



OPEN ACCESS

EDITED BY

David Sanchez,
Spanish National Research Council
(CSIC), Spain

REVIEWED BY

Lucia Loureiro-Porto,
University of the Balearic Islands, Spain
Angelica Sousa Da Mata,
Universidade Federal de Lavras, Brazil

*CORRESPONDENCE

Jonathan Dunn,
✉ jonathan.dunn@canterbury.ac.nz

RECEIVED 07 August 2023

ACCEPTED 15 September 2023

PUBLISHED 29 September 2023

CITATION

Dunn J (2023), Syntactic variation across
the grammar: modelling a complex
adaptive system.

Front. Complex Syst. 1:1273741.

doi: 10.3389/fcpxs.2023.1273741

COPYRIGHT

© 2023 Dunn. This is an open-access
article distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original author(s)
and the copyright owner(s) are credited
and that the original publication in this
journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Syntactic variation across the grammar: modelling a complex adaptive system

Jonathan Dunn*

Department of Linguistics and the New Zealand Institute for Language, Brain and Behaviour, University of Canterbury, Christchurch, New Zealand

While language is a complex adaptive system, most work on syntactic variation observes a few individual constructions in isolation from the rest of the grammar. This means that the grammar, a network which connects thousands of structures at different levels of abstraction, is reduced to a few disconnected variables. This paper quantifies the impact of such reductions by systematically modelling dialectal variation across 49 local populations of English speakers in 16 countries. We perform dialect classification with both an entire grammar as well as with isolated nodes within the grammar in order to characterize the syntactic differences between these dialects. The results show, first, that many individual nodes within the grammar are subject to variation but, in isolation, none perform as well as the grammar as a whole. This indicates that an important part of syntactic variation consists of interactions between different parts of the grammar. Second, the results show that the similarity between dialects depends heavily on the subset of the grammar being observed: for example, New Zealand English could be more similar to Australian English in phrasal verbs but at the same time more similar to UK English in dative phrases.

KEYWORDS

complex adaptive system, computational syntax, computational sociolinguistics, construction grammar, syntactic variation, dialect classification

1 Introduction

Within linguistics and cognitive science, language is increasingly viewed as a complex adaptive system (Bybee, 2007; Beckner et al., 2009). For example, usage-based theories of syntax like Construction Grammar (Goldberg, 2006; Langacker, 2008) view the grammar as a network which contains structures at different levels of abstraction. The network structure of the grammar is made up of inheritance relations (mother-child) and similarity relations (sibling-sibling) between constructions. A *construction* in this context is a symbolic mapping between form and meaning, where an individual construction is unique either syntactically or semantically. For example, there is an inheritance relationship between the schematic ditransitive construction, with examples like (1a), and idiomatic constructions, with examples like (1b) and (1c). While some of the properties of the ditransitive are inherited by these idiomatic children, they also retain unique and non-compositional meanings. It is this network structure which makes the grammar a complex system. As with any complex system, there are emergent properties of the grammar which cannot be described by looking at individual constructions in isolation.

- (1a) write the store a check
- (1b) *give me a hand*
- (1c) *give me a break*

The challenge is that most work on syntactic variation does exactly this: observing a few constructions that have been removed from the context of the larger grammar and modelled as discrete and independent variables. The contribution of this paper is to systematically evaluate whether the picture we get of syntactic variation changes depending on which sub-sets of the grammar we inspect. In other words, to what degree does our view of syntactic variation (for example, the similarity between New Zealand English and Australian English) depend on the sub-set of the grammar which we are observing?

This paper makes two significant contributions to models of dialectal variation: First, we examine dialect areas at three levels of spatial granularity. This includes *regional dialects* (like North American English), *national dialects* (like Canadian English), and *local dialects* (like Ontario English). This is the first computational study to systematically experiment with different levels of granularity when modelling dialectal variation. Second, we examine different nodes or clusters of constructions within the grammar. This includes macro-clusters which contain hundreds of constructions and smaller micro-clusters which contain dozens of constructions. This is the first computational study to systematically experiment with the distribution of spatial variation across an entire grammar.

In the first case, in order to understand syntactic variation we must view the population of speakers itself as a complex network. While most work on syntactic variation considers only a few segments of the population, this paper uses observations from 49 local populations distributed across 16 countries. Speakers of English within one regional dialect are in contact with speakers of other dialects through immigration, short-term travel, media, and digital communication. Thus, the first challenge is to conduct a dialect survey across all representative populations of English speakers in order to understand syntactic variation across the entire population network.

In the second case, in order to understand syntactic variation we must view the grammar itself as a complex network so that we can observe variation in its entirety rather than in isolated and disconnected portions of the grammar. In this paper we use Computational Construction Grammar (Dunn, 2017; Dunn, 2022) to provide an unsupervised network of constructions. For these experiments, this grammar network is learned using independent data from the same register as the geographic corpora used to represent dialects (tweets).

Our theoretical question is whether syntactic variation, as captured by systematic grammatical comparisons between dozens of regional populations, is influenced by the sub-set of the grammar network which is used to model variation. We operationalize a model of dialect as a classifier which learns to predict the dialect membership of held-out samples. A dialect classifier is able to handle high-dimensional spaces, important for viewing variation across an entire grammar. And, importantly, the quality of a dialect classifier can be measured using its prediction accuracy on held-out samples. Our goal here is not simply to find sub-sets of the grammar which are in

variation but rather to determine how accurate and robust these sub-sets of the grammar are for characterizing dialectal variation as a whole: for example, how accurately does a portion of the grammar characterize the difference between New Zealand and Australian English? To answer this question, we use prediction accuracy, error analysis, and feature pruning methods to determine the quality of dialect models that rely on different nodes within the grammar.

2 Related work

Current knowledge of large-scale linguistic variation (i.e., across many countries) consists of i) studies of lexical rather than syntactic variation, ii) quantitative corpus-based studies of syntactic variation, and iii) computational studies of syntactic variation. In the first case, lexical variation is the most approachable type of linguistic variation because it does not require any learning of grammars or representations. Thus, early large-scale corpus-based studies of variation focused on the usage of lexical items (Eisenstein et al., 2010; Mocanu et al., 2013; Eisenstein et al., 2014; Donoso et al., 2017; Rahimi et al., 2017). The challenge for lexical variation is to define the envelope of variation (i.e., discover the set of alternations to avoid topic-specific models). Two main approaches to this are, first, to rely on existing dialect surveys to provide hand-crafted alternations (Grieve et al., 2019) and, second, to use contextual embeddings to develop clusters of senses of words (Lucy and Bamman, 2021). In either case, lexical variation is a simpler phenomenon than syntactic variation because the number of potential alternations for each lexical item is limited. Recent work on semantic variation (Dunn, 2023c) has expanded this scope by looking at the conceptual rather than the lexical level and using participant-based measures like abstractness ratings and age-of-acquisition to determine what causes a concept to be subject to dialectal variation.

Most corpus-based approaches to syntactic variation choose a single construction to examine and then model variation within that construction alone (Buchstaller, 2008; Grieve, 2012; Schilk and Schaub, 2016; Calle-Martin and Romero-Barranco, 2017; Grafmiller and Szmrecsanyi, 2018; Deshors and Götz, 2020; Schneider et al., 2020; Rautionaho and Hundt, 2022; Xu et al., 2022; Larsson, 2023; Li et al., 2023). While this line of work can reveal small-scale *syntactic* variation and change, it can never account for *grammatical* variation. The difference between these two terms is important in this context: a single syntactic feature, like aspectual marking or noun pluralization, may be in variation but we cannot understand the variation without contextualizing it within the entire grammar. In other words, if the grammar is in fact a complex adaptive system, then measuring variation in a single construction is like assuming that the weather in Miami, Florida is independent of both the weather in Orlando and the current conditions of the Atlantic. By analogy, previous work has shown differences of behaviour in small-scale population networks vs. large-scale networks (Laitinen et al., 2020). Such network effects are important for establishing patterns of diffusion in addition to patterns of geographic variation (Laitinen and Fatemi, 2022) by modelling the social connections within each local population. This

paper examines the impact of the granularity or size of the network for both the underlying population (e.g., regional vs. local dialects) and the grammar itself (e.g., different clusters of constructions within the grammar).

The first computational work which viewed syntactic variation from the perspective of a complex adaptive system used 135 grammatical alternations in English, chosen manually to include features which can be extracted using regular expressions (Grieve, 2016). The alternations include examples like *anyone* vs. *anybody* and *hear of* vs. *hear about*. This study used a corpus of letters to the editor from 240 US cities, similar in spatial granularity to the local dialects in this paper. While this early work assumed a starting set of simple alternations, it was followed by work which focused on discovering the set of variants while instead assuming the spatial boundaries of dialect areas (Dunn, 2018a). The advantage of this approach is that it both expands the scope of the study (by including more complex constructional features) while also scaling across languages (Dunn, 2019b).

The other difference in these two approaches is that Grieve's early work relies on factor analysis to group together grammatical alternations according to their patterns of variation. In order to provide a measure of predictive accuracy on a held-out test set, by which a better model makes better predictions, more recent computational work has instead taken a classification approach (Dunn, 2019c; Dunn and Wong, 2022). As discussed further in the section on Computational Construction Grammar, constructions are organized into a network structure using similarity measures directly within the grammar. This means

that the nodes within which variation occurs are derived independently of the model of dialectal variation.

There are two questions which remain given this previous work: First, how much of the predictive accuracy of a dialect model is contained in different nodes within the grammar? One issue with surface-level alternations like *anyone* vs. *anybody* is that, while in variation, they do not capture more schematic differences between dialects and, overall, do not hold much predictive power given their relative scarcity. Second, previous work has always focused on a given size of spatial granularity, usually at the country or city level. This paper uses three levels of granularity to help understand complexity in the underlying population network as well.

3 Data

The data used for these experiments is drawn from geo-referenced social media posts (tweets), a source with a long history as an observation of dialectal production (c.f., Eisenstein et al., 2014; Gonçalves and Sánchez, 2014; Gonçalves et al., 2018; Dunn, 2019b; Grieve et al., 2019). The corpus is drawn from 16 English-speaking countries, as shown in Table 1. Countries are grouped into larger regional dialects (such as North American vs. South Asian English). And each country is divided into potentially many sub-areas using spatial clustering (for example, American English is divided into nine local dialect groups). Language identification is undertaken using two existing

TABLE 1 Distribution of Sub-Corpora by Region. Each sample is a unique sub-corpus with the same distribution of keywords, each approximately 3,910 words.

Region	Country		Areas	Corpus size	N. Samples
Africa, Southern	South Africa	ZA	2	9,205,166 words	2,299
	Zimbabwe	ZW	1	3,252,685 words	767
Africa, Sub-Saharan	Kenya	KE	3	5,537,772 words	1,310
	Nigeria	NG	3	6,139,948 words	1,492
North America	Canada	CA	4	16,801,386 words	4,261
	United States	US	9	26,050,840 words	5,802
Asia, South	Bangladesh	BD	2	6,287,670 words	1,649
	India	IN	7	28,935,606 words	7,045
	Pakistan	PK	2	12,765,491 words	3,271
Asia, Southeast	Indonesia	ID	1	2,392,074 words	546
	Malaysia	MY	1	8,580,789 words	2,052
	Philippines	PH	2	9,907,209 words	2,402
Europe	Ireland	IE	1	13,287,397 words	3,360
	United Kingdom	UK	5	20,307,094 words	4,890
Oceania	Australia	AU	4	23,163,447 words	5,914
	New Zealand	NZ	2	8,113,382 words	2,047
Total	16 Countries		49 Areas	200,727,956 words	49,107

models to make the corpus comparable with existing work on mapping digital language use (Dunn, 2020; Dunn and Nijhof, 2022). This corpus thus provides three levels of granularity: 7 regions, 16 countries, and 49 local areas.

The main challenge is to control for other sources of variation like topic or register that would lead to a successful classification model but would not be directly connected with the local population being observed. In other words, we need to constrain the production of the local populations to a specific set of topics: if New Zealand tweets are focused on economics and Australian tweets on rugby, the impact of register would be a potential confound. For this reason, we develop a set of 250 common lexical items (c.f., Table in Appendix 1) which are neither purely functional (like *the is*) nor purely topical (like *Biden is*). These keywords were selected by first creating a list of the most frequent words in a background corpus of English-language tweets and then removing words which would qualify as either too topical or too functional. The words are thus ordered by frequency, with the highest frequency words shown first. For each location we create sub-corpora which are composed of one unique tweet for each of these keywords. Thus, each location is represented by a number of sub-corpora which each have the same fixed distribution of key lexical items. This allows us to control for wide variations in topic or content or register, factors that would otherwise potentially contribute non-dialectal sources of variation.

Each sub-corpus thus contains 250 tweets, one for each keyword. This creates sub-corpora with an average of approximately 3,910 words. The distribution of sub-corpora (called *samples*) is shown in Table 1. For example, the US is represented by a corpus of 26 million words divided into 5,802 individual samples. Because these samples have the same distribution of lexical items, the prediction accuracy of the dialect classifier should not be influenced by topic-specific patterns in each region. The use of lexically-balanced samples, while important for forcing a focus on dialectal variation, reduces the overall size of the corpus that is available because tweets without the required keywords are discarded.

To form the local areas, we organize the data around the nearest airport (within a threshold of a 25 km radius) as a proxy for urban areas. We then use the density-based H-DBSCAN algorithm to cluster airports into groups that represent contiguous local areas (Campello et al., 2013; Campello et al., 2015; McInnes and Healy, 2017). The result is a set of local areas within a country, each of which is composed of multiple adjacent urban centers. For example, the nine areas within the United States are shown in Figure 1, where each color represents a different group. Manual adjustments of unclustered or borderline points is then undertaken to produce the final clusters. The complete set of local areas are documented in the Supplementary Material.¹ It is important to note that these local areas are entirely spatial in the sense that no linguistic information has been included in their formation. These areas represent local geographic groups, but not a linguistically-defined speech community.

Our experiments in this paper operate at three levels of spatial granularity: first, distinguishing between regional dialects, like North American English vs. South Asian English; second, distinguishing between country-level dialects, like American

English vs. Canadian English; third, distinguishing between local dialects within regions, like Midwestern American English vs. Central Canadian English. These different levels of granularity allow us to test how well different portions of the grammar are able to model increasingly fine distinctions between dialects.

4 Methods

The basic approach in this paper is to use an unsupervised grammar derived from the Construction Grammar paradigm (CxG) as a feature space for dialect classification. A dialect classifier is trained to predict the regional dialect of held-out samples using the frequency of constructions within the sample. This section describes in more detail both the grammar model and the dialect models.

4.1 Computational construction grammar

A construction grammar is a network of form-meaning mappings at various levels of schematicity. As discussed above, the grammar is a network with inheritance relationships and similarity relationships between pairs of constructions. CxG is a usage-based approach to syntax which, in practical terms, means that more item-specific constructions are learned first and then generalized into more schematic constructions (Nevens et al., 2022; Doumen et al., 2023). The grammar learning algorithm used in this paper is taken from previous work (Dunn, 2017; Dunn, 2018b; Dunn, 2019a; Dunn and Nini, 2021; Dunn and Tayyar Madabushi, 2021; Dunn, 2022), with the grammar trained from the same register as the dialectal data (tweets). Rather than describe the computational details of this line of work, this section instead analyzes constructions within the grammar as examples of the kinds of features used to model syntactic variation. The complete grammar together with examples is available in the Supplementary Material and the codebase for computational CxG is available as a Python package.²

A break-down of the grammar used in the experiments is shown in Figure 2, containing a total of 15,215 individual constructions. Constructions are represented as a series of slot-constraints and the first distinction between constructions involves the types of constraints used. Computational CxG uses three types of slot-fillers: lexical (LEX, for item-specific constraints), syntactic (SYN, for form-based or local co-occurrence constraints), and semantic (SEM, for meaning-based or long-distance co-occurrence constraints). As shown in (2), slots are separated by dashes in the notation used here. Thus, SYN in (2) describes the type of constraint and *determined-permitted* provides its value using two central exemplars of that constraint. Examples or tokens of the construction from a test corpus of tweets are shown in (2a) through (2d).

1 The DOI for this material is DOI 10.17605/OSF.IO/2AKMY.

2 <https://github.com/jonathandunn/c2xg/tree/v2.0>

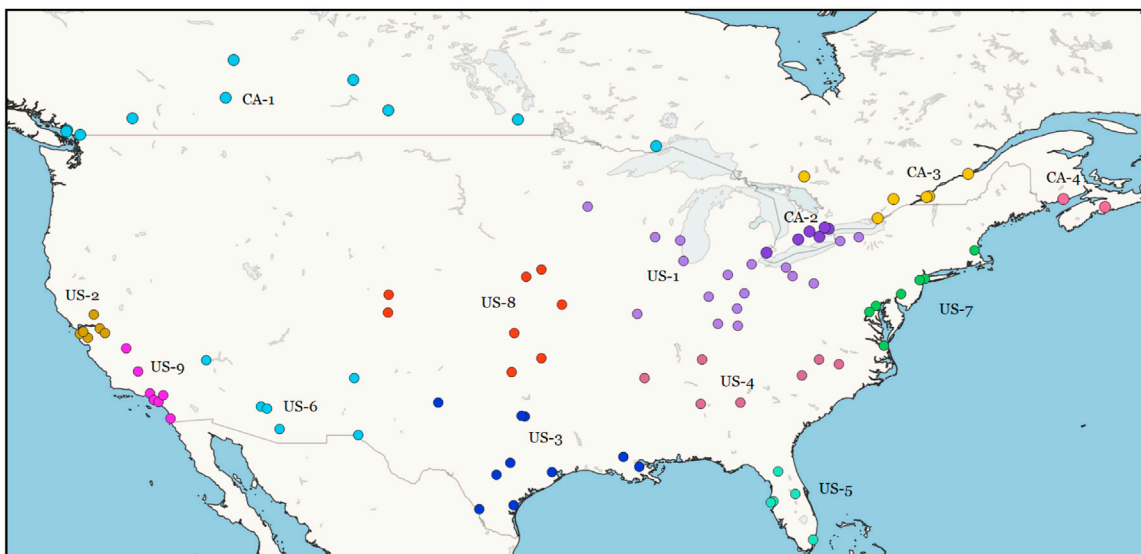


FIGURE 1
Distribution of local dialects in North America.

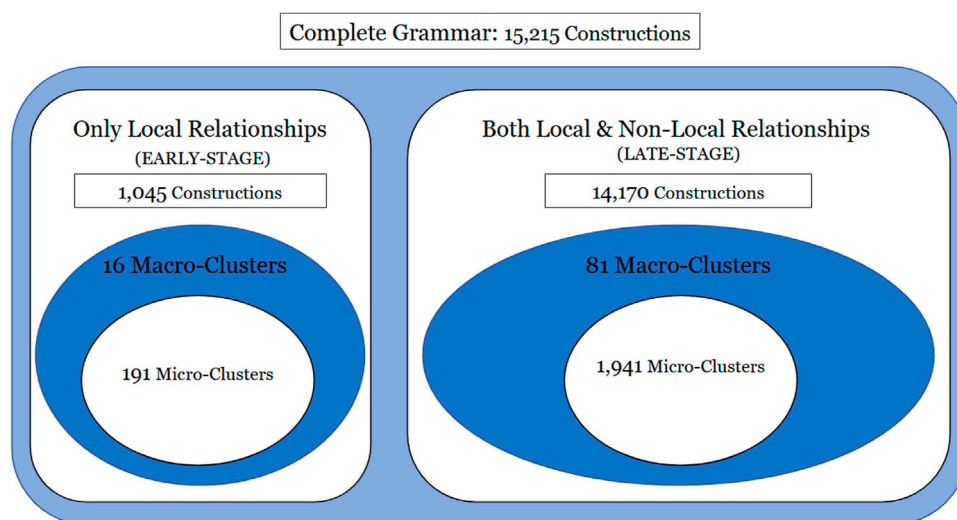


FIGURE 2
Break-down of the grammar used in the experiments by construction type.

- (2) [SYN: determined-permitted—SYN: to—SYN: pushover-backtrack]
 (2a) refused to play
 (2b) tried to watch
 (2c) trying to run
 (2d) continue to drive

Thus, the construction in (2) contains three slots, each defined using a syntactic constraint. These constraints are categories learned at the same time that the grammar itself is learned, formulated within an embedding space. An embedding that captures local co-

occurrence information is used for formulating syntactic constraints (a continuous bag-of-words fastText model with a window size of 1) while an embedding which instead captures long-distance co-occurrence information is used for formulating semantic constraints (a skip-gram fastText model with a window size of 5). Constraints are then formulated as centroids within that embedding space. Thus, the tokens for the construction in (2) are shown in (2a) through (2d). For the first slot-constraint, the name (*determined-permitted*) is derived from the lexical items closest to the centroid of the constraint. The proto-type structure of categories is modeled using cosine distance as a measure of how

well a particular slot-filler satisfies the constraint. Here the lexical items “reluctant,” “ready,” “refusal,” and “willingness” appear as fillers sufficiently close to the centroid to satisfy the slot-constraint. The construction itself is a complex verb phrase in which the main verb encodes the agent’s attempts to carry out the event encoded in the infinitive verb. This can be contrasted semantically with the construction in (3), which has the same form but instead encodes the agent’s preparation for carrying out the social action encoded in the infinitive verb.

- (3) [SYN: determined-permitted—SYN: to—SYN: demonstrate-reiterate]
 (3a) reluctant to speak
 (3b) ready to exercise
 (3c) refusal to recognize
 (3d) willingness to govern

An important idea in CxG is that structure is learned gradually, starting with item-specific surface forms and moving to increasingly schematic and productive constructions. This is called *scaffolded learning* because the grammar has access to its own previous analysis for the purpose of building more complex constructions (Dunn, 2022). In computational CxG this is modelled by learning over iterations with different sets of constraints available. For example, the constructions in (2) and (3) are learned with only access to the syntactic constraints, while the constructions in (4) and (5) have access to lexical and semantic constraints as well. This allows grammars to become more complex while not assuming basic structures or categorizations until they have been learned. In the dialect experiments below we distinguish between EARLY-STAGE grammars (which only contain syntactic constraints) and LATE-STAGE grammars (which contain lexical, syntactic, and semantic constraints).

- (4) [LEX: “the”—SEM: way—LEX: “to”]
 (4a) the chance to
 (4b) the way to
 (4c) the path to
 (4d) the steps to

Constructions have different levels of abstractness or schematicity. For example, the construction in (4) functions as a modifier, as in the X position in the sentence “Tell me [X] bake yeast bread.” This construction is not purely item-specific because it has multiple types or examples. But it is less productive than the location-based noun phrase construction in (5) which will have many more types in a corpus of the same size. CxG is a form of lexico-grammar in the sense that there is a continuum between item-specific and schematic constructions, exemplified here by (4) and (5), respectively. The existence of constructions at different levels of abstraction makes it especially important to view the grammar as a network with similar constructions arranged in local nodes within the grammar.

- (5) [LEX: “the”—SEM: streets]
 (5a) the street
 (5b) the sidewalk
 (5c) the pavement

- (5d) the avenues

A grammar or construction is not simply a set of constructions but rather a network with both taxonomic and similarity relationships between constructions. In computational CxG this is modelled by using pairwise similarity relationships between constructions at two levels: i) representational similarity (how similar are the slot-constraints which define the construction) and ii) token-based similarity (how similar are the examples or tokens of two constructions given a test corpus). Matrices of these two pairwise similarity measures are used to cluster constructions into smaller and then larger groups. For example, the phrasal verbs in (6) through (8) are members of a single cluster of phrasal verbs. Each individual construction has a specific meaning: in (6), focusing on the social attributes of a communication event; in (7), focusing on a horizontally-situated motion event; in (8), focusing on a motion event interpreted as a social state. These constructions each have a unique meaning but a shared form. The point here is that at a higher-order of structure, there are a number of phrasal verb constructions which share the same schema. These constructions have sibling relationships with other phrasal verbs and a taxonomic relationship with the more schematic phrasal verb construction. These phrasal verbs are an example of a *micro-cluster* referenced in the dialect experiments below (c.f., Dunn, 2023a)³.

- (6) [SEM: screaming-yelling—SYN: through]
 (6a) stomping around
 (6b) cackling on
 (6c) shouting out
 (6d) drooling over
- (7) [SEM: rolled-turned—SYN: through]
 (7a) rolling out
 (7b) slid around
 (7c) wiped out
 (7d) swept through
- (8) [SEM: sticking-hanging—SYN: through]
 (8a) poking around
 (8b) hanging out
 (8c) stick around
 (8d) hanging around

An even larger structure within the grammar is based on groups of these micro-clusters, structures which we will call *macro-clusters*. A macro-cluster is much larger because it contains many sub-clusters which themselves contain individual constructions. An example of a macro-cluster is given with five constructions in (9) through (13) which all belong to same neighborhood of the grammar. The partial noun phrase in (9) points to a particular sub-set of some entity (as in, “parts of the recording”). The partial adpositional phrase in (10) points specifically to the end of some temporal entity (as in, “towards the end of the show”). In contrast, the partial noun phrase in (11) points a particular sub-set of a spatial

³ Available at <https://www.jdunn.name/cxg>

location (as in, “the edge of the sofa”). A more specific noun phrase in (12) points to a sub-set of a spatial location with a fixed level of granularity (i.e., at the level of a city or state). And, finally, in (13) an adpositional phrase points to a location within a spatial object. The basic idea here is to use these micro-clusters and macro-clusters as features for dialect classification in order to determine how variation is distributed across the grammar.

- (9) [SEM: *part*—LEX: “of”—SYN: *the*]
 (9a) parts of the
 (9b) portion of the
 (9c) class of the
 (9d) division of the
- (10) [SYN: *through*—SEM: *which-whereas*—LEX: “end”—LEX: “of”—SYN: *the*]
 (10a) at the end of the
 (10b) before the end of the
 (10c) towards the end of the
- (11) [SEM: *which-whereas*—SEM: *way*—LEX: “of”]
 (11a) the edge of
 (11b) the side of
 (11c) the corner of
 (11d) the stretch of
- (12) [SEM: *which-whereas*—SYN: *southside-northside*—SYN: *chicagoland*]
 (12a) in north texas
 (12b) of southern california
 (12c) in downtown dallas
 (12d) the southside chicago
- (13) [LEX: “of”—SYN: *the*—SYN: *courtyard-balcony*]
 (13a) of the gorge
 (13b) of the closet
 (13c) of the room
 (13d) of the palace

The examples in this section have illustrated some of the fundamental properties of CxG and also provide a discussion of some of the features which are used in the dialect classification study. A more detailed linguistic examination of the contents of a grammar like this is available elsewhere (Dunn, 2023b). A breakdown of the contents of the grammar is shown in Figure 2. The 15,215 total constructions are first divided into different scaffolds (early-stage vs. late-stage), with a smaller number of local-only constructions which tend to be more schematic (1,045 vs. 14,170 constructions in the late-stage grammar). This grammar has a network structure and contains 2,132 micro-clusters (e.g., the phrasal verbs discussed above). At an even higher level of structure, there are 97 macro-clusters or neighborhoods within the grammar (e.g., the sub-set referencing constructions discussed above). We can thus look at variation across the entire grammar, across different levels of scaffolded structure, and across different levels of abstraction. The main reason for doing this is to determine whether all nodes within the grammar vary across dialects in the same way.

4.2 Dialect classification

A dialect classifier is a supervised discriminative approach to modelling dialects: given labelled training data, the model learns to distinguish between dialects like American and Canadian English using syntactic features from computational CxG. There are two advantages to taking a classification approach: First, classifiers work well in high-dimensional spaces while more traditional methods from quantitative sociolinguistics do not scale across tens of thousands of potentially redundant structural features. Second, dialect classifiers provide a ground-truth measure of quality in the form of prediction accuracy: we know how robust a classifier model is given how well it is able to distinguish between different dialects. The classification of dialects or varieties of a language is a robust area, although most work views this as an engineering challenge rather than as a way to learn about dialects themselves (Belinkov and Glass, 2016; Gamallo et al., 2016; Barbaresi, 2018; Kroon et al., 2018; Zampieri et al., 2020).

Following previous work on using dialect classification to model linguistic variation (Dunn, 2018a; Dunn, 2019b; Dunn, 2019c; Dunn and Wong, 2022), we use a Linear Support Vector Machine for classification. The advantage of an SVM over neural classifiers is that we can inspect the features which are useful for dialect classification; and the advantage over Logistic Regression or Naive Bayes is a better handling of redundant or correlated features.

The data is divided into a training set (80%) and a testing set (20%) with the same split used for each classification experiment. This means that the models for each level of spatial granularity (region, country, local area) are directly comparable across feature types. These dialect classifiers become our means of observing a grammar’s ability to capture spatial variation: a better grammar models dialectal variation with a higher prediction accuracy. This means that it is better able to describe the differences between dialects. For instance, the usage of phrasal verbs in (6) to (8) might differ significantly between Canada and New Zealand English, while at the same time accounting for only a minimal percentage of the overall syntactic difference between these dialects. A predictive model like a classifier, however, is evaluated precisely on how well it characterizes the total syntactic difference between each dialect.⁴

We can further explore each dialect model using a confusion matrix to examine the types of errors made. For instance, Figure 3 shows the distribution of errors for the country-level classification of dialects using the late-stage grammar. Correct predictions occur on the diagonal; given the weighted f-score of 0.97, most predictions are correct in this case. Yet the number of errors for each pair of dialects reflects their similarity. Thus, the most similar countries are i) New Zealand and Australia with 44+33 errors, ii) Ireland and the UK with 25+22 errors, iii) Pakistan and India with 25+21 errors, iv) Pakistan and Bangladesh with 21+8 errors, and v) Canada and the US with 15+1 errors. The confusion matrix also reveals the dominant variety, in the sense that only one sample of American English is mistakenly predicted to be Canadian (rows represent the true labels) while 15 samples of Canadian English are mistaken for American English. Thus, these are asymmetrical errors. The point here is that these models allow us to measure not only the overall quality of the

⁴ This holds true unless multiple models reach the same ceiling of accuracy, a situation which does not occur in this study.

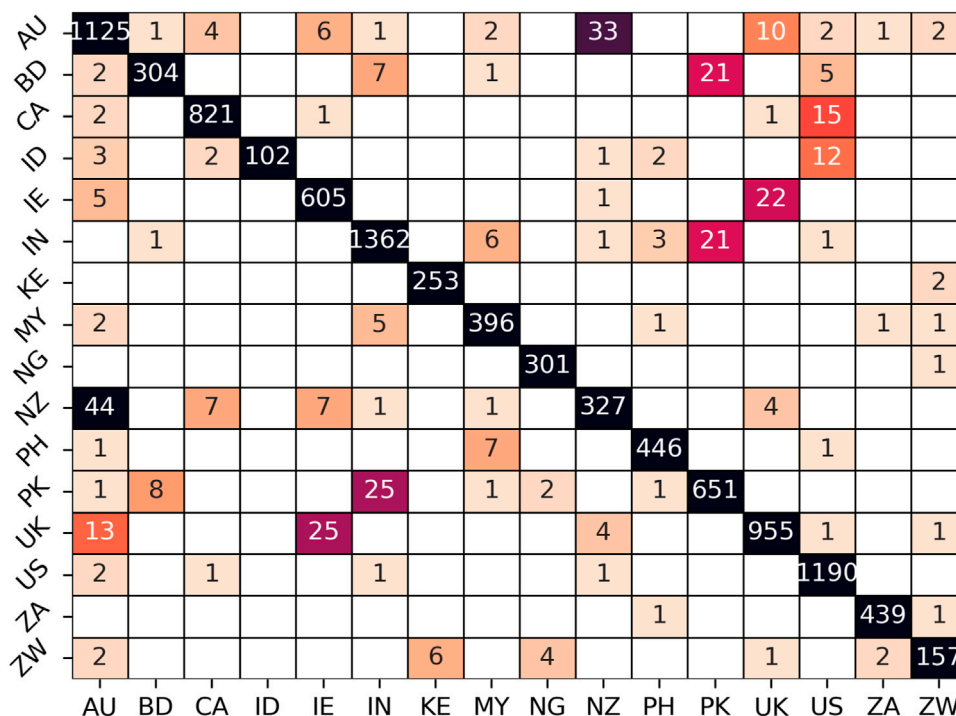


FIGURE 3 Distribution of Errors in Country-Level Dialect Model with Late-Stage Grammar. Rows represent ground-truth labels and columns represent predicted labels.

dialect classifier (its prediction accuracy represented by the f-score) but also determine which dialects are the most similar. This, in turn, means that we can measure the stability of dialect similarity across different nodes within the grammar. Note that this error-based similarity is different from feature-based similarity measures (Dunn, 2019b; Szmrecsanyi et al., 2019) which instead operate on the grammatical representations themselves. However, the more reliable a classifier is (i.e., the higher its prediction accuracy), the more its errors can be used directly as a similarity measure. Our focus here is on understanding the dialect model itself by examining its false positive and false negative errors.

5 Results

The basic question in this paper is whether different nodes within the grammar equally capture dialectal variation and whether the resulting picture of dialectal relations is the same regardless of which node we examine. Now that we have discussed the grammatical features used (from computational CxG) and the means of observing syntactic variation (a dialect classifier), this section analyzes the results. We start with region-level dialects before moving to country-level and then local dialects.

5.1 Regional dialects

The 16 countries used in this study are grouped into regions as shown in Table 1, resulting in seven larger macro-dialects as shown

in Table 2. The table shows three measures: precision (false positive errors), recall (false negative errors), and f-score. On the left the table shows the dialect performance with the late-stage grammar (i.e., with constructions containing all three types of slot-constraints) and on the right the early-stage grammar (i.e., with constructions containing only syntactic slot constraints). The late-stage grammar performs better (0.99 f-score vs. 0.93) but both perform well. In the early-stage grammar, two dialects are a particular source of the lower performance: Southern African English (South Africa and Zimbabwe) and Oceanic English (Australia and New Zealand).

The overall f-score of the late-stage grammar (0.99) tells us that, at this level of spatial granularity, the grammar as a single network is able to make almost perfect distinctions between the syntax of these regional dialects. But how does this ability to capture variation spread across nodes within the grammar? This is explored in Figure 4, which contrasts the f-score of individual nodes using the macro-clusters and micro-clusters discussed above. A macro-cluster is a fairly large group of constructions based on pairwise similarity between the constructions themselves. A micro-cluster is a smaller group within a macro-cluster, based on pairwise similarity between instances or tokens of constructions in a test corpus (Dunn, 2023a). Each point in this figure is a single node of the grammar (blue for macro-clusters and orange for micro-clusters). The position of the point on the y-axis reflects the prediction performance for regional dialect classification using only that portion of the grammar. The dotted horizontal line represents the majority baseline, at which classification performance is no better than chance.

TABLE 2 Performance of Dialect Classifier With Regional Dialects: Late-Stage Constructions (Left, with all constraint types) and Early-Stage Constructions (Right, with only syntactic constraints).

Region	Late-stage grammar			Early-stage grammar		
	Precision	Recall	F-Score	Precision	Recall	F-Score
Africa, Southern	0.99	0.97	0.98	0.88	0.81	0.84
Africa, Sub	0.98	0.99	0.99	0.92	0.91	0.91
North America	0.99	1.00	0.99	0.93	0.95	0.94
Asia, South	1.00	1.00	1.00	0.99	0.98	0.98
Asia, Southeast	0.99	0.98	0.98	0.92	0.92	0.92
Western Europe	0.98	0.99	0.99	0.94	0.94	0.94
Oceania	0.98	0.98	0.98	0.88	0.89	0.88
Weighted Avg	0.99	0.99	0.99	0.93	0.93	0.93

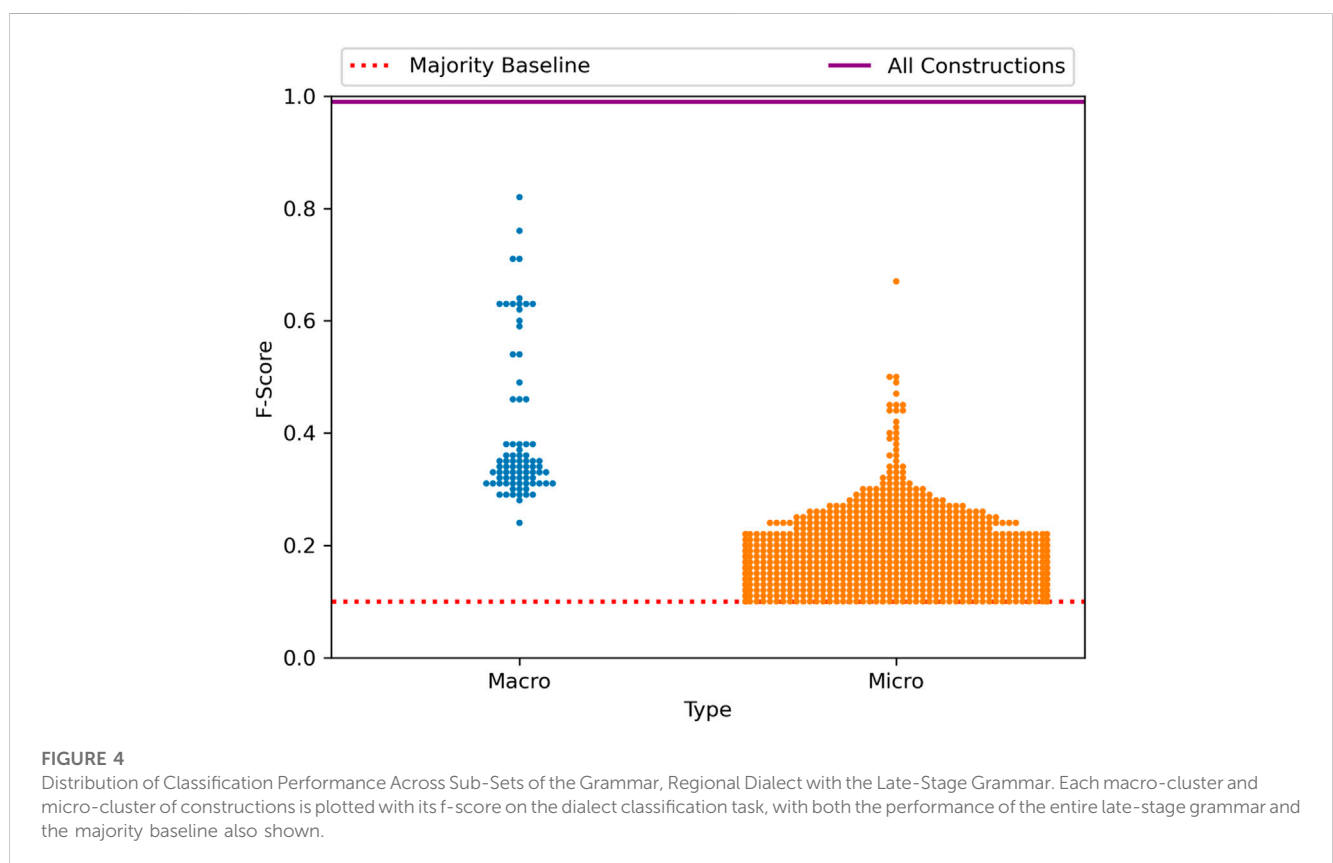


Figure 4 shows us, first, that the grammar as a whole is always better able to describe syntactic variation than individual nodes within that grammar. This is important in itself because, if language is a complex adaptive system, then it follows that variation in language is at least in part an emergent phenomenon in which the interaction between elements (here constructions) cannot be characterized by observing those elements in isolation. The second finding which this figure shows is that individual nodes vary greatly in the amount of dialectal variation which they are subject to. Thus, several nodes

in isolation are able to characterize between 70% and 80% of the overall variation, but others characterize very little. Within micro-clusters, especially, we see that many nodes within the grammar are essentially not subject to variation and thus provide little predictive power on their own.

The same type of graph is shown again in Figure 5, now for the early-stage grammar with only local constraints. This figure replicates the same findings: First, no individual node is able to capture the overall variation as well as the complete network. Second, there is a wide range across macro-clusters and micro-

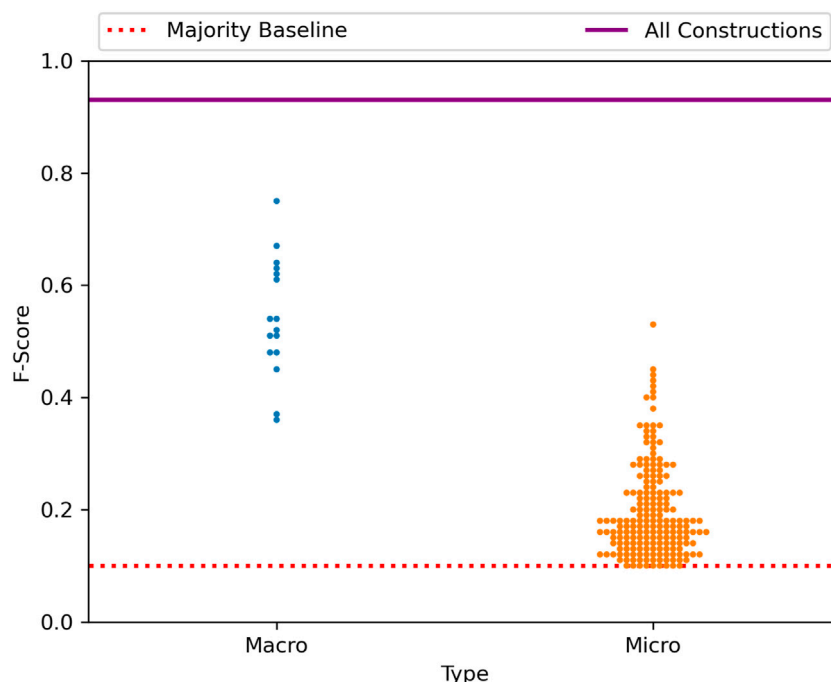


FIGURE 5 Distribution of classification performance across sub-sets of the grammar, regional dialects with the early-stage grammar.

TABLE 3 Distribution of errors in dialect classifier with regional dialects: Late-stage constructions (left) and early-stage constructions (right).

Late-stage grammar		Early-stage grammar	
Region Pairs	% of Errors	Region Pairs	% of Errors
Europe + Oceania	37.50%	Europe + Oceania	23.74%
America + Asia (Southeast)	10.83%	America + Oceania	20.33%
Africa (Southern) + Africa (Sub)	10.83%	Africa (Southern) + Africa (Sub)	10.09%
America + Oceania	10.00%	Asia (South) + Asia (Southeast)	6.38%
Asia (South) + Asia (Southeast)	8.33%	America + Asia (southeast)	5.93%
Africa (Southern) + Oceania	5.83%	Africa (Southern) + America	5.34%
Asia (Southeast) + Oceania	5.00%	Africa (Southern) + Asia (Southeast)	5.19%

clusters in terms of that node’s ability to characterize dialectal variation. What this tells us, in short, is that the complete picture of dialectal variation cannot be seen by observing small portions of the grammar in isolation. And yet this is precisely what almost all work on syntactic variation does.

The next question is which regional dialects are similar according to the representations learned in this classifier. This is shown in Table 3 using the relative number of errors between regional dialects. Thus, for example, European English (UK and Ireland) and Oceanic English (Australia and New Zealand) account for a plurality of errors: 37.5% in the late-stage grammar and 23.74% in the early-stage grammar. Thus, these are the most similar dialects because they are the most likely to be confused. Given the colonial spread of English, this is of course not

surprising. What this table also shows is that the similarity between dialects differs to some degree depending on which part of the overall grammar we are observing: for example, North American English and Southeast Asian English are much more similar if we observe constructions from one part of the grammar (late-stage, with 10.83% of the errors) than the other (early-stage, with only 5.93% of the errors). This is important because it means that observing part of the grammar in isolation will not only inadequately characterize the amount of dialectal variation but will also provide different characterizations of the dialects relative to one another.

Before we investigate changes in dialect similarity across nodes within the grammar, we first evaluate the degree to which the classifiers depend on a small number of very predictive

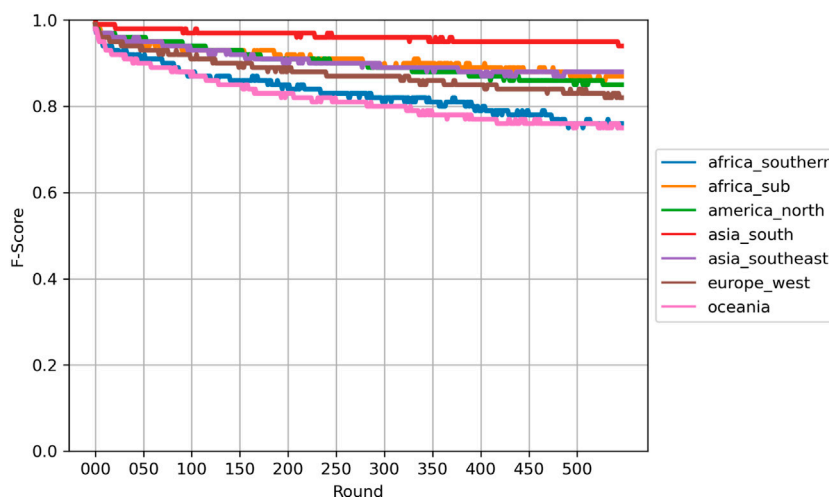


FIGURE 6
 Unmasking of Regional Dialects with the Late-Stage Grammar. Each 50 rounds removes approximately 2.5% of the grammar, so that round 500 includes only 75% of the original grammar.

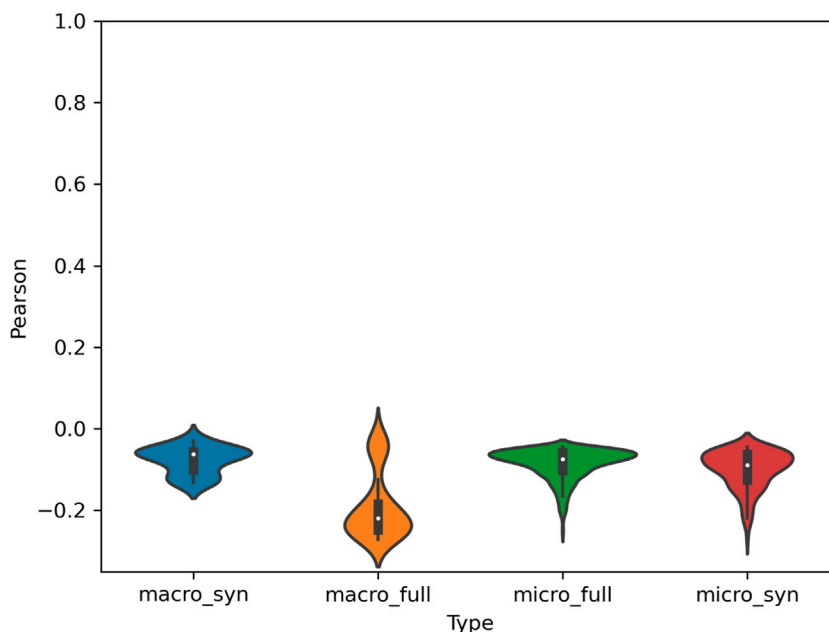


FIGURE 7
 Correlation of Error Distribution Between Late-Stage Grammar and Nodes within the Grammar for Regional Dialects. High Correlation indicates that the same dialects are similar in each model type while low correlation indicates that the relationships between dialects for a given component of the grammar differ from the late-stage grammar. Syn refers to the early-stage grammar and full to the late-stage grammar.

features. A possible confound with a prediction-based model of dialect is that a small number of constructions within each node could contribute almost all of the predictive power. This would distort our view of variation by implying that many constructions are in variation when, in fact, only a small number are. This is shown in Figure 6 for the late-stage grammar using the unmasking method from forensic linguistics (Koppel et al., 2007). The basic idea behind this

method is to train a linear classifier (in this case an SVM) over many rounds, removing the most predictive features for each dialect on each round. In the figure the f-score for each region is plotted on the y-axis over the course of 500 rounds of feature pruning, where each round removes the top features for each dialect. Overall, the prediction accuracy at round 500 represents the ability to characterize dialects when the top 25% of constructions have been removed. A robust classification

model exhibits a gentle downward curve on an unmasking plot while a brittle model (depending on just a few constructions) would exhibit a sharp decline in the first several rounds. This figure, then, shows that the regional dialect model is quite robust overall. This supports our further analysis by showing that the predictive power is not constrained to only a few constructions within each node.

Given that these dialect models do not implicitly depend on only a few constructions, we further examine the variation in similarity relationships across nodes within the grammar in Figure 7. This figure shows a violin plot of the distribution of correlation scores between i) the similarity relations derived from a node within the grammar and ii) the similarity relations derived from the high-performing late-stage grammar. The full late-stage grammar serves as a sort of ground-truth here because its prediction accuracies are nearly perfect. Thus, this measure can be seen as an indication of how close we would get to the actual relationships between dialects by observing only one node within the grammar (a macro-cluster or micro-cluster). The figure shows that, in all cases, there is no meaningful relationship between the two. If our goal is to characterize syntactic dialect variation as a whole, this means that studies which observe only isolated sets of features will not be able to

make predictions beyond those features. In short, language is a complex adaptive system and observing only small pieces of this system is inadequate for capturing emergent interactions between constructions.

5.2 National dialects

The previous section analyzed variation across regional dialects in order to show that variation is spread across many nodes within the grammar and that our view of syntactic variation depends heavily on the specific portions of the grammar under observation. This section continues this analysis with finer-grained country-level dialects in order to determine whether these same patterns emerge. The same methods of analysis are used, starting with the classification performance in Table 4. As before, this is divided into late-stage constructions and early-stage constructions. This table represents a single dialect model but for convenience it is divided into inner-circle varieties on the top and outer-circle varieties on the bottom (Kachru, 1990). As before, the late-stage grammar provides a better characterization of dialects than the early-stage grammar, 0.96 f-score vs. 0.83 f-score. While the late-stage grammar’s performance is comparable to the regional-dialect

TABLE 4 Performance of Dialect Classifier With National Dialects: Late-Stage Constructions (Left) and Early-Stage Constructions (Right). This is a single model for all dialects, with Inner-Circle Varieties shown at the top and Outer-Circle Varieties at the bottom. South African English does not fit neatly into the inner vs. outer circle distinction and would better be seen as partly belonging to both. This is due to the country’s internal demographic variation. The categorization here is for convenience in presenting the classification results and should not be taken as ignoring the intermediate status of this dialect.

Country	Late-stage grammar			Early-stage grammar		
	Precision	Recall	F-Score	Precision	Recall	F-Score
Inner-Circle Varieties						
Australia (au)	0.94	0.95	0.94	0.81	0.83	0.82
Canada (ca)	0.98	0.98	0.98	0.71	0.73	0.72
Ireland (ie)	0.94	0.96	0.95	0.88	0.88	0.88
New Zealand (nz)	0.89	0.84	0.86	0.64	0.59	0.61
United Kingdom (uk)	0.96	0.96	0.96	0.85	0.86	0.86
United States (us)	0.97	1.00	0.98	0.87	0.87	0.87
South Africa (za)	0.99	1.00	0.99	0.80	0.81	0.80
Outer-Circle Varieties						
Bangladesh (bd)	0.97	0.89	0.93	0.76	0.69	0.72
Indonesian (id)	1.00	0.84	0.91	0.57	0.49	0.53
India (in)	0.97	0.98	0.97	0.94	0.96	0.95
Kenya (ke)	0.98	0.99	0.98	0.92	0.87	0.89
Malaysia (my)	0.96	0.98	0.97	0.82	0.74	0.78
Nigeria (ng)	0.98	1.00	0.99	0.90	0.83	0.87
Philippines (ph)	0.98	0.98	0.98	0.83	0.87	0.85
Pakistan (pk)	0.94	0.94	0.94	0.84	0.89	0.87
Zimbabwe (zw)	0.95	0.91	0.93	0.48	0.53	0.51
Weighted Avg	0.96	0.96	0.96	0.83	0.83	0.83

TABLE 5 Distribution of errors in dialect classifier with national dialects: Late-stage constructions (left) and early-stage constructions (right).

Late-stage grammar		Early-stage grammar	
Country Pairs	% of Errors	Country Pairs	% of Errors
Australia + New Zealand	19.85%	Canada + United States	11.76%
Ireland + United Kingdom	12.11%	Australia + New Zealand	6.77%
India + Pakistan	11.86%	Bangladesh + Pakistan	5.58%
Bangladesh + Pakistan	7.47%	Australia + Canada	5.34%
Australia + United Kingdom	5.93%	Ireland + United Kingdom	5.29%
Canada + United States	4.12%	Australia + United Kingdom	4.69%

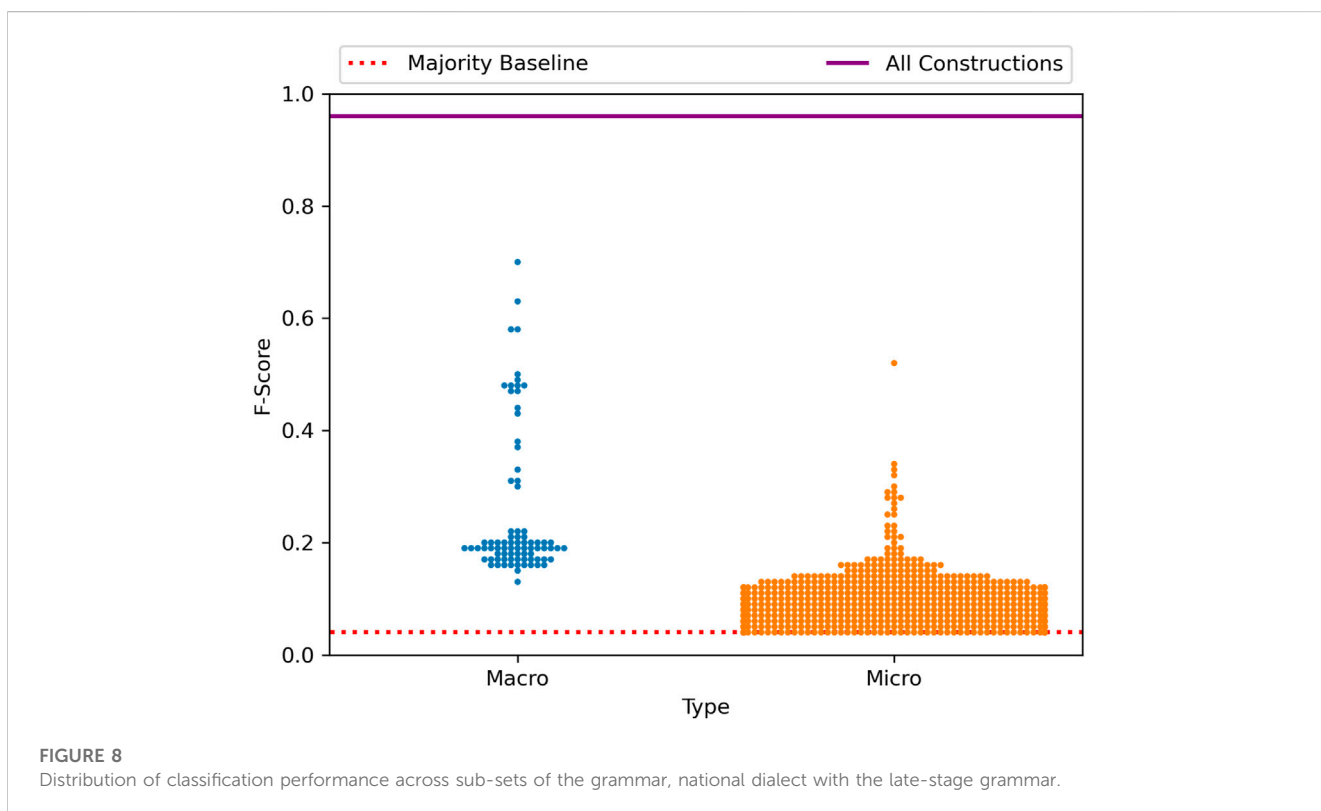


FIGURE 8 Distribution of classification performance across sub-sets of the grammar, national dialect with the late-stage grammar.

model, the early-stage model shows a sharp decline. Dialects with lower performance are less distinct and thus more influenced by other dialects; for example, we see that New Zealand English in both models is much less distinct or entrenched than other country-level dialects. It is thus more difficult to identify a sample of New Zealand English (c.f., [Dunn and Wong, 2022](#)).

Continuing with the error analysis in [Table 5](#), we see that almost 20% of the errors in the late-stage model are confusions between New Zealand and Australia and over 12% between Ireland and the UK (which includes Northern Ireland here). As before, the similarity between dialects derived from classification errors reflects the similarity between these countries in geographic, historical, and social terms.

The distribution of predictive accuracy across nodes within the grammar is shown in [Figure 8](#). Once again we see that a few portions

of the grammar have relatively high performance on their own (capturing upwards of 70% of the predictive power), but that no individual nodes perform as well as the grammar as a whole. As before this means that variation is spread throughout the grammar and that interactions between nodes is important for characterizing syntactic variation at the country level.

Continuing with the distribution of similarity values across sub-sets of the grammar, [Figure 9](#) shows the correlation between macro- and micro-clusters and the ground-truth of the high-performing late-stage grammar. Here the correlation is above chance but remains quite low (below 0.2). This again indicates that different nodes of the grammar provide different views of the similarity between country-level dialects, agreeing with the different errors ranks shown in [Table 5](#). For instance, New Zealand English might be close to Australian in one part of the grammar but close to UK

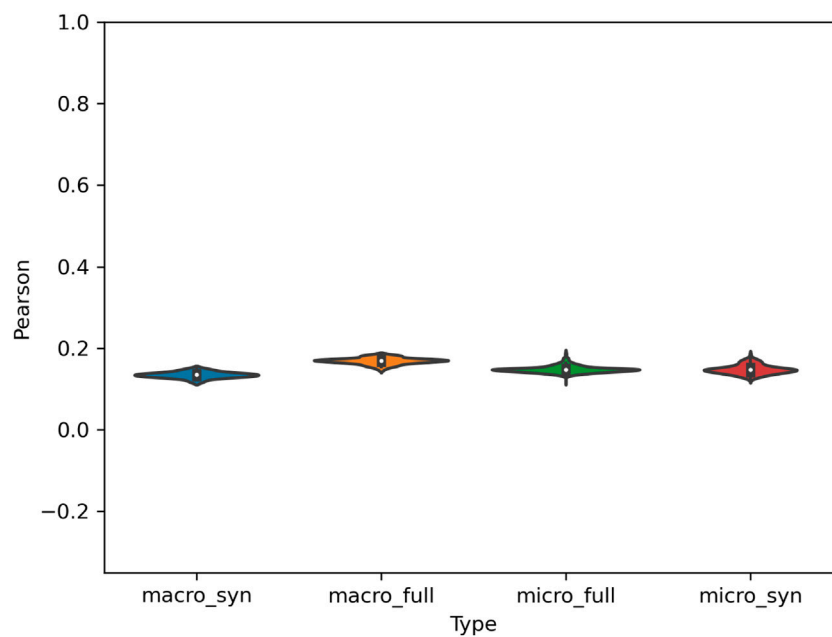


FIGURE 9

Correlation of Error Distribution Between Late-Stage Grammar and Nodes within the Grammar for National Dialects. High Correlation indicates that the same dialects are similar in each model type. Syn refers to the early-stage grammar and full to the late-stage grammar.

English in another. As a high-dimensional and complex system, grammatical variation must be viewed from the perspective of the entire grammar.

National dialects are more fine-grained than regional dialects and also have a higher number of categories (from 7 to 16), making the classification task more difficult because it must now distinguish between similar dialects (like American and Canadian English). Given the importance of the nation-state in modern mobility (i.e., in terms of ease of travel and immigration), these country-level dialects are more reflective of the underlying population than cross-country aggregations. In other words, the social network of countries is more coherent than that of larger regions because it is the nation which creates boundaries around human mobility. Since we are viewing both the grammar and the population of speakers of complex networks, it is important to go further and analyze local populations in the form of dialect areas within countries.

5.3 Local dialects

This section takes a closer look at dialectal variation by modelling the differences between local populations within the same region. As discussed above, data is collected from areas around airports as a consistent proxy for urban areas. These urban areas are then clustered into local groups using spatial but not linguistic information. Thus, as shown in the map of North American dialect areas in [Figure 1](#), a country-level dialect is divided into many smaller areas and then the differences between these local dialects are modelled using the same classification methods. This section examines the North American, European, and South Asian models in more

detail. The full results are available in the [Supplementary Material](#).⁵

5.3.1 North America

We start with North American dialects, mapped in [Figure 1](#) and listed by name in [Table 6](#), again with the late-stage grammar on the left and the early-stage grammar on the right. The distinction between same-country dialects is much smaller which means that the prediction task is much harder: here the f-scores drop to 0.69 and 0.40 (both of which remain much higher than the majority baseline). A classification model thus also provides a measure of the amount of dialectal variation between varieties: here there is less variation overall because the local populations are more similar. Local dialects with lower performance are again less distinguishable and thus less unique in their grammar: for example, Midwestern and Plains American English have the lowest f-scores at 0.57 and 0.52, respectively.

The distribution of classification errors is shown in [Table 7](#). All major sources of errors in the late-stage grammar are within the same country and close geographically. For example, over 10% of the errors are between Ontario and Quebec. The distribution of errors also shows that the lower performance of the Midwest and Plains areas is a result of their similarity to one another and, in the case of the Midwest, to the Eastern dialect. Thus, the types of errors here are as important as the number of errors.

The distribution of nodes within the grammar in terms of classification performance is shown in [Figure 10](#); here most micro-clusters have minimal predictive adequacy on their own but some

⁵ The DOI for this material is DOI 10.17605/OSF.IO/2AKMY.

TABLE 6 Performance of Dialect Classifier With Local Dialects: Late-Stage Constructions (Left) and Early-Stage Constructions (Right). This is a single model for all dialects, with American Varieties shown at the bottom and Canadian at the top.

Local Area	Late-stage grammar			Early-stage grammar		
	Precision	Recall	F-Score	Precision	Recall	F-Score
Western Canada (CA-1)	0.82	0.83	0.82	0.64	0.58	0.61
Ontario (CA-2)	0.65	0.69	0.67	0.43	0.45	0.44
Quebec (CA-3)	0.65	0.61	0.63	0.35	0.33	0.34
Nova Scotia (CA-4)	0.57	0.56	0.57	0.19	0.24	0.21
Midwest (US-1)	0.57	0.57	0.57	0.32	0.26	0.29
Central California (US-2)	0.80	0.81	0.81	0.35	0.34	0.35
Texas (US-3)	0.66	0.67	0.67	0.30	0.32	0.31
Southern (US-4)	0.68	0.56	0.61	0.26	0.27	0.27
Florida (US-5)	0.77	0.78	0.78	0.23	0.30	0.26
West Texas (US-6)	0.60	0.55	0.57	0.25	0.24	0.25
East Coast (US-7)	0.77	0.82	0.79	0.57	0.56	0.56
Plains (US-8)	0.55	0.49	0.52	0.11	0.14	0.12
Southwestern (US-9)	0.66	0.71	0.68	0.39	0.41	0.40
Weighted Avg	0.69	0.69	0.69	0.40	0.40	0.40

TABLE 7 Distribution of errors in dialect classifier with north American dialects: Late-stage constructions (left) and early-stage constructions (right).

Late-stage grammar		Early-stage grammar	
Country Pairs	% of Errors	Country Pairs	% of Errors
Ontario (CA-2) + Quebec (CA-3)	10.31%	Western (CA-1) + Quebec (CA-3)	6.09%
Western (CA-1) + Ontario (CA-2)	8.05%	Ontario (CA-2) + Quebec (CA-3)	5.12%
Midwest (US-1) + East (US-7)	7.73%	Ontario (CA-2) + East (US-7)	4.47%
Western (CA-1) + Quebec (CA-3)	6.44%	Western (CA-1) + Ontario (CA-2)	4.39%
Midwest (US-1) + Plains (US-8)	5.48%	Midwest (US-1) + Plains (US-8)	3.82%

macro-clusters retain meaningful predictive power. The correlation between similarity relations across nodes within the grammar is visualized in Figure 11. While the overall prediction accuracy is lower, the similarity relationships are more stable. In other words, the representation of which dialects have similar grammars is less dependent here on the sub-set of the grammar being observed.

5.3.2 Europe

The map of European dialects (within the UK and Ireland) is shown in Figure 12. As before, these areas are formed using spatial clustering with the dialect classifier then used to characterize the syntactic differences between dialect areas. As listed in Table 8, there are six local dialects with an f-score of 0.77 (late-stage grammar) and 0.64 (early-stage grammar). The range of f-scores across dialect areas is wider than in the North American model, with a high of 0.94 (Ireland) and a low of 0.40 (Scotland). This is partly driven by the number of samples per dialect, with Scotland particularly under-

represented. This means, then, that the characterizations this model makes about Irish English are much more reliable than the characterizations made about Scottish English.

The distribution of performance across nodes in the grammar is shown in Figure 13. The figure shows that many nodes are meaningfully above the majority baseline in terms of predictive power but all fall short of the grammar as a whole. In fact, some micro-clusters have nearly as much power as the best macro-clusters. This indicates that the variation in UK and Irish English is spread throughout many individual parts of the grammar, a fact that would be overlooked if we had focused instead on a few constructions in isolation.

The main take-away from the European model, then, is once again that an accurate characterization of the grammatical differences between these local dialects requires access to the entire grammar. Focusing on smaller nodes within the grammar does not provide strong predictive power. As before, this has two

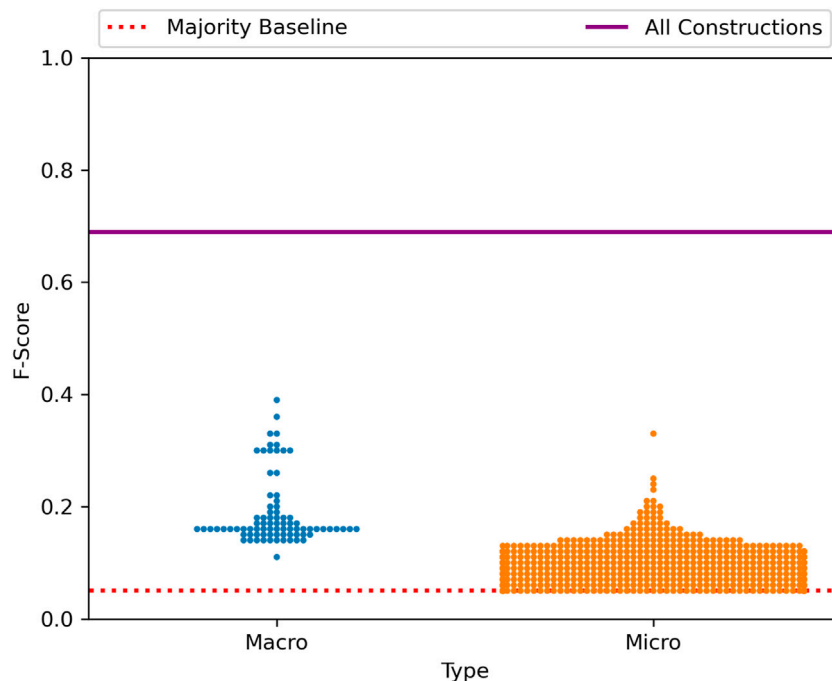


FIGURE 10
Distribution of Performance across the Late-Stage grammar, North America.

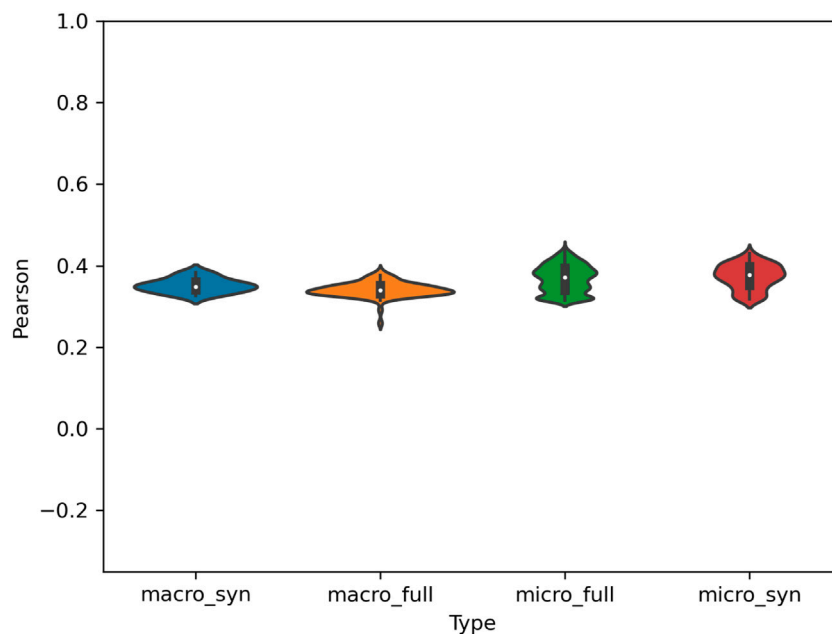


FIGURE 11
Correlation of Error Distribution Between Late-Stage Grammar and Nodes within the Grammar for North American Dialects. High Correlation indicates that the same dialects are similar in each model. Syn refers to the early-stage grammar and full to the late-stage grammar.

implications: first, that variation is distributed across the grammar and, second, that emergent relationships between constructions are essential for depicting syntactic variation in the aggregate.

5.3.3 South Asia

The final model of local dialects we will investigate is the outer-circle varieties from South Asia. These differ from the

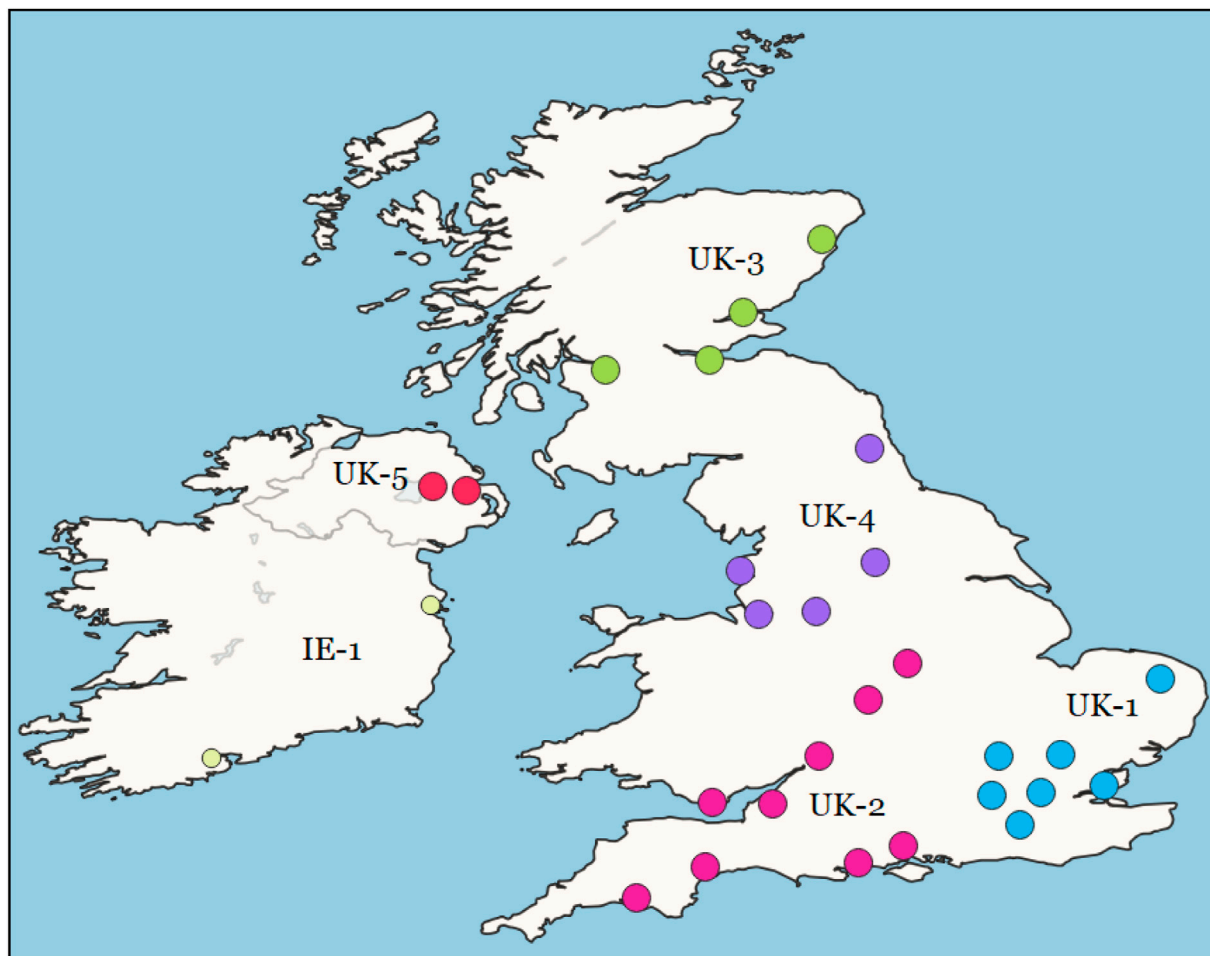


FIGURE 12
Distribution of Local Dialects in Europe. Each point is an airport; thus, for example, IE has two airport locations but only one local dialect area.

TABLE 8 Performance of Dialect Classifier With Local Dialects: Late-Stage Constructions (Left) and Early-Stage Constructions (Right). This is a single model for all local dialects.

Local Area	Late-stage grammar			Early-stage grammar		
	Precision	Recall	F-Score	Precision	Recall	F-Score
Ireland (IE-1)	0.92	0.97	0.94	0.84	0.88	0.86
Southeast England (UK-1)	0.70	0.74	0.72	0.59	0.61	0.60
Southwest England (UK-2)	0.60	0.58	0.59	0.43	0.43	0.43
Scotland (UK-3)	0.54	0.32	0.40	0.23	0.23	0.23
Northern England (UK-4)	0.75	0.71	0.73	0.59	0.57	0.58
Northern Ireland (UK-5)	0.75	0.46	0.57	0.38	0.23	0.29
Weighted Avg	0.77	0.78	0.77	0.64	0.65	0.64

inner-circle dialects in a number of ways. In the first case, these speakers use English for different purposes than inner-circle speakers, almost always being highly multilingual with other languages used in the home (while, for example, North

American speakers of English are often monolingual). In the second case, there is a socio-economic skew in the sense that higher class speakers are more likely to use English. The impact of socio-economic status is even more important when we consider

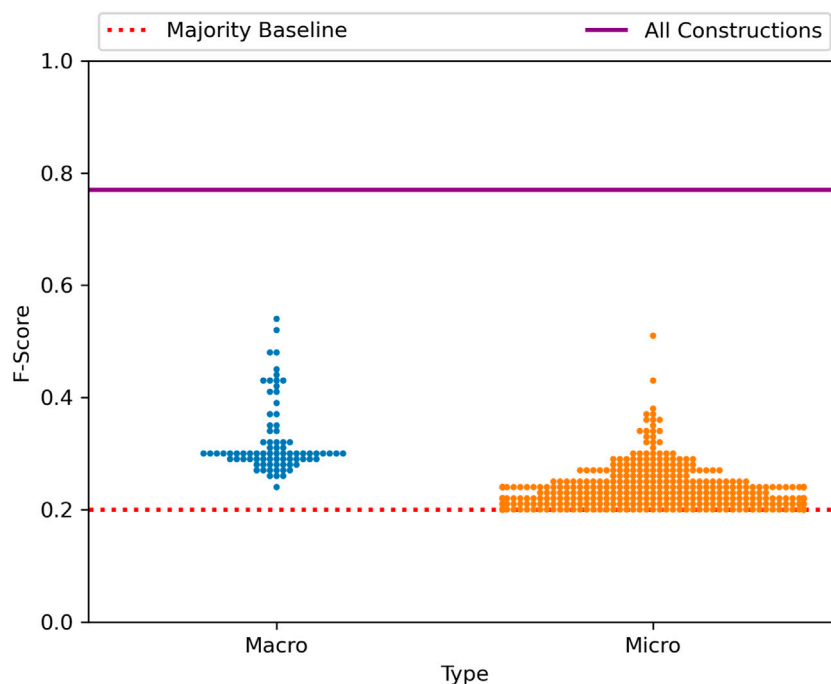


FIGURE 13
Distribution of Performance across the Late-Stage grammar, Europe.

that this data is collected from digital sources (tweets) and that the population of Twitter users in South Asia is less representative of the larger population than it is in North America and Europe. Thus, we use South Asia as a case-study in variation within outer-circle dialects. The map of local dialect areas is shown in Figure 14. This encompasses three countries: India, Pakistan, and Bangladesh.

The performance of the dialect model by local dialect is shown in Table 9, divided by country and by grammar type. The overall f-scores here are more similar to the North American model, at 0.68 and 0.56. As before, it is more difficult to distinguish between these local dialects because they are more similar overall. There is also a larger range of f-scores than before: the Bangladesh dialects are the highest performing, at 0.97 and 0.82. But within India, with many adjacent dialect areas, the f-scores fall as low as 0.14 (Uttar Pradesh) and 0.24 (Bihar). One conclusion to be drawn from these low values for specific local dialects is that the areas posited by the spatial clustering step do not contain unique and predictable dialectal variants. This is a clear case where a joint spatial/linguistic approach to forming local areas would be preferable. In other words, the spatial organization suggests a boundary which the linguistic features do not support.

The distribution of predictive power across the grammar is shown in Figure 15. As before, predictive power is distributed across nodes within the grammar and no single node nears the performance of the grammar as a whole. This means that interactions between constructions are an important part of variation and that the variability in different nodes is not simply redundant information.

6 Discussion and conclusion

The main focus of this paper has been to investigate the assumption, common to most studies of grammatical variation across dialects, that individual constructions within the grammar can be viewed in isolation. The results show conclusively that this is simply not the case: language is a complex adaptive system and observing variation in an arbitrary sub-set of that system leads to incomplete models. The paper has repeated the same experimental paradigm at three levels of spatial granularity (regional, national, and local dialects) in order to show that the inadequacy of individual nodes within the grammar is a consistent phenomenon. Across all levels of spatial granularity, these experiments tell us three important things about variation within a complex system:

Finding 1: No individual node within the grammar captures dialectal variation as accurately as the grammar as a whole. The basic approach taken here is to first cluster constructions within the grammar into macro- and micro-clusters using similarity relations between the constructions themselves (thus, independently of their patterns of variation). As shown in Figure 2, this leads to 1,941 micro-clusters in the late-stage grammar and 191 in the early-stage grammar. Even larger nodes are captured by macro-clusters, including 81 in the late-stage grammar and 16 in the early-stage grammar. The dialect classification experiments are repeated with each of these clusters alone, with the results shown in Figures 4, 5, 8, 10, 13, 15. Additional figures are available in the Supplementary Material. In each case the grammar as a whole performs much better than any individual sub-set or node within the grammar.

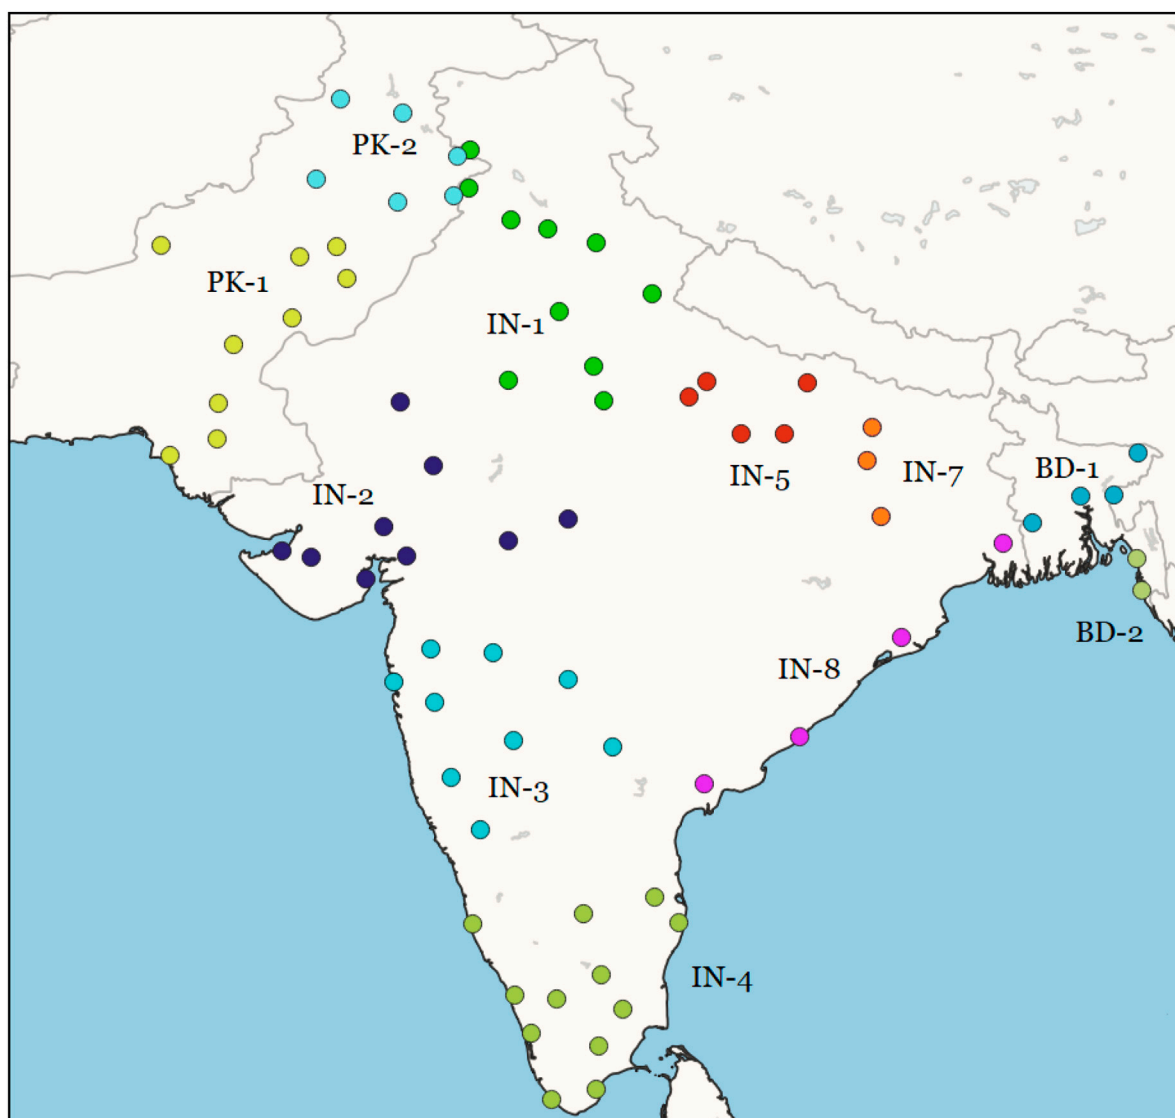


FIGURE 14
Distribution of local dialects in South Asia.

Why? Our interpretation is that there are emergent interactions between constructions in different nodes across the grammar. This means, for instance, that the use of plural nouns interacts with the use of certain adpositional phrases which interacts in turn with the use of phrasal verbs. Language is a complex adaptive system and the huge number of such minor interactions provides information about syntactic variation which is not redundant with the information contained in local nodes within the grammar. In other words, a significant part of syntactic variation is contained within emergent interactions which cannot be localized to specific families of constructions.

Finding 2: Individual nodes within the grammar vary widely in the degree to which they are subject to dialectal variation. Following from the above finding, the degree to which individual macro- and micro-clusters are able to predict dialect membership represents the degree to which they themselves are subject to spatial variation. Thus, the results shown in [Figures 4, 5, 8, 10, 13, 15](#) also

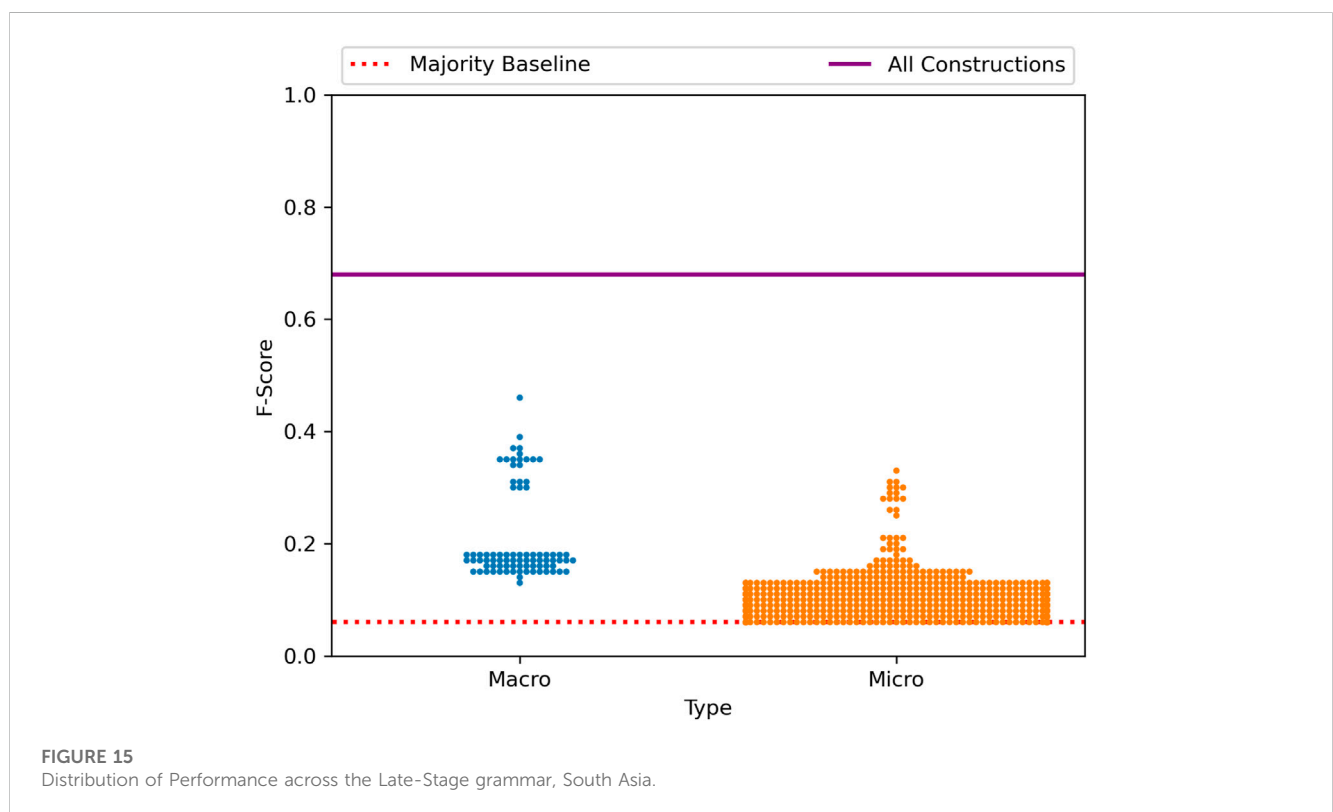
mean that not all clusters of constructions are subject to the same level of variation.

Why? Our interpretation is that dialect variation is distributed across the grammar but in uneven portions. For instance, the unmasking experiment ([Figure 6](#) and the [Supplementary Material](#)) shows that, even if we disregard the network structure of the grammar, the performance of the model is distributed across many features: performance remains rather high even with the top 25% of constructions removed. Thus, we know that variation is distributed widely across the grammar (regardless of cluster assignments) and that macro- and micro-clusters also vary widely in predictive power. This is important because it means that a grammar's variation is not simply the sum of the variation of its component constructions.

Finding 3: Similarity relations between dialects diverge widely from the best-performing model depending on the subset of the grammar being observed. Given that the entire late-stage

TABLE 9 Performance of Dialect Classifier With Local Dialects in South Asia: Late-Stage Constructions (Left) and Early-Stage Constructions (Right). This is a single model for all local dialects.

Local Area	Late-stage grammar			Early-stage grammar		
	Precision	Recall	F-Score	Precision	Recall	F-Score
Dhaka (BD-1)	0.96	0.98	0.97	0.88	0.89	0.89
Chattogram (BD-2)	0.87	0.79	0.82	0.61	0.58	0.60
New Delhi (IN-1)	0.56	0.64	0.60	0.51	0.53	0.52
Gujarat (IN-2)	0.55	0.55	0.55	0.33	0.39	0.36
Maharashtra (IN-3)	0.67	0.71	0.69	0.54	0.53	0.54
Tamil Nadu (IN-4)	0.72	0.71	0.71	0.55	0.58	0.56
Uttar Pradesh (IN-5)	0.17	0.11	0.14	0.42	0.22	0.29
Bihar (IN-7)	0.39	0.17	0.24	0.28	0.17	0.21
Andhra Pradesh + Odisha (IN-8)	0.69	0.61	0.65	0.54	0.49	0.51
Southern (PK-1)	0.56	0.51	0.53	0.43	0.40	0.42
Northern (PK-2)	0.72	0.75	0.74	0.63	0.66	0.65
Weighted Avg	0.68	0.68	0.68	0.56	0.56	0.56



grammar retains high classification performance, we extract similarity relationships between dialects by looking at the errors which the classifier makes. For example, our reasoning is that Midwestern American English and Plains American English have more errors precisely because their grammar is more similar and thus easier to confuse. But the question is whether all sub-sets of the

grammar lead to the same similarity errors. The answer is a resounding no, as shown in Tables 4 and 6, and 7 and in Figures 7, 9, 11. Taking similarity measures as a characterization of the dialects relative to one another, this means that the overall story of variation depends heavily on which sub-set of the grammar we observe.

Why is this a fundamental problem for previous work? Our interpretation is that language is a complex adaptive system so that examining any sub-set in isolation leads to incorrect conclusions. Thus, if we observed a small part of the grammar and tried to predict the most similar grammars overall, we would rarely reach the same relationships. By treating parts of the grammar as independent and self-contained systems, we fail to capture the interactions and complexities of the grammar as a single functioning system. The implication of these findings are that all studies which are based on isolating one or two constructions are fundamentally unable to provide an accurate representation of syntactic variation.

Data availability statement

While the corpora analyzed for this study cannot be shared directly, the constructional features used for the classification models can be found <https://osf.io/2akmy/>. This allows the replication of these findings. The full supplementary materials contain the raw classification results, the full grammar and composition of macro- and micro-clusters, and additional figures and maps for each level of spatial granularity. These additional results repeat the main findings of the analysis presented in the paper. The supplementary material can be found <https://osf.io/2akmy/>.

Author contributions

JD: Writing—original draft, Writing—review and editing.

References

- Barbaresi, A. (2018). Computationally efficient discrimination between language varieties with large feature vectors and regularized classifiers. *Proc. Fifth Workshop NLP Similar Lang. Var. Dialects*, 164–171. <https://aclanthology.org/W18-3918>.
- Beckner, C., Blythe, R., Bybee, J., Christiansen, M. H., Croft, W., Ellis, N. C., et al. (2009). Language is a complex adaptive system: position paper. *Lang. Learn.* 59, 1–26. doi:10.1111/j.1467-9922.2009.00533.x
- Belinkov, Y., and Glass, J. (2016). “A character-level convolutional neural network for distinguishing similar languages and dialects,” in *Proceedings of the third workshop on NLP for similar languages, varieties and dialects* (United States: Association for Computational Linguistics), 145–152.
- Buchstaller, I. (2008). The localization of global linguistic variants. *Engl. World-Wide* 29, 15–44. doi:10.1075/eww.29.1.03buc
- Bybee, J. (2007). *Frequency of use and the organization of language*. Oxford: Oxford University Press.
- Calle-Martin, J., and Romero-Barranco, J. (2017). Third person present tense markers in some varieties of English. *Engl. World-Wide* 38, 77–103. doi:10.1075/eww.38.1.05cal
- Campello, R. J. G. B., Moulavi, D., and Sander, J. (2013). “Density-based clustering based on hierarchical density estimates,” in *Advances in knowledge discovery and data mining* (Germany: Springer Berlin Heidelberg), 160–172. doi:10.1007/978-3-642-37456-2/_/14
- Campello, R., Moulavi, D., Zimek, A., and Sander, J. (2015). Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Trans. Knowl. Discov. Data* 10, 1–51. doi:10.1145/2733381
- Deshors, S. C., and Götz, S. (2020). Common ground across globalized English varieties: A multivariate exploration of mental predicates in world englishes. *Corpus Linguistic Theory* 16, 1–28. doi:10.1515/clt-2016-0052
- Donoso, G., Sánchez, D., and Sanchez, D. (2017). “Dialectometric analysis of language variation in Twitter,” in *Proceedings of the fourth workshop on NLP for similar languages, varieties and dialects (VarDial)* (Valencia, Spain: Association for Computational Linguistics), 4, 16–25. doi:10.18653/v1/W17-1202
- Doumen, J., Beuls, K., and Van Eecke, P. (2023). “Modelling language acquisition through syntactico-semantic pattern finding,” in Findings of the Association for computational linguistics: EACL 2023 (*dubrovnik, Croatia* (China: Association for Computational Linguistics), 1347–1357.
- Dunn, J. (2023a). *Computational construction grammar: A usage-based approach*. available at <https://www.jdunn.name/cxg>.
- Dunn, J. (2017). Computational learning of construction grammars. *Lang. Cognition* 9, 254–292. doi:10.1017/langcog.2016.7
- Dunn, J. (2023b). “Exploring the construction: linguistic analysis of a computational CxG,” in *Proceedings of the first international workshop on construction grammars and NLP (CxGs+NLP, GURT/SyntaxFest 2023)* (Washington, D.C.: Association for Computational Linguistics), 1–11.
- Dunn, J. (2022). *Exposure and emergence in usage-based grammar: Computational experiments in 35 languages*. Cognitive Linguistics.
- Dunn, J. (2018a). Finding variants for construction-based dialectometry: A corpus-based approach to regional cxgs. *Cogn. Linguist.* 29, 275–311. doi:10.1515/cog-2017-0029
- Dunn, J. (2019a). “Frequency vs. Association for constraint selection in usage-based construction grammar,” in *Proceedings of the workshop on cognitive modeling and computational linguistics* (United States: Association for Computational Linguistics).
- Dunn, J. (2019b). Global syntactic variation in seven languages: toward a computational dialectology. *Front. Artif. Intell.* 2, 15. doi:10.3389/frai.2019.00015
- Dunn, J. (2020). Mapping languages: the corpus of global language use. *Lang. Resour. Eval.* 54, 999–1018. doi:10.1007/s10579-020-09489-2
- Dunn, J. (2019c). “Modeling global syntactic variation in English using dialect classification,” in *Proceedings of the sixth workshop on NLP for similar languages, varieties and dialects* (Ann Arbor, Michigan: Association for Computational Linguistics), 42–53. doi:10.18653/v1/W19-1405

Funding

The author declares that no financial support was received for the research and authorship of this article. The publication was supported by the University of Canterbury.

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author JD declared that they were an editorial board member of *Frontiers* at the time of submission. This had no impact on the peer review process and the final decision.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fcpxs.2023.1273741/full#supplementary-material>

- Dunn, J. (2018b). Modeling the complexity and descriptive adequacy of construction grammars. *Proc. Soc. Comput. Linguistics*, 81–90. doi:10.7275/R59P2ZTB
- Dunn, J., and Nijhof, W. (2022). “Language identification for austronesian languages,” in *Proceedings of the 13th international conference on language resources and evaluation* (European: European Language Resources Association), 6530–6539.
- Dunn, J., and Nini, A. (2021). “Production vs perception: the role of individuality in usage-based grammar induction,” in *Proceedings of the workshop on cognitive modeling and computational linguistics* (USA: Association for Computational Linguistics), 149–159.
- Dunn, J., and Tayyar Madabushi, H. (2021). “Learned construction grammars converge across registers given increased exposure,” in *Conference on natural language learning* (USA: Association for Computational Linguistics).
- Dunn, J. (2023c). “Variation and instability in dialect-based embedding spaces,” in Tenth Workshop on NLP for similar languages, Varieties and dialects (VarDial 2023) (dubrovnik, Croatia (Germany: Association for Computational Linguistics), 67–77.
- Dunn, J., and Wong, S. (2022). “Stability of syntactic dialect classification over space and time,” in *Proceedings of the 29th international conference on computational linguistics (gyeongju, Republic of Korea* (Korea: International Committee on Computational Linguistics), 26–36.
- Eisenstein, J., O’Connor, B., Smith, N., and Xing, E. (2010). A latent variable model for geographic lexical variation, Proceedings of the Conference on Empirical Methods in Natural Language Processing, Abu Dhabi, December 2010, 287. Association for Computational Linguistics, 221–227.
- Eisenstein, J., O’Connor, B., Smith, N., and Xing, E. (2014). Diffusion of lexical change in social media. *PLoSOne* 10, 1371. <https://aclanthology.org/D10-1124>.
- Gamallo, P., Pichel, J., Algeria, I., and Agirrezabal, M. (2016). Comparing two basic methods for discriminating between similar languages and varieties. Proceedings of the third workshop on NLP for similar languages. *Var. Dialects*, 170–177. <https://aclanthology.org/W16-4822>.
- Goldberg, A. (2006). *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.
- Gonçalves, B., Loureiro-Porto, L., Ramasco, J. J., and Sánchez, D. (2018). Mapping the americanization of English in space and time. *PLOS ONE* 13, 01977411–e197815. doi:10.1371/journal.pone.0197741
- Gonçalves, B., and Sánchez, D. (2014). Crowdsourcing dialect characterization through twitter. *PLOS ONE* 9, e112074–e112076. doi:10.1371/journal.pone.0112074
- Grafmiller, J., and Szmrecsanyi, B. (2018). Mapping out particle placement in Englishes around the world A study in comparative sociolinguistic analysis. *Lang. Var. Change* 30, 385–412. doi:10.1017/s0954394518000170
- Grieve, J. (2012). A statistical analysis of regional variation in adverb position in a corpus of written Standard American English. *Corpus Linguistics Linguistic Theory* 8, 39–72. doi:10.1515/cllt-2012-0003
- Grieve, J., Montgomery, C., Nini, A., Murakami, A., and Guo, D. (2019). Mapping lexical dialect variation in British English using twitter. *Front. Artif. Intell.* 2, 11. doi:10.3389/frai.2019.00011
- Grieve, J. (2016). *Regional variation in written American English*. Cambridge, UK: Cambridge University Press.
- Kachru, B. (1990). *The Alchemy of English the spread, functions, and models of non-native Englishes*. Urbana-Champaign: University of Illinois Press.
- Koppel, M., Schler, J., and Bonchek-Dokow, E. (2007). Measuring differentiability: unmasking pseudonymous authors. *J. Mach. Learn. Res.* 8, 1261–1276. <http://jmlr.org/papers/v8/koppel07a.html>.
- Kroon, M., Medvedeva, M., and Plank, B. (2018). When simple n-gram models outperform syntactic approaches discriminating between Dutch and Flemish. *Proc. Fifth Workshop NLP Similar Lang. Var. Dialects*, 225–244. <https://aclanthology.org/W18-3928/>.
- Laitinen, M., Fatemi, M., and Lundberg, J. (2020). Size matters: digital social networks and language change. *Front. Artif. Intell.* 3, 46–15. doi:10.3389/frai.2020.00046
- Laitinen, M., and Fatemi, M. (2022). Social and regional variation in world englishes: local and global perspectives (routledge), chap. *Big and rich social networks in computational sociolinguistics*, 1–25.
- Langacker, R. (2008). *Cognitive grammar: A basic introduction*. Oxford: Oxford University Press.
- Li, Y., Szmrecsanyi, B., and Zhang, W. (2023). The theme-recipient alternation in Chinese: tracking syntactic variation across seven centuries. *Corpus Linguistics Linguistic Theory* 19, 207–235. doi:10.1515/cllt-2021-0048
- Lucy, L., and Bamman, D. (2021). Characterizing English variation across social media communities with bert. *Trans. Assoc. Comput. Linguistics* 9, 538–556. doi:10.1162/tacl_a_00383
- Marttinen Larsson, M. (2023). Modelling incipient probabilistic grammar change in real time: the grammaticalisation of possessive pronouns in european Spanish locative adverbial constructions. *Corpus Linguistics Linguistic Theory* 19, 177–206. doi:10.1515/cllt-2021-0030
- McInnes, L., and Healy, J. (2017). “Accelerated hierarchical density clustering,” in *IEEE international conference on data mining workshops (ICDMW)* (Germany: IEEE). doi:10.1109/ICDMW.2017.12
- Mocanu, D., Baronchelli, A., Perra, N., Gonçalves, B., Zhang, Q., Vespignani, A., et al. (2013). The twitter of babel: mapping world languages through microblogging platforms. *PLOSOne* 10, 1371. doi:10.1371/journal.pone.0061981
- Nevens, J., Doumen, J., Van Eecke, P., and Beuls, K. (2022). Language acquisition through intention reading and pattern finding, Proceedings of the 29th International Conference on Computational Linguistics (Gyeongju, Republic of Korea, October 2022, Korea, 15–25. International Committee on Computational Linguistics.
- Rahimi, A., Cohn, T., and Baldwin, T. (2017). A neural model for user geolocation and lexical dialectology. *Proc. 55th Annu. Meet. Assoc. Comput. Linguistics (Volume 2 Short Pap. 2)*, 209–216. doi:10.18653/v1/P17-2033
- Rautionaho, P., and Hundt, M. (2022). Primed progressives? Predicting aspectual choice in world englishes. *Corpus Linguistics Linguistic Theory* 18, 599–625. doi:10.1515/cllt-2021-0012
- Schilk, M., and Schaub, S. (2016). Noun phrase complexity across varieties of English: focus on syntactic function and text type. *Engl. World-Wide* 37, 58–85. doi:10.1075/eww.37.1.03sch
- Schneider, G., Hundt, M., and Schreier, D. (2020). Pluralized non-count nouns across englishes: A corpus-linguistic approach to variety types. *Corpus Linguistics Linguistic Theory* 16, 515–546. doi:10.1515/cllt-2018-0068
- Szmrecsanyi, B., Grafmiller, J., and Rosseel, L. (2019). Variation-based distance and similarity modeling: A case study in world englishes. *Front. Artif. Intell.* 2, 23. doi:10.3389/frai.2019.00023
- Xu, H., Jiang, M., Lin, J., and Huang, C.-R. (2022). Light verb variations and varieties of Mandarin Chinese: comparable corpus driven approaches to grammatical variations. *Corpus Linguistics Linguistic Theory* 18, 145–173. doi:10.1515/cllt-2019-0049
- Zampieri, M., Nakov, P., and Scherrer, Y. (2020). Natural language processing for similar languages, varieties, and dialects: A survey. *Nat. Lang. Eng.* 26, 595–612. doi:10.1017/S1351324920000492

Appendix

TABLE A1 Keywords used to create lexically-balanced samples. Words are listed in order of frequency.

know	time	people	day	love	new	see	think
why	here	want	go	really	need	today	make
still	because	first	very	best	after	than	never
got	much	back	please	going	great	right	then
life	thank	well	way	always	year	over	world
most	take	man	say	last	let	into	work
where	other	look	many	said	off	same	years
which	game	video	better	come	something	happy	thanks
via	yes	down	hope	god	stop	give	ever
feel	everyone	big	team	help	live	getting	while
use	keep	things	another	long	week	sure	days
watch	real	looking	shit	against	actually	doing	money
free	show	since	home	lot	nothing	bad	find
already	read	through	part	tell	without	won	such
start	little	play	thought	everything	morning	old	support
person	call	done	check	mean	news	put	both
wait	women	end	believe	used	top	around	night
looks	family	name	country	yeah	anyone	between	gonna
trying	says	hard	guys	maybe	friends	point	beautiful
remember	win	full	sorry	follow	government	high	during
amazing	yet	making	school	under	anything	coming	state
post	away	guy	change	try	house	open	might
season	whole	makes	left	song	media	saying	few
using	president	different	enough	black	called	talk	trump
place	talking	friend	care	power	once	wrong	city
working	nice	ready	times	business	understand	set	music
join	buy	vote	hate	heart	future	girl	mind
wish	face	seen	tomorrow	found	needs	watching	party
though	playing	least	problem	stay	covid	project	head
kind	white	group	health	until	food	story	cause
literally	soon	men	congratulations	job	ask	police	human
saw	waiting	far	—	—	—	—	—