the amount and pattern of individual variation (Dingemanse & Doehrmann, 2013; Kliegl, Wei, Dambacher,

Yan, & Zhou, 2011). The present paper not only tests the effect of predictors but also focuses on individual variation disclosed by mixed-effects models, and through the process, it demonstrates that this technique can model systematicity and individuality simultaneously.

Nonlinearity in SLA

Another recent trend in SLA is that it emphasizes the process of learning than on the product (Atkinson, 2011). The learning process, however, is never linear. There is ample empirical evidence that demonstrates nonlinearity in L2 development. Perhaps the best-known nonlinearity in SLA is U-shaped development (e.g., Lightbown, 1983) and power-law development (e.g., DeKeyser, 1997; Ellis & Schmidt, 1998). In U-shaped development, accuracy is high at the beginning, and temporarily decreases before becoming high again. In power-law development, decrement in error becomes progressively smaller as the learner develops. Because power-law development covers the entire span of development and does not exhibit a systematic decrease in accuracy in the process, U-shaped and power-law development are mutually exclusive.

Despite the prevalence of nonlinearity in SLA, researchers are not fully equipped with appropriate statistical tools to analyze it. Classical statistical analysis is generally incapable of analyzing the learning process, including nonlinearity (Larsen-Freeman, 2011; see also Baayen, 2010b). For instance, if we want to investigate the effect of a treatment on the linguistic complexity of learners' writings while controlling for their proficiency, there is no sufficient evidence to assume a particular functional form between proficiency and linguistic complexity; thus it is not straightforward to statistically control proficiency. As in individual differences analysis, however, recent development in statistics allows analysts to model nonlinearity. Although the technique — the generalized additive model (Hastie & Tibshirani, 1990) — is new

in SLA, it has been used in other areas of (applied) linguistics including psycholinguistics (e.g., Baayen, 2010a; Baayen, Milin, Đurđević, Hendrix, & Marelli, 2011) and sociolinguistics (Wieling, Nerbonne, & Baayen, 2011). The present paper illustrates its utility in SLA.

Aim and Research Questions

This paper aims to introduce to the SLA community two types of statistical modeling techniques that take into account systematicity, individual variation, and nonlinearity. I do so by modeling the accuracy of L2 English grammatical morphemes. Grammatical morphemes were targeted in this exposition because their acquisition has been extensively studied in SLA since its early days (e.g., Dulay & Burt, 1973), and we already know much about the variables that affect their accuracy. This enables us to focus on what the new techniques can contribute to the field.

The status of the morpheme has been challenged as a functional unit of representation (e.g., Baayen et al., 2011; Ellis & Schmidt, 1998; Plaut & Gonnerman, 2000). Bybee (1985, 2010), for instance, demonstrates the essentially gradient nature of grammatical morphemes. The historical account further shows that the word, and not the morpheme, has been regarded as the smallest unit of a grammatical system (Blevins, 2013). Given the methodological focus of this paper, however, the issue is rather marginal.

This paper poses two demonstrative research questions:

1. How large is individual variation in the developmental pattern of morphemes?
2. Do their cross-sectional and longitudinal developmental patterns vary depending on morphemes and on whether learners' native languages (L1s) have an equivalent morpheme?

The background of Research Question 1 is that while SLA has identified prototypical developmental patterns, individual learners are hypothesized to exhibit a variety of learning

curves. This paper investigates the extent to which individual variation is observed in the developmental patterns of morphemes.

With regard to Research Question 2, in addition to individual variation, the present paper addresses the systematic effect of L1, which is known to affect nearly every aspect of L2 development (Jarvis & Pavlenko, 2007; Odlin, 1989), including grammatical morphemes (Luk & Shirai, 2009; Murakami & Alexopoulou, *in press*). It is not clear, however, how L1 influence emerges or changes during the acquisitional process (Jarvis, 2000).

The paper further investigates whether the developmental pattern differs across morphemes. Prior research often draws distinctions between free vs bound and verbal vs nominal morphemes (e.g., Brown, 1973; Goldschneider & DeKeyser, 2001). Slobin (1996) further distinguishes between the morphemes that encode language-independent concepts (e.g., number as expressed by plural *-s*) and those that encode language-dependent concepts (e.g., definiteness as expressed by articles). It is, therefore, natural to observe differences in developmental patterns between morphemes as well. By modeling both systematicity and individuality simultaneously, this paper aims to gain a more comprehensive view of morpheme accuracy development.

This paper presupposes no knowledge of generalized additive (mixed) models. It, assumes, however, that readers are familiar with the basic ideas of regression modeling including generalized linear models and model comparison based on information-theoretic measures such as AIC. It further assumes that readers have a basic knowledge of mixed-effects models. Online Supporting Document 1 provides an introduction to general ideas in regression modeling that are necessary for this paper.

Data Source and Analysis

Corpus

This paper employed EF-Cambridge Open Language Database (EFCAMDAT; Alexopoulou, Geertzen, Korhonen, & Meurers, 2015). The learner corpus includes writings at Englishtown, an online school run by Education First. A course in Englishtown consists of 16 levels with eight units each. Although learners are free to go back or skip units, they usually progress from lower to higher levels unit by unit. A placement test suggests an appropriate level at which learners begin their coursework. At the end of each unit is a free composition task on a variety of topics (e.g., self-introduction, making requests). A sample writing is provided for each writing task, and learners can consult the sample and other external resources such as dictionaries in the process of writing. Each writing task specifies length, with assignments ranging from 20-40 words in Level 1 Unit 1 to 150-180 words in Level 16 Unit 8. Teachers provide feedback on writings, including the correction of erroneous grammatical morphemes. The present study used teacher feedback as error tags and collected necessary information to calculate accuracy by exploiting them. Error tags are not annotated in all of the writings, however. Apart from learners' writings, EFCAMDAT includes, for each writing, such metadata as the ID of the learner, his/her country of residence, and the date and time of submission. This information allows researchers to track the longitudinal development of individual learners. EFCAMDAT is publicly available at <http://corpus.mml.cam.ac.uk/efcamdat/>.

Target morphemes. The initial target was six English grammatical morphemes: articles, past tense *-ed*, plural *-s*, possessive *'s*, progressive *-ing*, and third person *-s*. These are the morphemes that have often been targeted in SLA literature (cf. Goldschneider & DeKeyser, 2001). However, possessive *'s* was dropped as it did not occur frequently enough to allow for the investigation of individual variation or longitudinal development. Furthermore, progressive *-ing* and third person *-s* were dropped because their accuracy was close to 100% throughout learners'

development. High accuracy rates make the inspection of development difficult because we cannot distinguish learners who barely achieve 100% accuracy from those who do so effortlessly (i.e., the ceiling effect). Thus, the final set of target morphemes was composed of articles, past tense *-ed*, and plural *-s*. Articles included both definite and indefinite articles. Past tense *-ed* included only regular past tense forms (e.g., *opened*) and not irregular ones (e.g., *thought*). Similarly, plural *-s* included only regular forms (e.g., *cups*) and not irregular ones (e.g., *mice*).

Target L1 groups and proficiency levels. The present paper targeted the following 10 L1 groups with the largest amount of data in EFCAMDAT: Brazilian Portuguese, Mandarin Chinese, German, French, Italian, Japanese, Korean, Russian, Spanish, and Turkish. As EFCAMDAT does not provide direct information as to learners' L1s, such information was inferred from the countries in which learners reside as a close approximation. Accordingly, L1 Brazilian Portuguese, German, French, Italian, Korean, Russian, and Turkish learners correspond to those living in Brazil, Germany, France, Italy, Korea, Russia, and Turkey, respectively. L1 Mandarin Chinese learners included those living in Mainland China and in Taiwan, and L1 Spanish learners included those living in Spain and Mexico. L1 Mandarin Chinese is referred to as L1 Chinese and L1 Brazilian Portuguese as L1 Brazilian to save space. English town Levels are aligned with the Common European Framework of Reference (CEFR), as shown in Table 1.

Subcorpus. The present paper only targeted learners whose sum of obligatory contexts and overgeneralization errors in error-tagged texts was 10 or more for each of the three morphemes. In addition, due to the high computational cost of part of the analysis, it was necessary to limit the data to a maximum of 20 learners from each L1 group. The 20 learners selected were those with the largest number of writings within the L1 group. Because the L1 French, L1 Japanese, L1 Korean and L1 Turkish groups included 20 or fewer learners after

applying the first selection criterion (i.e., obligatory contexts plus overgeneralization errors \geq 10), the second criterion was not relevant to the four groups.

Figure 1 shows the distribution of learners and error-tagged writings across L1 groups and Englishtown levels. Learner level was operationalized as the learner's mean level in Englishtown. Note that the Total panel has a different y-axis scale from the other panels. In all, there were 3,323 writings by 158 learners. The subcorpus included 315,141 words in total, and the average number of words per writing was 94.8 (SD = 50.0).

Accuracy Measure and Data Extraction

As a measure of accuracy, this paper employed the ratio between correct uses and errors. The number of correct uses was obtained by subtracting the number of omission and misinformation errors from that of obligatory contexts. Obligatory contexts were operationalized as morpheme use in the corrected text, which is the text wherein erroneous portions were replaced with the corresponding corrected forms based on error tags. For instance, if a learner wrote, *She has a big nose and small mouth*, and it was corrected to read, *She has a big nose and a small mouth*, there are two obligatory contexts of articles as the article occurs twice in the corrected sentence. Errors combined omission, misinformation, and overgeneralization errors. This accuracy measure is conceptually equivalent to target-like use scores (TLU; Pica, 1983). In visualizing accuracy, the study used TLU scores, which are calculated by dividing the number of correct uses by the sum of the numbers of obligatory contexts and overgeneralization errors. R scripts were written to count the frequency of obligatory contexts and each type of errors in error-tagged texts.¹

There is no guarantee that errors are exhaustively annotated in EFCAMDAT. However, a manual given to Englishtown teachers asks them to be complete in providing feedback, and it

explicitly raises articles, plural -s, and verb tense among the features teachers should pay attention to. This briefing should make error annotation of the target morphemes fairly comprehensive and reliable. While the use of teacher feedback as error annotation can introduce noise to the data, the information provided by the annotation leads to intriguing insights into patterns of accuracy development, as will be discussed later.

Variables and Analysis

In this paper, I modeled accuracy as a function of several variables and explored the models to address the research questions. The dependent variable was accuracy in the form of odds. In the variants of logistic regression models employed in this study, the number of correct uses was entered as the number of successes, and the number of errors was entered as the number of failures.

There were four independent variables: proficiency, writing number (writingnum), morpheme, and L1 type (L1type):

- Proficiency was represented by the average Englishtown level in terms of unit at which the learner submitted his/her writings. The value is unchanged within learners, and the variable is meant to capture between-learner, cross-sectional development. Proficiency was standardized to facilitate interpretation. The mean and standard deviation of proficiency were 51.8 (Level 7 Unit 4) and 22.6, respectively.
- Writing number represented the within-learner writing order. One indicates the first writing of a learner, two indicates the second writing, and so forth. Writing numbers were assigned to both error-tagged and untagged writings so that we can interpolate development over untagged writings. This variable was meant to capture within-learner, longitudinal development, and was standardized over learners after its values were

centered within each learner. Accordingly, zero in the standardized writing number indicates the mean writing number within each learner. The standard deviation of the writing number was 15.6.

- Morpheme was a categorical variable with three levels: one for each morpheme with articles as the reference level.
- L1 type was a dichotomous variable representing L1 influence and indicating whether an L1 has an equivalent morpheme. The L1 type had two levels: ABSENT and PRESENT. The ABSENT group was the reference level. The ABSENT group include L1 groups that lack the equivalent linguistic features in their L1s, and L1 groups wherein the marking of the feature is optional. By contrast, the PRESENT group must mark equivalent features. For instance, L1 Japanese was considered to be in the ABSENT group in the article because it is not obligatory in Japanese to express definiteness, the central concept of the English article system. Conversely, Japanese was considered to belong to the PRESENT group in past tense *-ed* because a Japanese morpheme, *-ta*, roughly corresponds to past tense *-ed* in English, and it is difficult to express past-ness without the use of this morpheme in Japanese. This approach to representing the effect of L1 is rather crude and oversimplified, but as will be shown, it is useful to capture L1 influence. The ABSENT group included L1 Chinese, Japanese, Korean, Russian, and Turkish for articles; L1 Chinese for past tense *-ed*; and L1 Chinese, Japanese, and Korean for plural *-s*. The rest were included in the PRESENT group.

In addition to these variables, some of their interactions were entered into the model as well in a stepwise manner. Treatment contrasts were used for categorical variables throughout this paper.

Before running the analyses, observations without any obligatory contexts or overgeneralization errors were removed. There were 7,247 non-zero observations across the three morphemes. Table 2 shows the average number and standard deviation of non-zero observations, obligatory contexts, omission errors, and overgeneralization errors per learner. Naturally, the data size is larger for articles and for plural -s than for past tense -ed due to their higher frequency.

All of the statistical analyses in this paper were performed with R (version 3.2.1; R Core Team, 2015; cf. Mizumoto & Plonsky, in press). The R codes and data used in this paper are available at the following repository at the Open Science Framework: <https://osf.io/dbuh4>. Before moving on to the main analysis, a cross-sectional view of the data is presented herein.

Cross-Sectional View of Morpheme Development

Figure 2 illustrates the cross-sectional development of the three morphemes across L1 types. Each line shows the cross-sectional development in each L1 type. Unlike typical cross-sectional data, however, a learner contributes multiple data points to the figure as he/she produces multiple writings. C2 level was dropped out of the figure due to its small data size but is included in modeling that follows.

The fluctuation of accuracy in the graph, which is partially due to the small data size of several observations, makes the close examination of the data difficult. Employing variants of logistic regression models that weigh each observation according to its data size, the present paper investigates whether we can observe a significant difference in the developmental pattern across groups and the extent to which individual variation is present in the development.

Taking into Account Individual Variation: Generalized Linear Mixed-Effects Models

This section analyzes the extent to which accuracy developmental patterns vary across individual learners. To quantify individual variation, the paper employs a generalized linear mixed-effects model (GLMM). Mixed-effects models can handle both systematicity and individuality because they can deal not only with usual within- and between-learner fixed-effects variables such as morphemes and proficiency (i.e., systematicity) but also remaining variance across and within learners (i.e., individuality). Partly for this reason, mixed-effects models have been widely used in longitudinal data analysis (Long, 2012), including in SLA (Barkaoui, 2014; Kozaki & Ross, 2011).

Model Specification and Model Selection

The study employed a mixed-effects logistic regression model to analyze the relationship between accuracy, proficiency, longitudinal development, and morpheme. The model included L1 and learner as random-effects factors. Writings were nested within individual learners, who were in turn nested within L1 groups. The model thus had nested random-effects structure where variance was partitioned into between-L1, between-learner, and between-writing levels (cf. Gries, 2015). Although it was possible to construct yet another level by viewing data points as nested within writings, this was not attempted to avoid further complexity of the model. By-L1 random intercepts allow overall accuracy to vary across L1 groups. Variables can also be entered as random contrasts and random slopes. By-L1 random contrasts and random slopes, however, were not entered because the small number of L1 levels (10) may result in unstable models.

The role of each random-effects parameter is as follows. When the by-morpheme random contrasts are present, the by-learner random intercepts allow article accuracy to vary across individual learners. By-morpheme random contrasts represent individual variation in the accuracy difference between morphemes. The by-writingnum random slope similarly represents

individual variation in morpheme-independent learning rates implying whether some learners are naturally quicker in learning than others. We are interested in the extent to which we can observe such individual differences and whether — and to what extent — systematic variables (e.g., proficiency) can account for these differences.

The study constructed multiple models and found the most plausible model by comparing them. There has been no agreement on how best to perform model selection in mixed-effects modeling (Gries, 2013). It has been suggested that we should start with the maximal model, or the model with all possible predictors and the largest random-effects structure, and drop the variables and/or the random-effects component that are not supported by the data (Barr, Levy, Scheepers, & Tily, 2013; Gelman & Hill, 2007; see Bates, Kliegl, Vasishth, & Baayen, submitted for a counter-argument). However, because the present model only had 10 L1 groups, I opted for an approach where initially the simplest model was built and predictors were added to the model one at a time only if it improves the model. More specifically, the following forward selection approach was used (cf. James, Witten, Hastie, & Tibshirani, 2013). I first built the so-called unconditional model (Bates, 2010) that only includes by-L1 and by-learner random intercepts but no fixed-effects predictor. I then added a predictor one by one that decreases the model's AIC the most, repeating the procedure until no predictor could improve the model further (but see Whittingham, Stephens, Bradbury, & Freckleton, 2006 for criticisms of the stepwise approach in general). Interaction terms were considered only when the model already included the main effects constituting the interactions. Random contrasts were considered only when the variable was already in the fixed-effects component of the model. Although restricted maximum likelihood (REML) procedures are often used for linear mixed-effects models, the present study employed maximum likelihood estimation because REML does not allow for the

comparison of models with different fixed-effects structure (Bolker et al., 2009) and also because REML estimates are not well-defined for GLMMs (Bates, 2009). All of the statistical analyses in this section were carried out with the `lme4` package (version 1.1-8; Bates, Mächler, Bolker, & Walker, submitted) in R. To avoid convergence failure, the BOBYQA algorithm was used as the optimizer, as suggested by Bolker (2014).

Let us call the unconditional model Model 1. Model 2 added morpheme to the fixed-effects part because the predictor most dramatically decreased the AIC. A comparison between Model 1 and Model 2 tested whether different morphemes are of different accuracy levels. Model 3 added by-morpheme random contrasts to Model 2. A comparison between Model 2 and Model 3 tells us whether it is worth allowing the accuracy difference between morphemes to vary across learners. Likewise, Model 4 and Model 5 tested the effects of `L1type` and `writingnum`, respectively. Model 6 further added the by-`writingnum` random slope. Models 7 through 9 examined the effects of proficiency, morpheme-proficiency interaction, and `L1type`-`writingnum` interaction, respectively.

Table 3 shows the summary of model comparison. The first three columns give the model number and the variables included in the fixed effects and random effects of the model. The fourth column lists the AICs of the model, and the fifth column shows the difference of AIC in comparison to the previous model. A negative value means that this model has better predictive accuracy than the model above. The last two columns show the results of likelihood ratio tests comparing the model with the previous model. The table indicates that AIC categorically decreased until Model 9, and likelihood ratio tests similarly suggested steady improvement until the same model. No other term (e.g., morpheme-`writingnum` interaction) further decreased AIC.

Although Model 9 may appear to be the most plausible model, Model 8 was selected as the final model because the decrease of AIC from Model 8 to Model 9 (-2.8) is fairly small and the p -value of the added parameter (0.029 for the L1type-writingnum interaction) is not as low as the p -values of other parameters, either. Given that both models are already highly complex and that AIC tends to prefer more complex models in general (Held & Bové, 2014), I opt for the less complex model herein. As a reference, I also constructed a model that has the same structure as Model 8 but does not include L1type, proficiency, or any interaction terms involving them. A comparison between Model 8 and this reference model informs us of the extent to which L1 type and proficiency explain the variance.

The forward-selection procedure employed above may result in underspecified models because there can be a model that is better than the final model and includes a combination of parameters untested in the model selection procedure. To mitigate the potential effect of the procedure, a series of models were built in the following manner². Instead of adding one variable at a time, I added two variables that sequentially decreased AIC the most. I then deleted one variable that resulted in the minimum increase in AIC. This procedure was repeated until no further iteration decreased AIC further. This process partially alleviates the potential underspecification issue because the procedure explores part of the parameter combination space that is not tested in the pure forward-selection procedure. This 2-in-1-out procedure resulted in Model 9 as the final model, thereby partially confirming the robustness of our model. For the reason described above, however, we take Model 8 as the final model.

Interpretation of the Model

Interpreting random effects. Table 4 presents the random-effects components of the mixed-effects model. It also shows the random effects of the Reference Model, against which the

effects of predictors in Model 8 are tested. In Table 4, the intercept rows represent the standard deviation of random intercepts for L1 and learner, and the other rows show the standard deviation of by-morpheme random contrasts and by-writingnum random slopes.

The Reference Model tells us that the standard deviation of the by-L1 random intercept (Row 2 in the table) is 0.300, which indicates the dispersion of L1 groups in absolute accuracy in the logit scale. Similarly, the by-learner random intercept (Row 4) is 0.495, which is the magnitude of individual differences in article accuracy within each L1 group after progressing the mean number of writings (i.e., standardized writingnum = 0). The standard deviations of the by-morpheme random contrasts (Row 5-7) are 0.716 for past tense *-ed* and 0.582 for plural *-s*, and denote individual differences in the accuracy difference between articles and the morphemes. The standard deviation of the by-writingnum random slope (Row 8) is 0.192, which represents the magnitude of individual variation in the overall learning rate. When the values in Model 8 are examined, we notice a fair amount of decrease in the by-learner random intercept ($0.495 \rightarrow 0.412$, or -20.0%). This shows the extent to which learners' overall proficiency and L1type explain individual variation in article accuracy. The by-morpheme random contrast similarly decreases in Model 8 ($0.716 \rightarrow 0.613$, or -16.8%, for past tense *-ed* and $0.582 \rightarrow 0.481$, or -21.1%, for plural *-s*). This represents the degree to which proficiency (but not L1type due to the absence of L1type-morpheme interaction in the fixed-effects structure) explains individual variation in between-morpheme accuracy difference.

Surprisingly, the by-writingnum random slope increases from the Reference Model to Model 8 ($0.192 \rightarrow 0.197$, or +2.3%). This is rooted in the fact that some of the within-learner variance can be reflected as between-learner variance in mixed-effects modeling (Hox, 2002; Snijders & Bosker, 1994). As a result, a within-learner predictor may explain both within-learner

and between-learner variance. Maximum likelihood estimation (MLE) generally corrects this. However, when a predictor is centered or standardized within learners as in the present case, it results in smaller between-learner variation in the average predictor value than embedded in the correcting mechanism of MLE. This invites overcorrection by MLE, and random effects may increase as a result. This, therefore, does not mean model misspecification.

Because the value is larger in past tense *-ed* random contrast than in plural *-s* random contrast, a larger individual variation remains in the accuracy difference between articles and past tense *-ed* than in the accuracy difference between articles and plural *-s*. Note the caveat, however, that because between-learner and within-learner variability are not completely independently quantified even in mixed effects models, random-effects components of different models are not strictly comparable. Comparison, however, is a common practice (e.g., Hox, 2002) and is still a useful strategy by which to examine the effect of predictors on random-effects components.

Interpreting fixed effects. Let us now turn to the fixed-effects part (Table 5). P-values indicated by asterisks are only approximate, but a parametric bootstrap — a resampling technique that compares the target model with the reduced model that does not include the interested parameters (Pinheiro & Bates, 2000) — agreed with the significance of parameters in the table with the significance level of $p < 0.05$ based on 1,000 samples. We can thus make the following observations.

- The main effect of morpheme (Row 2-4 in the table) is significant. At the mean proficiency level, the accuracy of plural *-s* is generally higher than that of articles.
- The main effect of the L1 type (Row 5-6) is also significant. The PRESENT group overall outperforms the ABSENT group at standardized writingnum = 0.

- The main effect of writing number (Row 7) is significant. As learners write, morpheme accuracy increases.
- The main effect of proficiency (Row 8) is significant and positive. Article accuracy tends to be higher in learners of higher proficiency.
- The morpheme-proficiency interaction (Row 9-11) shows that the accuracy increase over proficiency is smaller in plural *-s* than in articles. In plural *-s*, accuracy increase per standard deviation of proficiency is nearly negligible ($0.238 - 0.224 = 0.015$)³.

Let us further examine some of the terms that did not turn out to be significant.

- The L1type-proficiency and L1type-writingnum interactions were not present in the final model. This means that there is no evidence showing different cross-sectional or longitudinal developmental patterns between the PRESENT and ABSENT groups. This is interesting because the PRESENT group generally outperforms the ABSENT group and they could be more likely to have approached the ceiling.
- The morpheme-writingnum interaction was not retained in the final model. This outcome shows that the rate of longitudinal development is similar across morphemes.

To look into the magnitude of individual variation, it is interesting to compare random effects in Table 4 with the corresponding fixed effects in Table 5. The fact that the random contrast for past tense *-ed* is 0.613 and its estimate in the fixed effects structure is 0.141 means that at the mean proficiency level, the standard deviation of individual variation in the accuracy difference between articles and past tense *-ed* is much larger than the mean accuracy difference between the two morphemes, which in turn indicates that although past tense *-ed* is more accurate than articles on average in this sample, the accuracy order between the two morphemes depends heavily on learners. The case of the by-writingnum random slope is similar. The

standard deviation in Table 4 is 0.197, while the coefficient in the fixed effects 0.082. This indicates that while on average learners' longitudinal development is characterized by increased accuracy, a great proportion exhibits a decreased accuracy overall. This is not the case for the difference between articles and plural *-s*, however. Because its random slope (0.481) is smaller than the fixed-effects coefficient (0.787), plural *-s* is usually (though not necessarily always) more accurate than articles in individual learners. At higher proficiency levels, however, the mean difference between the two morphemes decreases, as reflected on the negative coefficient of the interaction between proficiency and plural *-s*. The proportion of the learners whose accuracy is higher in articles than in plural *-s* is expected to increase. The discussion here illustrates that it is possible to quantify individual variation through GLMMs.

Summary of the GLMM Approach

The present section demonstrated systematicity (e.g., plural *-s* is on average more accurate than articles) and individual variation in the L2 accuracy of grammatical morphemes. In addition to its ability to model systematicity and individuality simultaneously, a particular strength of the GLMM is its feature of quantifying individual variation through random effects. Variance in random effects is informative as to (i) the extent to which individual variation is present in a certain effect (e.g., the standard deviation of the individual variation in article accuracy is 0.412 in logit scale), (ii) whether it is larger or smaller compared to individual variation in another effect (e.g., individual variation in accuracy difference between articles and past tense *-ed* is larger than the variation in the difference between articles and plural *-s*), and (iii) the degree to which predictors explain variation (e.g., proficiency decreases the accuracy difference between articles and plural *-s* by 21.1%).⁴

Accounting for Nonlinearity and Individuality: Generalized Additive Mixed Models

In the previous section, the analysis assumed a linear change of accuracy in both cross-sectional and longitudinal development. The assumption, however, is unwarranted, particularly in light of prior SLA research demonstrating nonlinear learning curves (DeKeyser, 1997; Lightbown, 1983). The present section examines whether the developmental path varies depending on learners' L1 types and on morphemes when nonlinear development is assumed.

Brief Overview of Generalized Additive Models

Generalized additive models (GAMs) extend generalized linear models (GLMs) and model nonlinear relationships between independent and dependent variables. They achieve nonlinearity through the use of *splines*. The following explanation of splines is largely based on James et al. (2013) and Hastie, Tibshirani, and Friedman (2009).

A traditional way of modeling nonlinearity is by using polynomial functions. However, they cannot model flexible shapes without spending a large number of degrees of freedom, and doing so renders the resulting model unstable. In regression splines, one polynomial function models only part of the data, and multiple functions are used to cover the entire data. Those functions are smoothly connected so that there is no wide jump in the predicted value. This point is illustrated in the upper two panels in Figure 3. Figure 3A demonstrates morpheme development in hypothetical learners. The dashed line represents the predicted values of accuracy based on a cubic function of proficiency (i.e., $TLU = \beta_0 + \beta_1 \times proficiency + \beta_2 \times proficiency^2 + \beta_3 \times proficiency^3$, where β 's are estimated from the data)⁵. Here, we observe relatively large differences between observed (i.e., small circles) and fitted (i.e., dashed line) values. A cubic function is thus inadequate for modeling accuracy development in this dataset. The solid line is a piecewise cubic function. Data points were horizontally divided into five equally spaced regions, and a cubic function was fitted to each region. We see that the

predicted function is absurd as a whole: The lines are not connected and there are jumps in the fitted value as a result. Thus, simply employing multiple piecewise polynomial functions is insufficient for modeling nonlinearity.

To achieve more natural modeling of nonlinearity, certain constraints can be imposed on the piecewise polynomial functions. Specifically, it is common to constrain piecewise cubic functions so that the values of the function and its first and second derivatives are continuous at *knots*, the points at which cubic pieces connect. This way, the function is not only continuous throughout but also smooth at the knots (Zuur, Ieno, Walker, Saveliev, & Smith, 2009). In Figure 3B, the same data points are modeled by a *smoothing spline*. Though conceptually somewhat different, it is mathematically a variant of the cubic splines discussed above. Based largely on cubic functions, the smoothing spline models the data well in the present case.

The spline balances difference between fitted and observed values and the roughness or wiggleness of the curve. If it is allowed to be infinitely wiggly, it goes through all of the observed data points and would clearly overfit the data by modeling noise in addition to the underlying shape, thereby making it difficult to generalize to new datasets. If, on the other hand, the spline is not allowed to be wiggly at all, it would end up being a straight line that models nonlinearity poorly. The smoothing spline achieves this bias-variance trade-off through a procedure called generalized cross validation, which is an approximation of leave-one-out cross validation commonly employed to evaluate statistical models. Conceptually, it fits to all but one data point a spline function with a certain degree of smoothness, and calculates the difference between the observed value of the omitted data point and its predicted value based on the spline (i.e., error). This process is repeated as many times as there are data points. The average error in this procedure indexes the goodness of the degree of smoothness. This whole process is then repeated

for a wide range of smoothness values, and the optimal wiggleness is found in which the average error is minimized (Wood, 2009; Zuur et al., 2009).

GAMs are a semi-parametric technique that combines the smooths discussed above with parametric terms, thereby allowing a statistical test of the significance of some terms while controlling for the nonlinear effects of other terms. The lower two panels of Figure 3 illustrate the importance of accounting for nonlinearity through GAMs. These two figures show hypothetical accuracy development in two L1 groups: L1 Japanese and L1 Spanish. Both groups show clear U-shaped developmental patterns, but the data for L1 Spanish learners were generated to mark higher accuracy overall than those for L1 Japanese learners throughout development. In Figure 3C, the pattern is modeled by a linear function. It forces linearity on the nonlinear shape, resulting in large differences between observed and predicted values. Because the model hardly explains variance and residuals are large, the accuracy difference between the two L1 groups is non-significant ($t = 1.643$, $p = 0.102$) when proficiency is (mis)controlled for, despite the consistently higher accuracy of the L1 Spanish learners. Figure 3D models the same data with a GAM based on a *thin plate regression spline* (Wood, 2003), an approximation to a thin plate spline (Wood, 2010), which is a generalized form of the cubic spline discussed earlier. The model was constructed with the `mgcv` package (version 1.8-6; Wood, 2006) in R. Here, without pre-specifying shape, the GAM accurately models the U-shape. This in turn results in much smaller residuals than in Figure 3C, and this time, the effect of L1 is correctly identified ($t = 7.858$, $p < 0.001$). Therefore, the GAM was able to model the usual parametric term (L1) and nonparametric smooth (nonlinear effect of proficiency) simultaneously.

An exciting recent development is the incorporation of random effects into GAMs, making the model capable of accounting for nonlinear patterns of individual learners. The model,

referred to as a generalized additive mixed model (GAMM; Baayen, 2014, in preparation, Chapter 8; Wood, 2004, 2006), can construct separate wiggly curves for each learner by *penalized factor smooths*, which achieve the interaction between smooths and factors with the same degree of smoothness across learners (Wood, 2014). GAMMs have been used in psycholinguistics (e.g., Balling & Baayen, 2012; Mulder, Dijkstra, Schreuder, & Baayen, 2014), sociolinguistics (Wieling, Montemagni, Nerbonne, & Baayen., 2014), and SLA (Ning, Shih, & Loucks, 2014).

Model Specification and Model Selection

Models assumed binomial error distribution and employed a logit link function. The dependent variable and the potential independent variables were the same as the GLMM's, except that nonlinear terms were also considered. The interaction between two nonlinear terms (i.e., proficiency-writingnum) was entered as a *tensor product smooths*. Tensor product smooths extend nonlinearity to more than one dimension and model wiggly surfaces between the variables of naturally different scales (Hastie et al., 2009; Wood, 2010).

A separate smooth was constructed for each factor level when L1 type or morpheme interacted with proficiency and/or writingnum. For example, in the specification of L1type-proficiency interaction, separate proficiency curves were created for each L1 type. Thus, unlike interactions in typical regression models, factor-smooth interactions in GAMs also account for the main effects of the continuous variables included in the interaction. Due to a centering constraint, factors need to be specified in the model separately.

Due to the high computational cost of GAMMs, building a model takes a relatively long time, and it was impractical to run the forward selection process in model selection that requires building multiple models at each step. Instead, I started with a model that was conceptually

equivalent to the final GLMM constructed earlier, and tested whether all of the parameters included in the model were necessary and whether including additional terms improved the model. Model 1, thus, included

- L1type and morpheme as fixed effects,
- as random effects by-L1 random intercepts, by-learner random intercepts, and by-morpheme random contrasts at the level of individual learners,
- (standardized) writingnum and by-morpheme (standardized) proficiency as smooth terms to capture their potentially nonlinear effects, and
- by-writingnum random wiggly curves at the learner level.

Smooth terms were specified with thin plate regression splines. Random wiggly curves are similar to random slopes but also allow nonlinearity in the longitudinal developmental patterns of individual learners. Maximum likelihood estimation was employed. This model is different from the final GLMM in that nonlinear effects are assumed in proficiency and writingnum, and random wiggly curves are assumed instead of random slopes for individual learners.

With this model as the starting point, I first tested whether any additional terms improve the model. For this purpose, five candidate models were built:

1. Model 1 + L1type-morpheme interaction in the fixed-effects structure ($B = -0.004$, $p = 0.985$ for PRESENT - past tense *-ed*; $B = -0.125$, $p = 0.355$ for PRESENT - plural *-s*)
2. Model 1 - writingnum smooth + writingnum smooth for each morpheme (i.e., writingnum-morpheme interaction; $\chi^2 = 0.115$, $p = 0.735$ for the writingnum curve for articles; $\chi^2 < 0.001$, $p = 0.999$ for the writingnum curve for past tense *-ed*; $\chi^2 = 0.011$, $p = 0.915$ for the writingnum curve for plural *-s*)

3. Model 1 + proficiency smooth for each L1 type ($\chi^2 < 0.001$, $p = 0.999$ for the proficiency curve for the ABSENT group; $\chi^2 = 0.029$, $p = 0.864$ for the proficiency curve for the PRESENT group)
4. Model 1 + proficiency-writingnum interaction realized as a tensor-product interaction ($\chi^2 = 5.327$, $p = 0.419$)
5. Model 1 + - writingnum smooth + writingnum smooth for each L1type (i.e., writingnum-L1type interaction; $\chi^2 = 0.027$, $p = 0.871$ for the writingnum curve for the ABSENT group; $\chi^2 = 19.416$, $p < 0.001$ for the writingnum curve for the PRESENT group)

The model selection procedure, based on p values in the parentheses (Wood, 2013a, 2013b), suggests that Candidate Model 5 is better than Model 1, and AIC-based model comparison supports the decision as well ($\Delta\text{AIC} = -8.2$). This model is referred to as Model 2.

The next step is to test whether it is worth adding further terms. The same procedure was again followed, except that candidate terms were added to Model 2 this time. The added terms were the same as Candidate Models 1 through 4 above. The process indicated that none of the terms improves the model ($p > 0.105$ for all of the terms).

I then examined whether we need all of the terms in Model 2. The p values of Model 2 parameters indicated that while some parameters were non-significant (e.g., $\chi^2 = 2.668$, $p = 0.102$ for the proficiency curve for past tense *-ed*), they were restricted to the levels of the factors or the levels of the interaction terms involving the factors whose other levels were significant (e.g., $\chi^2 = 21.869$, $p < 0.001$ for the proficiency curve for articles). This indicates that all of the terms should be kept in the model.

Model 2, however, suggested that the effect of proficiency is linear (EDF = 1.000 for all of the morphemes). The proficiency term, therefore, was moved to the parametric part: Model 3

included L1type, proficiency, morpheme, and the proficiency-morpheme interaction as fixed-effects parametric terms. This did not affect AIC ($\Delta\text{AIC} = -0.002$). To further test whether random wiggly curves are necessary, another model was constructed in which random wiggly curves in Model 3 were replaced with by-writingnum random slopes. In other words, the model assumes linear effects of writingnum at the level of individual learners. Model comparison indicated that we need to keep random wiggly curves ($\Delta\text{AIC} = 112.5$), suggesting that the learning curve is nonlinear at the level of individual learners.

The above did not directly tell us whether we need separate writingnum curves for the two L1 types. To analyze this, a separate curve was estimated on top of the curve for the reference level. In other words, to examine whether L1 type affects the longitudinal developmental pattern, two separate curves were constructed: one for the ABSENT learners and the other for the PRESENT learners on top of the curve for the ABSENT group (Baayen, in preparation, Chapter 8; Wieling, 2015; Wood, 2014). If the latter is significant, it suggests that it is worth having an additional curve for the PRESENT group on top of the ABSENT group curve, which in turn means that the longitudinal developmental pattern differs across L1 types. The results suggested that we need a separate writingnum curve for the PRESENT group ($\chi^2 = 13.472, p = 0.012$). I, thus, select Model 3 as our final model and will explore it below.

Interpretation of the Model

Tables 6 through 8 show the results of the final model. Parametric terms (Table 6) suggest that (i) PRESENT learners generally outperform ABSENT learners (Row 3), (ii) higher proficiency learners tend to be more accurate in using articles than lower proficiency learners (Row 4), (iii) learners are more accurate in the use of plural -s than articles at the mean proficiency level (Row 7), and (iv) cross-sectional accuracy increase is smaller in plural -s than

in articles (Row 10). The non-significance of L1type-proficiency interaction indicates that the cross-sectional developmental pattern can be assumed to be similar across L1 types. Table 7 shows estimated degrees of freedom (EDF), reference degrees of freedom (Ref.df), χ^2 , and p -values for the splines. If the EDF is close to one as in the effect of writingnum in the ABSENT group, the relationship between independent and dependent variables is close to linear in logit scale (Baayen, 2010a), and the larger its value, the wigglier the curve is. The table shows linearity in the partial effect of writingnum for the ABSENT group (EDF = 1.001 at Row 2) but nonlinearity for the PRESENT group (EDF = 3.503 at Row 3). The table also indicates significant individual variation in longitudinal development (Row 4). Table 8 indicates the standard deviation of random effects.⁶ As in the GLMM, we can observe between-L1 variation in absolute accuracy and individual variation in the accuracy difference between articles and other morphemes. Drawing inferences from the above tables, however, is not necessarily straightforward: Smooth terms in Table 7 make interpretation especially difficult. I turn now to one strategy that can assist us in drawing inferences from the results: visualizing the fitted values.

Figure 4 shows the fitted nonlinear accuracy development in individual learners. The upper panel represents adjustments to logit TLU scores for individual learners across standardized writing numbers. If there is no individual variation within each L1 type, morpheme, and proficiency level, all of the lines should completely overlap. As we can see, however, large individual variation is present both in terms of absolute accuracy and developmental shape. The figure demonstrates large individual variation well, but it does not show how learners develop in the scale of TLU scores in a particular morpheme. The bottom four panels in Figure 4, therefore, show the fitted values of article accuracy in individual learners divided into two proficiency groups (higher vs lower) and two L1 types (ABSENT vs PRESENT). The cut-off proficiency

level for the two proficiency groups was learners' mean proficiency. The thick lines in each panel are locally weighted scatterplot smoothing lines (LOESS; Larson-Hall & Herrington, 2010; Singer & Willett, 2003) showing the overall trend. Although the parametric terms in Table 6 indicate that the PRESENT group outperforms the ABSENT group, this is hardly visible in Figure 4 due to individual variation within each L1 type. Furthermore, whereas higher proficiency learners use articles more accurately than lower proficiency learners on average, this is merely a tendency and only characterizes the development of the hypothetical 'average' learner. Individual variation definitely outweighs the typological difference in L1 and can also have a larger impact than general proficiency. Moreover, the developmental pattern slightly differs between the ABSENT and PRESENT groups, as Table 7 indicates. However, the figure also suggests that this difference is marginal compared to the scale of individual variation.

Summary of the GAMM Approach

The GAMM took into account individual variation and nonlinearity and modeled accuracy development as a function of proficiency, longitudinal development, and L1 type. The final model demonstrated (i) individual variation in absolute accuracy and in nonlinear development, (ii) systematic L1 influence and proficiency effects on absolute accuracy, and (iii) L1 influence on longitudinal developmental patterns. The empirical and quantified demonstration of nonlinearity, individual variation, and systematicity was only achievable through GAMMs.

Contrasting GLMM/GAMM with GLM/GAM

Now that both types of models have been explored, they are compared against each other and against GLMs and GAMs, the models that do not account for individual variation. The only difference between the GLMM and the GAMM is that the GAMM includes the L1type-

writingnum interaction while the GLMM does not. Recall that the term was at the borderline in the model selection process of GLMMs. As mentioned earlier, Figure 4 based on the final GAMM also shows that the difference in the developmental curve between the two L1 types is minute, especially in view of large individual variation. Thus, although it is worth including the interaction term in the model when nonlinearity is accounted for, I conclude that its effect is nearly negligible from a practical perspective.

It is also interesting to compare GLMM/GAMM with GLM/GAM as such a comparison highlights the importance of taking individual variation into account when modeling L2 development. As in the GLMM, GLMs and GAMs were constructed based on the forward selection approach. Both GLMs and GAMs used the logit link function and assumed binomial error distribution. The final GLM included morpheme, L1 type, proficiency, writingnum, morpheme-proficiency interaction, and proficiency-L1type interaction. The final GAM included morpheme and L1 type as parametric terms, and as smooths terms separate wiggly proficiency curves for each morpheme, separate writingnum curves across L1 types, and a proficiency-writingnum wiggly surface.⁷

The results showed a few disputed findings between GLMM/GAMM and GLM/GAM. More specifically, the GLM supported the L1type-proficiency interaction and the GAM included the proficiency-writingnum interaction, while GLMM and GAMM supported neither. In addition, the GAM suggested nonlinear cross-sectional development, while the GAMM demonstrated linear development. Notice that the findings of the GLMM/GAMM were more conservative than those of the GLM/GAM: The GLM/GAM either pointed toward more significant parameters than the GLMM/GAMM or suggested nonlinear effects when the GAMM indicated linear effects. These are all likely to be rooted in whether individual variation is taken into account (GLMM

and GAMM) or not (GLM and GAM). Generally speaking, ignoring the nested structure of data results in unfairly small standard errors (Hox, 2002; Long, 2012), leading to narrower confidence intervals (cf. McKeown & Sneddon, 2014; Wieling, 2015). In the present context, because the GLM and the GAM ignore the dependency of data within individual learners, their standard errors turned out to be unfairly small, inviting spurious significant results.

The difference between the models is illustrated in Figure 5. The figure visualizes the predicted cross-sectional and longitudinal development of article accuracy in two learners: one L1 Russian and one L1 Brazilian who contributed the largest number of error-tagged writings among the ABSENT and PRESENT learners, respectively. The point of the figure is the magnitude of uncertainty represented by the width of shaded 95% confidence intervals, which are clearly wider in the GLMM and the GAMM panels than in the GLM and the GAM panels. The wider confidence intervals of the GLMM and GAMM are brought about by their ability to account for individual variation. The GLM suggested that cross-sectional developmental patterns vary across L1 types because models are (erroneously) certain of the trajectory of each L1 type and the trajectories differ, while the GLMM and the GAMM are much less certain that the two trajectories are different. Similarly, the GAM judged cross-sectional developmental patterns to be nonlinear because the narrow confidence intervals and a relatively fixed trajectory as their results suggest nonlinearity, while the wide confidence intervals of the GAMM and the resulting uncertainty in the trajectory do not support it. Thus, the GLMM/GAMM results are more trustworthy, and the illustration here demonstrates the significance of accounting for individual variation in modeling L2 development.

Discussion

GLMMs and GAMMs in this paper demonstrated nonlinearity and individual variation in the L2 development of English grammatical morphemes. SLA researchers have shown interest in these phenomena but were previously unequipped with the analytical tools to investigate them. With sophisticated statistical models of the type employed in this paper, however, we can model complex phenomena such as L2 development while losing much less information than traditional statistical techniques.

More specifically, the present paper showed that (i) plural *-s* is more accurate than articles in general, (ii) learners with an equivalent feature in their L1 outperform those whose L1s lack the feature, (iii) article accuracy increases as learners' proficiency rises, (iv) cross-sectional developmental patterns vary across morphemes, and (v) large individual variation is present in absolute accuracy, the accuracy difference between morphemes, and longitudinal developmental patterns. There was no disagreement in the above findings between GLMM and GAMM. Thus, we can safely conclude them.

The cross-sectional developmental pattern varies between articles and past tense *-ed* on the one hand and plural *-s* on the other. Articles and past tense *-ed* undergo more rapid increase in accuracy than do plural *-s*, whose accuracy remains relatively unchanged throughout development. This difference is likely due to the higher accuracy of plural *-s* and, as a result, the ceiling effect. It is interesting, however, that we did not observe a significant difference in the developmental pattern between articles and past tense *-ed* despite the fact that the article is a nominal free morpheme that encodes a language-dependent concept (i.e., definiteness) and past tense *-ed* is a verbal bound morpheme that encodes tense, a fairly language-independent concept. This finding shows that the classic distinctions between morphemes may not strongly influence

the developmental trajectory of morpheme accuracy. Because the paper targeted only three morphemes, this observation is merely suggestive rather than conclusive.

This paper demonstrated systematicity, individuality, and nonlinearity in L2 development. L1 type consistently exerted influence on accuracy in the present paper, demonstrating that accuracy is not determined randomly. However, as has been repeatedly emphasized throughout the paper, large individual variation was present as well both in the absolute accuracy and developmental patterns of morphemes. Together with the complex nonlinear patterns discussed earlier, I echo Baayen (2014, p.361):

The results obtained with GAMs can be embarrassingly rich, in the sense that the results are far more complex than expected given current models. GAMs will often challenge the state of the art of current theories, and the author's intuition is that they may force the field to move more into the direction of dynamic systems approaches to language.

Although the claim was made in the context of GAMs and GAMMs, it fully applies also to GLMMs.

I now briefly summarize features of the models discussed in this paper and note their potential weaknesses. The defining property of GLMMs is that they incorporate both fixed-effects and random-effects variables. This allows us to model systematicity and individuality simultaneously. GLMMs, however, are not very flexible in modeling nonlinearity. GAMMs can model nonlinearity and individual variation simultaneously. They cannot, however, currently handle a correlation parameter in random effects (Wieling, 2015). A further potential drawback is that they are less interpretable than other simpler models like GLMs/GLMMs (cf. James et al., 2013). It is worth noting, however, that interpretability of simpler models may come at the cost of less precision.

The present paper is not without its limitations. Accuracy in this paper was calculated by aggregating all error types. However, different mechanisms may operate between omission, misinformation, and overgeneralization errors or between definite and indefinite article uses. Ideally, error type should be incorporated into the model. Additionally, models in this paper only included developmental measures (i.e., proficiency and writing number) and L1-related variables (i.e., L1 and L1 type) as predictors of accuracy. Many more variables are certain to affect accuracy, such as tasks and linguistic contexts. Further investigation into the sources of variability should shed light on why cross-sectional and longitudinal developmental patterns take the form they do. The dataset itself is also a source of limitations. For example, it is worth looking into the potential effects of tasks, teaching materials in Englishtown, and varying progress rates across learners on the results (Alexopoulou et al., 2015).

Conclusion

The present paper introduced statistical models that capture systematicity, individuality, and nonlinearity and illustrated their potential in SLA research with the L2 accuracy development of English grammatical morphemes as an example. In light of the nonlinear and variable nature of L2 development, these techniques help researchers to better model L2 development and provide insights into the complex, dynamic, and nonlinear process of development.

References

- Alexopoulou, T., Geertzen, J., Korhonen, A., & Meurers, D. (2015). Exploring big educational learner corpora for SLA research: Perspectives on relative clauses. *International Journal of Learner Corpus Research*, 1(1), 96–129. doi: 10.1075/ijlcr.1.1.04ale
- Atkinson, D. (2011). A sociocognitive approach to second language acquisition: How mind, body, and world work together in learning additional languages. In D. Atkinson (Ed.), *Alternative approaches to second language acquisition* (pp. 143–166). Abingdon: Routledge.
- Baayen, R. H. (2010a). Demythologizing the word frequency effect: A discriminative learning perspective. *The Mental Lexicon*, 5(3), 436–461.
- Baayen, R. H. (2010b). A real experiment is a factorial experiment? *The Mental Lexicon*, 5(1), 149–157. doi: 10.1075/ml.5.1.06baa
- Baayen, R. H. (2014). Multivariate statistics. In R. Podesva & D. Sharma (Eds.), *Research methods in linguistics* (pp. 337–372). Cambridge: Cambridge University Press.
- Baayen, R. H. (in preparation). *Analyzing linguistic data: A practical introduction to statistics using R (Second edition)*. Cambridge: Cambridge University Press.
- Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on Naive Discriminative Learning. *Psychological Review*, 118(3), 438–481. doi: 10.1037/a0023851
- Balling, L. W., & Baayen, R. H. (2012). Probability and surprisal in auditory comprehension of morphologically complex words. *Cognition*, 125(1), 80–106. doi: 10.1016/j.cognition.2012.06.003

- Barkaoui, K. (2014). Quantitative approaches for analyzing longitudinal data in second language research. *Annual Review of Applied Linguistics*, 34, 65–101. doi: 10.1017/S0267190514000105
- 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. doi: 10.1016/j.jml.2012.11.001
- Bates, D. M. (2009). *[R-sig-ME] lmer: ML and REML estimation*. Retrieved from <https://stat.ethz.ch/pipermail/r-sig-mixed-models/2009q1/002096.html>
- Bates, D. M. (2010). *Lme4: Mixed-effects modeling with R*. Retrieved from <http://lme4.r-forge.r-project.org/IMMwR/lrgprt.pdf>
- Bates, D. M., Kliegl, R., Vasishth, S., & Baayen, R. H. (submitted). *Parsimonious mixed models*. Retrieved from <http://arxiv.org/abs/1506.04967>
- Bates, D. M., Mächler, M., Bolker, B., & Walker, S. (submitted). *Fitting linear mixed-effects models using lme4*. Retrieved from <http://arxiv.org/abs/1406.5823>
- Blevins, J. P. (2013). Word-based morphology from Aristotle to modern WP (Word and Paradigm models). In K. Allan (Ed.), *The oxford handbook of the history of linguistics* (pp. 375–395). Oxford: Oxford University Press.
- Bolker, B. M. (2014, January 26). *Convergence error for development version of lme4* [A response at a Q&A website]. Retrieved from <http://stackoverflow.com/a/21370041/3237850>
- Bolker, B. M., Brooks, M. E., Clark, C. J., Geange, S. W., Poulsen, J. R., Stevens, M. H. H., & White, J.-S. S. (2009). *Generalized linear mixed models: A practical guide for ecology*

- and evolution. *Trends in Ecology & Evolution*, 24(3), 127–35. doi: 10.1016/j.tree.2008.10.008
- Brown, R. (1973). *A first language: The early stages*. London: George Allen & Unwin.
- Bybee, J. (1985). *Morphology: A study of the relation between meaning and form*. Amsterdam: John Benjamins.
- Bybee, J. (2010). *Language, usage and cognition*. Cambridge: Cambridge University Press.
- Cunnings, I. (2012). An overview of mixed-effects statistical models for second language researchers. *Second Language Research*, 28(3), 369–382. doi: 10.1177/0267658312443651
- DeKeyser, R. (1997). Beyond explicit rule learning: Automatizing second language morphosyntax. *Studies in Second Language Acquisition*, 19(2), 195–221. doi: 10.1017/S0272263197002040
- Dingemanse, N. J., & Dochtermann, N. A. (2013). Quantifying individual variation in behaviour: Mixed-effect modelling approaches. *The Journal of Animal Ecology*, 82(1), 39–54. doi: 10.1111/1365-2656.12013
- Dörnyei, Z. (2009). *The psychology of second language acquisition*. Oxford: Oxford University Press.
- Dulay, H. C., & Burt, M. K. (1973). Should we teach children syntax? *Language Learning*, 23(2), 245–258.
- Ellis, N. C., & Schmidt, R. (1998). Rules or associations in the acquisition of morphology? The frequency by regularity interaction in human and PDP learning or morphosyntax. *Language and Cognitive Processes*, 13(2&3), 307–336. doi: 10.1080/016909698386546

- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. New York, NY: Cambridge University Press.
- Goldschneider, J., & DeKeyser, D. (2001). Explaining the "natural order of L2 morpheme acquisition" in English: A meta-analysis of multiple determinants. *Language Learning*, 51(1), 1–50. doi: 10.1111/1467-9922.00147
- Gries, S. Th. (2013). *Statistics for linguistics with R: A practical introduction (second edition)*. Berlin: De Gruyter Mouton.
- Gries, S. Th. (2015). The most under-used statistical method in corpus linguistics: Multi-level (and mixed-effects) models. *Corpora*, 10(1), 95–125. doi: 10.3366/cor.2015.0068
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction (second edition)*. New York, NY: Springer.
- Hastie, T., & Tibshirani, R. J. (1990). *Generalized additive models*. London: Chapman and Hall.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, 7(2), 185–207. doi: 10.3758/BF03212979
- Held, L., & Bové, D. S. (2014). *Applied statistical inference: Likelihood and Bayes*. Heidelberg: Springer.
- Hox, J. (2002). *Multilevel analysis: Techniques and applications*. New York: Lawrence Erlbaum Associates.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning: With applications in R*. New York, NY: Springer.

- Jarvis, S. (2000). Methodological rigor in the study of transfer: Identifying L1 influence in the interlanguage lexicon. *Language Learning*, 50(2), 245–309. doi: 10.1111/0023-8333.00118
- Jarvis, S., & Pavlenko, A. (2007). *Crosslinguistic influence in language and cognition*. New York: Routledge.
- Kliegl, R., Wei, P., Dambacher, M., Yan, M., & Zhou, X. (2011). Experimental effects and individual differences in linear mixed models: Estimating the relationship between spatial, object, and attraction effects in visual attention. *Frontiers in Psychology*, 1, Article 238. doi: 10.3389/fpsyg.2010.00238
- Kozaki, Y., & Ross, S. J. (2011). Contextual dynamics in foreign language learning motivation. *Language Learning*, 61(4), 1328–1354. doi: 10.1111/j.1467-9922.2011.00638.x
- Larsen-Freeman, D. E. (2011). A complexity theory approach to second language development/acquisition. In D. Atkinson (Ed.), *Alternative approaches to second language acquisition* (pp. 48–72). Abingdon: Routledge.
- Larson-Hall, J., & Herrington, R. (2010). Improving data analysis in second language acquisition by utilizing modern developments in applied statistics. *Applied Linguistics*, 31(3), 368–390. doi: 10.1093/applin/amp038
- Lightbown, P. (1983). Exploring relationships between developmental and instructional sequences in L2 acquisition. In H. Seliger & M. H. Long (Eds.), *Classroom oriented research in second language acquisition* (pp. 217–243). Rowley, MA: Newbury House.
- Linck, J. A., & Cunnings, I. (2015). The utility and application of mixed-effects models in second language research. *Language Learning*, 65(S1), 185–207. doi: 10.1111/lang.12117

- Long, J. D. (2012). *Longitudinal data analysis for the behavioral sciences using R*. Thousand Oaks, CA: Sage Publications.
- Luk, Z. P., & Shirai, Y. (2009). Is the acquisition order of grammatical morphemes impervious to L1 knowledge? Evidence from the acquisition of plural -s, articles, and possessive 's. *Language Learning*, 59(4), 721–754. doi: 10.1111/j.1467-9922.2009.00524.x
- McKeown, G. J., & Sneddon, I. (2014). Modeling continuous self-report measures of perceived emotion using generalized additive mixed models. *Psychological Methods*, 19(1), 155–174. doi: 10.1037/a0034282
- Mizumoto, A., & Plonsky, L. (in press). R as a lingua franca : Advantages of using R for quantitative research in applied linguistics. *Applied Linguistics*. doi: 10.1093/applin/amv025
- Mulder, K., Dijkstra, T., Schreuder, R., & Baayen, H. R. (2014). Effects of primary and secondary morphological family size in monolingual and bilingual word processing. *Journal of Memory and Language*, 72, 59–84. doi: 10.1016/j.jml.2013.12.004
- Murakami, A., & Alexopoulou, T. (in press). L1 influence on the acquisition order of English grammatical morphemes: A learner corpus study. *Studies in Second Language Acquisition*.
- Ning, L.-H., Shih, C., & Loucks, T. M. (2014). Mandarin tone learning in L2 adults: A test of perceptual and sensorimotor contributions. *Speech Communication*, 63-64, 55–69. doi: 10.1016/j.specom.2014.05.001
- Odlin, T. (1989). *Language transfer*. Cambridge: Cambridge University Press.

- Pica, T. (1983). Methods of morpheme quantification: Their effect on the interpretation of second language data. *Studies in Second Language Acquisition*, 6(1), 69–78. doi: 10.1017/S0272263100000309
- Pinheiro, J. C., & Bates, D. M. (2000). *Mixed-effects models in S and S-PLUS*. New York, NY: Springer New York.
- Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*, 15(4-5), 445–485. doi: 10.1080/01690960050119661
- R Core Team. (2015). *R: A language and environment for statistical computing*. Vienna, Austria. Retrieved from <http://www.r-project.org/>
- Singer, J. D., & Willett, J. B. (2003). *Applied longitudinal data analysis: Modeling change and event occurrence*. New York: Oxford University Press.
- Slobin, D. I. (1996). From "thought to language" to "thinking for speaking". In J. J. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 70–96). Cambridge: Cambridge University Press.
- Snijders, T. A. B., & Bosker, R. J. (1994). Modeled variance in two-level models. *Sociological Methods & Research*, 22(3), 342–363. doi: 10.1177/0049124194022003004
- Tremblay, A., Derwing, B., Libben, G., & Westbury, C. (2011). Processing advantages of lexical bundles: Evidence from self-paced reading and sentence recall tasks. *Language Learning*, 61(2), 569–613. doi: 10.1111/j.1467-9922.2010.00622.x
- van Geert, P. (2008). The dynamic systems approach in the study of L1 and L2 acquisition : An introduction. *Modern Language Journal*, 92(2), 179–199. doi: 10.1111/j.1540-4781.2008.00713.x

- van Geert, P., & van Dijk, M. (2002). Focus on variability: New tools to study intra-individual variability in developmental data. *Infant Behavior & Development*, 25(4), 340–374. doi: 10.1016/S0163-6383(02)00140-6
- Verspoor, M., Lowie, W., & van Dijk, M. (2008). Variability in second language development from a dynamic systems perspective. *Modern Language Journal*, 92(2), 214–231. doi: 10.1111/j.1540-4781.2008.00715.x
- Whittingham, M. J., Stephens, P. A., Bradbury, R. B., & Freckleton, R. P. (2006). Why do we still use stepwise modelling in ecology and behaviour? *The Journal of Animal Ecology*, 75(5), 1182–9. doi: 10.1111/j.1365-2656.2006.01141.x
- Wieling, M. (2015). *Analyzing EEG data using GAMs: Lecture 4 of advanced regression for linguists*. Retrieved 16 January, 2015, from <http://martijnwieling.nl/statscourse/lecture4/presentation.pdf>
- Wieling, M., Montemagni, S., Nerbonne, J., & Baayen, R. H. (2014). Lexical differences between Tuscan dialects and standard Italian: Accounting for geographic and socio-demographic variation using generalized additive mixed modeling. *Language*, 90(3), 669–692.
- Wieling, M., Nerbonne, J., & Baayen, R. H. (2011). Quantitative social dialectology: Explaining linguistic variation geographically and socially. *PLoS ONE*, 6(9), e23613. doi: 10.1371/journal.pone.0023613
- Wood, S. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1), 95–114. doi: 10.1111/1467-9868.00374

- Wood, S. (2004). Stable and efficient multiple smoothing parameter estimation for generalized additive models. *Journal of the American Statistical Association*, 99(467), 673–686. doi: 10.1198/016214504000000980
- Wood, S. (2006). *Generalized additive models: An introduction with R*. Boca Raton, FL: Chapman & Hall/CRC.
- Wood, S. (2009). GAMs: Semi-parametric GLMs. Retrieved 18 June, 2014, from <http://people.bath.ac.uk/sw283/mgcv/gam-theory.pdf>
- Wood, S. (2010). *A toolbox of smooths*. Retrieved 18 June, 2014, from <http://people.bath.ac.uk/sw283/mgcv/tampere/smooth-toolbox.pdf>
- Wood, S. (2013a). On p-values for smooth components of an extended generalized additive model. *Biometrika*, 100(1), 221–228. doi: 10.1093/biomet/ass048
- Wood, S. (2013b). A simple test for random effects in regression models. *Biometrika*, 100(4), 1005–1010. doi: 10.1093/biomet/ast038
- Wood, S. (2014). *Package ‘mgcv’*. Retrieved 30 November, 2014, from <http://cran.r-project.org/web/packages/mgcv/mgcv.pdf>
- Zuur, A. F., Ieno, E. N., Walker, N. J., Saveliev, A. A., & Smith, G. M. (2009). *Mixed effects models and extensions in ecology with R*. New York, NY: Springer.

Footnotes

¹ The accuracy of the R scripts is reported in Online Supporting Document 2.

² I thank an anonymous reviewer for suggesting this procedure.

³ The calculation may not look correct due to rounding, but the value is accurate.

⁴ There are two further features of GLMMs that merit discussion but the space does not allow to elaborate: the correlation structure of random effects and shrinkage. They are demonstrated in Online Supporting Document 3.

⁵ Although accuracy is proportional, logistic regression was not employed in order to avoid confusion between linearity in probability scale and linearity in logit scale. The same follows for the remaining panels.

⁶ The values were calculated with the `getSD.gam` function in the paper package of Wieling et al. (2014), available at <http://openscience.unileipzig.de/index.php/mr2/article/view/41>.

⁷ The detailed model selection procedure is provided in Online Supporting Document 4.

Tables

Table 1

Alignment of Englishtown Levels and the CEFR

Englishtown level number	CEFR levels
1-3	A1
4-6	A2
7-9	B1
10-12	B2
13-15	C1
16	C2

Table 2

Data Size per Learner

Morpheme	Non-zero observations (SD)	Obligatory contexts (SD)	Omission errors (SD)	Overgeneralization errors (SD)
Articles	20.08 (6.01)	140.92 (71.20)	12.23 (8.47)	5.53 (4.17)
Past tense <i>-ed</i>	6.89 (3.07)	16.09 (7.64)	0.66 (0.96)	1.04 (1.27)
Plural <i>-s</i>	18.91 (5.92)	90.53 (47.53)	3.44 (3.21)	1.94 (2.00)

Table 3

Comparison of GLMMs

Model	Model description		AIC	Δ AIC	Likelihood ratio test against the previous model	
	Fixed-effects	Random-effects			Statistic	p value
Model 1	None	By-L1 + by-learner random-intercepts	13799.8			
Model 2	Model 1 + morpheme	Same as Model 1	13356.4	-443.4	$\chi^2(2) = 447.39$	< 0.001
Model 3	Same as Model 2	Model 1 + by-morpheme random-contrasts at learner level	13276.5	-80.0	$\chi^2(5) = 89.97$	< 0.001
Model 4	Model 2 + L1type	Same as Model 3	13253.7	-22.8	$\chi^2(1) = 24.80$	< 0.001
Model 5	Model 4 + writingnum (standardized)	Same as Model 3	13240.3	-13.4	$\chi^2(1) = 15.41$	< 0.001
Model 6	Same as Model 5	Model 3 + by-writingnum random-slope at learner level	13212.7	-27.5	$\chi^2(4) = 35.51$	< 0.001
Model 7	Model 5 + proficiency (standardized)	Same as Model 6	13197.9	-14.8	$\chi^2(1) = 16.81$	< 0.001
Model 8	Model 7 + morpheme-proficiency interaction	Same as Model 6	13188.3	-9.7	$\chi^2(2) = 13.67$	0.001
Model 9	Model 8 + L1type-writingnum interaction	Same as Model 6	13185.5	-2.8	$\chi^2(1) = 4.75$	0.029
Reference Model	Morpheme + writingnum (standardized)	Same as Model 6	13229.5			

Table 4

Random Effects Structure of GLMMs

Factor	Random Effects	SD in Model 8	SD in Reference Model
L1			
	Intercept	0.295	0.300
Learner			
	Intercept	0.412	0.495
	Morpheme		
	Past tense <i>-ed</i>	0.613	0.716
	Plural <i>-s</i>	0.481	0.582
	Writingnum (standardized)	0.197	0.192

Table 5

Fixed Effects Structure of GLMM Model 8

Parameter		B		SE
Intercept		1.561	***	0.123
Morpheme				
	Past tense <i>-ed</i>	0.141		0.098
	Plural <i>-s</i>	0.787	***	0.063
L1type				
	PRESENT	0.679	***	0.123
Writingnum (standardized)		0.082	**	0.027
Proficiency (standardized)		0.238	***	0.043
Proficiency (standardized) : Morpheme				
	Proficiency (standardized) : Past tense <i>-ed</i>	-0.115		0.089
	Proficiency (standardized) : Plural <i>-s</i>	-0.224	***	0.059

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 6

Parametric Terms of GAMM Model 3

Parameter		B		SE
Intercept		1.532	***	0.127
L1type				
	PRESENT	0.685	***	0.118
Proficiency (standardized)		0.236	***	0.050
Morpheme				
	Past tense <i>-ed</i>	0.049		0.086
	Plural <i>-s</i>	0.741	***	0.062
Proficiency (standardized) : Morpheme				
	Proficiency (standardized) : Past tense <i>-ed</i>	-0.100		0.084
	Proficiency (standardized) : Plural <i>-s</i>	-0.220	***	0.059

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 7

Smooths Terms of GAMM Model 3

Term	EDF	Ref.df	χ^2	<i>p</i> -value
Writingnum (standardized) : L1type				
Writingnum (standardized) : ABSENT	1.001	1.002	0.026	0.872
Writingnum (standardized) : PRESENT	3.503	4.300	19.830	0.001
By-writingnum random wiggly curve for individual learners	233.053	1415.000	867.258	0.002

Table 8

Random Effects of GAMM Model 3

Random effects	SD	<i>p</i> -value
By-L1 random intercepts	0.285	< 0.001
By-morpheme random slopes for individual learners	0.176	< 0.001