

Modeling Global Syntactic Variation in English Using Dialect Classification

Jonathan Dunn

University of Canterbury

Department of Linguistics

jonathan.dunn@canterbury.ac.nz

Abstract

This paper evaluates global-scale dialect identification for 14 national varieties of English as a means for studying syntactic variation. The paper makes three main contributions: (i) introducing data-driven language mapping as a method for selecting the inventory of national varieties to include in the task; (ii) producing a large and dynamic set of syntactic features using grammar induction rather than focusing on a few hand-selected features such as function words; and (iii) comparing models across both web corpora and social media corpora in order to measure the robustness of syntactic variation across registers.

1 Syntactic Variation Around the World

This paper combines grammar induction (Dunn, 2018a, 2018b, 2019) and text classification (Joachims, 1998) to model syntactic variation across national varieties of English. This classification-based approach is situated within the task of dialect identification (Section 2) and evaluated against other baselines for the task (Sections 7 and 8). But the focus is modelling syntactic variation on a global-scale using corpus data. On the one hand, the problem is to use a model of syntactic preferences to predict an author’s dialect membership (Dunn, 2018c). On the other hand, the problem is to take a spatially-generic grammar of English that is itself learned from raw text (c.f., Zeman, et al., 2017; Zeman, et al., 2018) and adapt that grammar using dialect identification as an optimization task: which constructions are more likely to occur in a specific regional variety?

Because we want a complete global-scale model, we first have to ask: how many national varieties of English are there? This question, considered in Sections 3 and 4, is essential for determining the inventory of regional varieties that need to be included in the dialect identification task. This

paper uses data-driven language mapping to find out where English is consistently used, given web data and Twitter data, in order to avoid the arbitrary selection of dialect areas. This is important for ensuring that each construction in the grammar receives the best regional weighting.

What syntactic features are needed to represent variation in English? As discussed in Section 6, this paper uses grammar induction on a large background corpus to provide a replicable and dynamic feature space in order to avoid arbitrary limitations (e.g., lists of function words). The other side of this problem is to optimize grammar induction for regional dialects by using an identification task to learn regional weights for each part of the grammar: how much does a single generic grammar of English vary across dialects? To what degree does it represent a single dominant dialect?

Finally, a corpus-based approach to variation is restricted to the specific domains or registers that are present in the corpus. To what degree is such a model of variation limited to a specific register? This paper uses both web-crawled corpora and social media corpora to explore the robustness of dialect models across domains (Section 8). Along these same lines, how robust is a model of syntactic variation to the presence of a few highly predictive features? This paper uses unmasking, a method from authorship verification (Koppel, et al., 2007), to evaluate the stability of dialect models over rounds of feature pruning (Section 9).

2 Previous Work

Because of its long history as a colonial language (Kachru, 1990), English is now used around the world by diverse national communities. In spite of the global character of English, dialectology and sociolinguistics continue to focus largely on sub-national dialects of En-

glish within so-called *inner-circle* varieties (for example, Labov, et al., 2016; Strelluf, 2016; Schreier, 2016; Clark & Watson, 2016). This paper joins recent work in taking a global approach by using geo-referenced texts to represent national varieties (e.g., Dunn, 2018c; Tamaredo, 2018; Calle-Martin & Romero-Barranco, 2017; Szmrecsanyi, et al., 2016; Sanders, 2010, 2007; c.f., Davies & Fuchs, 2015). For example, this study of dialect classification contains *inner-circle* (Australia, Canada, United Kingdom, Ireland, New Zealand, United States), *outer-circle* (India, Malaysia, Nigeria, Philippines, Pakistan, South Africa), and *expanding-circle* (Switzerland, Portugal) varieties together in a single model.

The problem is that these more recent approaches, while they consider more varieties of English, have arbitrarily limited the scope of variation by focusing on a relatively small number of features (Grafmiller & Szmrecsanyi, 2018; Kruger & van Rooy, 2018; Schilk & Schaub, 2016; Collins, 2012). In practical terms, such work uses a smaller range of syntactic representations than comparable work in authorship analysis (c.f., Grieve, 2007; Hirst & Feiguina, 2007; Argamon & Koppel, 2013).

From a different perspective, we could view the modelling of dialectal variation as a classification task with the goal of predicting which dialect a sample belongs to. Previous work has drawn on many representations that either directly or indirectly capture syntactic patterns (Gamallo, et al., 2016; Barbaresi, 2018; Kreutz & Daelemans, 2018; Kroon, et al., 2018). Given a search for the highest-performing approach, other work has shown that methods and features without a direct linguistic explanation can still achieve impressive accuracies (McNamee, 2016; Ionescu & Popescu, 2016; Belinkov & Glass, 2016; Ali, 2018).

On the other hand, there is a conceptual clash between potentially topic-based methods for dialect identification and other tasks that explicitly model place-specific language use. For example, text-based geo-location can use place-based topics to identify where a document is from (c.f., Wing & Baldrige, 2014; Hulden, et al., 2015; Lourentzou, et al., 2017). And, at the same time, place-based topics can be used for both characterizing the functions of a location (c.f., Adams & McKenzie, 2018; Adams, 2015) and disambiguating gazeteers (c.f., Ju, et al., 2016). This raises an

Region	CC	TW
Africa, North	123,859,000	85,552,000
Africa, Southern	59,075,000	87,348,000
Africa, Sub	424,753,000	254,200,000
America, Brazil	218,119,000	118,138,000
America, Cen.	886,610,000	383,812,000
America, North	236,590,000	350,125,000
America, South	1,163,008,000	402,150,000
Asia, Cen.	965,090,000	102,794,000
Asia, East	2,201,863,000	95,704,000
Asia, South	448,237,000	331,192,000
Asia, Southeast	2,011,067,000	245,181,000
Europe, East	4,553,101,000	322,460,000
Europe, Russia	101,444,000	105,045,000
Europe, West	2,422,855,000	823,807,000
Middle East	660,732,000	222,985,000
Oceania	164,025,000	213,064,000
TOTAL	16.65 billion	4.14 billion

Table 1: Background Corpus Size in Words by Region

important conceptual problem: when does predictive accuracy reflect *dialects* as opposed to either place-references or place-based content? While geo-referenced corpora capture both types of information, syntactic representations focus specifically on *linguistic* variation while place-references and place-based topics are part of document content rather than linguistic structure.

3 Where Is English Used?

The goal of this paper is to model syntactic variation across all major or robust varieties of English. But how do we know which varieties should be included? Rather than select some set of varieties based on convenience, we take a data-driven approach by collecting global web-crawled data and social media data to determine where English is used. This approach is biased towards developed countries with access to digital technologies. As shown in Table 1, however, enough global language data is available from both sources to determine where national varieties of English exist.

Data comes from two sources of digital texts: web pages from the Common Crawl¹ and social media from Twitter.² Both types of data have been used previously to study dialectal and spatial variation in language. More commonly, geo-referenced Twitter data has been taken to repre-

¹<http://commoncrawl.org>

²<http://twitter.com>

sent language-use in specific places (e.g., Eisenstein, et al., 2010; Roller, et al., 2012; Kondor, et al., 2013; Mocanu, et al., 2013; Eisenstein, et al., 2014; Graham, et al., 2014; Donoso & Sanchez, 2017); regional variation in Twitter usage was also the subject of a shared task at PAN-17 (Rangel, et al., 2017). Web-crawled data has also been curated and prepared for the purpose of studying spatial variation (Goldhahn, et al., 2012; Davies & Fuchs, 2015), including the use of country-level domains for geo-referencing (Cook & Brinton, 2017). This paper builds on such previous work by systematically collecting geo-referenced data from both sources on a global scale. The full web corpus is available for download.³

For the Common Crawl data (abbreviated as CC), language samples are geo-located using country-specific top-level domains. The assumption is that a language sample from a web-site under the *.ca* domain originated from Canada (c.f., Cook & Brinton, 2017). This approach to regionalization does not assume that whoever produced that language sample was born in Canada or represents a traditional Canadian dialect group; rather, the assumption is only that the sample represents someone in Canada who is producing language data. Some countries are not available because their top-level domains are used for other purposes (i.e., *.ai*, *.fm*, *.io*, *.ly*, *.ag*, *.tv*). Domains that do not contain geographic information are also removed from consideration (e.g., *.com* sites). The Common Crawl dataset covers 2014 through the end of 2017, totalling 81.5 billion web pages. As shown in Table 1, after processing this produces a corpus of 16.65 billion words.

The basic procedure for processing the Common Crawl data is to look at text within paragraph tags: any document with at least 40 words within paragraph tags from a country-level domain is processed. Noise like navigational items, boilerplate text, and error messages is removed using heuristic searches and also using deduplication: any text that occurs multiple times on the same site or multiple times within the same month is removed. A second round of deduplication is used over the entire dataset to remove texts in the same language that occur in the same country. Its limited scope makes this final deduplication stage possible. For reproducibility, the code used for collecting and

³https://labbcats.canterbury.ac.nz/download/?jonathandunn/CGLU_v3

processing the Common Crawl data is also made available.⁴

The use of country-level domains for geo-referencing raises two questions: First, are there many domains that are not available because they are not used or are used for non-geographic purposes? After removing irrelevant domains like *.tv*, the CC dataset covers 166 countries (30 of which are not included in the Twitter corpus) while the Twitter corpus covers 169 countries (33 of which are not included in the CC corpus). Thus, while the use of domains does remove some countries from consideration, the effect is limited. Second, does the amount of data for each country domain reflect the actual number of web pages from that country? In other words, some countries like the United States are less likely to use their top-level codes. However, the United States is still well-represented in the model. The bigger worry is that regional varieties from Africa or East Asia, both of which are under-represented in these datasets, might be missing from the model.

For the Twitter corpus, a spatial search is used to collect Tweets from within a 50km radius of 10k cities.⁵ Such a search avoids biasing the selection by using language-specific keywords or hashtags. The Twitter data covers the period from May of 2017 until early 2019. This creates a corpus containing 1,066,038,000 Tweets. The language identification component, however, only provides reliable predictions for samples containing at least 50 characters. Thus, the corpus is pruned to include only those Tweets above that length threshold. As shown in Table 1, this produces a corpus containing 4.14 billion words with a global distribution. Language identification (LID) is important here because a failure to identify some regional varieties of English will ultimately bias the model. The LID system used is available for testing.⁶ But given that the focus is a major language, English, the performance of LID is not a significant factor in the overall model of syntactic variation.

The datasets summarized in Table 1 include many languages other than English. The purpose is to provide background information about where robust varieties of English are found: where is

⁴https://github.com/jonathandunn/common_crawl_corpus

⁵<https://github.com/datasets/world-cities>

⁶<https://github.com/jonathandunn/idNet/>

Country	CC	TW
South Africa	53,447,000	57,017,000
Nigeria	113,957,000	29,390,000
Canada	149,882,000	97,835,000
United States	42,890,000	220,947,000
India	71,219,000	80,038,000
Pakistan	140,190,000	34,044,000
Malaysia	198,566,000	18,296,000
Philippines	209,476,000	19,705,000
England	62,811,000	43,376,000
Ireland	43,975,000	46,045,000
Portugual	20,960,000	23,333,000
Switzerland	15,459,000	17,788,000
Australia	29,129,000	98,955,000
New Zealand	87,951,000	37,428,000
TOTAL	1.23 billion	0.82 billion

Table 2: English Varieties by Dataset in N. Words

English discovered when the search is not biased by looking only for English? On the one hand, some regions may be under-represented in these datasets; if national varieties are missing from a region, it could be (i) that there is no national variety of English or (ii) that there is not enough data available from that region. On the other hand, Table 1 shows that each region is relatively well-represented, providing confidence that we are not missing other important varieties.

4 How Many Varieties of English?

We take a simple threshold-based approach to the question of which regional varieties to include: any national variety that has at least 15 million words in both the Common Crawl and Twitter datasets is included in the attempt to model all global varieties of English. This threshold is chosen in order to ensure that sufficient training/testing/development samples are available for each variety. The inventory of national varieties in Table 2 is entirely data-driven and does not depend on distinctions like dialects vs. varieties, inner-circle vs. outer-circle, or native vs. non-native. Instead, the selection is empirical: any area with a large amount of observed English usage is assumed to represent a regional variety. Since the regions here are based on national boundaries, we call these national varieties. We could just as easily call them national dialects.

Nevertheless, the inventory (sorted by region) contains within it some important combinations.

	CC	TW
Training Samples	327,500	308,000
Testing Samples	66,500	64,000

Table 3: Samples by Function and Dataset

There are two African varieties, two south Asian varieties, two southeast Asian varieties, two native-speaker European varieties and two non-native-speaker European varieties. Taken together, these pairings provide a rich ground for experimentation. Are geographically closer varieties more linguistically similar? Is there an empirical reality to the distinction between inner-circle and outer-circle varieties (e.g., American English vs. Malaysian English)? The importance of this language-mapping approach is that it does not assume the inventory of regions.

5 Data Preparation and Division

The goal of this paper is to model syntactic variation using geo-referenced documents taken from web-crawled and social media corpora. Such geo-referenced documents represent language use *in* a particular place but, unlike traditional dialect surveys, there is no assurance that individual authors are native speakers *from* that place. We have to assume that most language samples from a given country represent the native English variety of that country. For example, many non-local residents live in Australia; we only have to assume that *most* speakers observed in Australia are locals.

In order to average out the influence of out-of-place samples, we use random aggregation to create samples of exactly 1,000 words in both corpora. For example, in the Twitter corpus this means that an average of 59 individual Tweets from a place are combined into a single sample. First, this has the effect of providing more constructions per sample, making the modeling task more approachable. Second and more importantly, individual out-of-place Tweets are reduced in importance because they are aggregated with other Tweets presumably produced by local speakers.

The datasets are formed into training, testing, and development sets as follows: First, 2k samples are used for development purposes regardless of the amount of data from a given regional variety. Depending on the size of each variety, at least 12k training and 2.5k testing samples are available. Because some varieties are represented by

much larger corpora (i.e., Tweets from American English), a maximum of 25k training samples and 5k testing samples are allowed per variety per register. This creates a corpus with 327,500 training and 66,500 testing samples (CC) and a corpus with 308,000 training and 64,000 testing samples (TW). As summarized in Table 3, these datasets contain significantly more observations than have been used in previous work (c.f., Dunn, 2018c).

6 Learning the Syntactic Feature Space

Past approaches to syntactic representation for this kind of task used part-of-speech n-grams (c.f., Hirst & Feiguina, 2007) or lists of function words (c.f., Argamon & Koppel, 2013) to indirectly represent grammatical patterns. Recent work (Dunn, 2018c), however, has introduced the use of a full-scale syntactic representations based on grammar induction (Dunn, 2017, 2018a, 2019) within the Construction Grammar paradigm (CxG: Langacker, 2008; Goldberg, 2006). The idea is that this provides a replicable syntactic representation.

A CxG, in particular, is useful for text classification tasks because it is organized around complex *constructions* that can be quantified using frequency. For example, the ditransitive construction in (1) is represented using a sequence of slot-constraints. Some of these slots have syntactic fillers (i.e., NOUN) and some have joint syntactic-semantic fillers (i.e., *V:transfer*). Any utterance, as in (2) or (3), that satisfies these slot-constraints counts as an example or instance of the construction. This provides a straight-forward quantification of a grammar as a one-hot encoding of construction frequencies.

- (1) [NOUN – *V:transfer* – *N:animate* – NOUN]
- (2) “He mailed Mary a letter.”
- (3) “She gave me a hand.”

This paper compares two learned CxGs: first, the same grammar used in previous work (Dunn, 2018c); second, a new grammar learned with an added association-based transition extraction algorithm (Dunn, 2019). These are referred to as CxG-1 (the frequency-based grammar in Dunn, 2019) and CxG-2 (the association-based grammar), respectively. Both are learned from web-crawled corpora separate from the corpora used for modeling regional varieties (from Baroni, et al., 2009; Majljs & Žabokrtský, 2012; Benko,

Country	CxG-1 (CC)	CxG-2 (CC)
South Africa	+4.42%	+4.62%
Nigeria	-0.93%	-0.78%
Canada	+4.03%	+5.17%
United States	-0.98%	-1.90%
India	-3.15%	-10.38%
Pakistan	-4.76%	-17.25%
Malaysia	-3.39%	-11.51%
Philippines	-4.48%	-17.39%
England	+4.59%	+13.98%
Ireland	+4.26%	+18.62%
Portugual	-5.82%	-4.70%
Switzerland	+0.98%	+13.96%
Australia	+3.75%	+8.15%
New Zealand	+1.83%	-0.59%

Table 4: Relative Average Feature Density

2014; and the data provided for the CoNLL 2017 Shared Task: Ginter, et al., 2017). The exact datasets used are available.⁷

In both cases a large background corpus is used to represent syntactic constructions that are then quantified in samples from regional varieties. The grammar induction algorithm itself operates in folds, optimizing grammars against individual test sets and then aggregating these fold-specific grammars at the end. This creates, in effect, one large umbrella-grammar that potentially over-represents a regional dialect. From the perspective of the grammar, we can think of false positives (the umbrella-grammar contains constructions that a regional dialect does not use) and false negatives (the umbrella-grammar is missing constructions that are important to a regional dialect). For dialect identification as a task, only missing constructions will reduce prediction performance.

How well do CxG-1 and CxG-2 represent the corpora from each regional variety? While prediction accuracies are the ultimate evaluation, we can also look at the average frequency across all constructions for each national dialect. Because the samples are fixed in length, we would expect the same frequencies across all dialects. On the other hand, false positive constructions (which are contained in the umbrella-grammar but do not occur frequently in a national dialect) will reduce the overall feature density for that dialect. Because the

⁷https://labbcac.canterbury.ac.nz/download/?jonathandunn/CxG_Data_FixedSize

classification results do not directly evaluate false positive constructions, we investigate this in Table 4 using the average feature density: the total average frequency per sample, representing how many syntactic constructions from the umbrella-grammar are present in each regional dialect. This is adjusted to show differences from the average for each grammar (i.e., CxG-1 and CxG-2 are each calculated independently).

First, CxG-1 has a smaller range of feature densities, with the lowest variety (Portugal English) being only 10.41% different from the highest variety (UK English). This range is much higher for CxG-2, with a 36.01% difference between the lowest variety (Philippines English) and the highest variety (Irish English). One potential explanation for the difference is that CxG-2 is a better fit for the inner-circle dominated training data. This is a question for future work. For now, both grammars pattern together in a general sense: the highest feature density is found in UK English and varieties more similar to UK English (Ireland, Australia). The lowest density is found in under-represented varieties such as Portugal English or Philippines English. Any grammar-adaptation based on dialect identification will struggle to add unknown constructions from these varieties.

7 Modeling National Varieties

The main set of experiments uses a Linear Support Vector Machine (Joachims, 1998) to classify dialects using CxG features. Parameters are tuned using the development data. Given the general robust performance of SVMs in the literature relative to other similar classifiers on variation tasks (c.f., Dunn, et al., 2016), we forego a systematic evaluation of classifiers.

We start, in Table 5, with an evaluation of baselines by feature type and dataset. We have two general types of features: purely syntactic representations (CxG-1, CxG-2, Function words) and potentially topic-based features (unigrams, bigrams, trigrams). The highest performing feature on both datasets is simple lexical unigrams, at 30k dimensions. We use a hashing vectorizer to avoid a region-specific bias: the vectorizer does not need to be trained or initialized against a specific dataset so there is no chance that one of the varieties will be over-represented in determining which n-grams are included. But this has the side-effect of preventing the inspection of individual features. Vec-

Features	Prec.	Recall	F1
CxG-1 (CC)	0.80	0.80	0.80
CxG-1 (TW)	0.75	0.76	0.76
CxG-2 (CC)	0.96	0.96	0.96
CxG-2 (TW)	0.92	0.92	0.92
Funct. (CC)	0.65	0.65	0.65
Funct. (TW)	0.56	0.57	0.55
Unigrams (CC)	1.00	1.00	1.00
Unigrams (TW)	1.00	1.00	1.00
Bigrams (CC)	0.98	0.98	0.98
Bigrams (TW)	0.97	0.97	0.97
Trigrams (CC)	0.87	0.87	0.87
Trigrams (TW)	0.82	0.82	0.82

Table 5: Classification Performance By Feature Set

tors for all experiments are available, along with the trained models that depend on these vectors.⁸

As n increases, n -grams tend to represent structural rather than topical information. In this case, performance decreases as n increases. We suggest that this decrease provides an indication that the performance of unigrams is based on location-specific content (e.g., “Chicago” vs. “Singapore”) rather than on purely linguistic lexical variation (e.g., “jeans” vs. “denim”). How do we differentiate between predictions based on place-names, those based on place-specific content, and those based on dialectal variation? That is a question for future work. For example, is it possible to identify and remove location-specific content terms? Here we focus instead on using syntactic representations that are not subject to such interference.

Within syntactic features, function words perform the worst on both datasets with F1s of 0.65 and 0.55. This is not surprising because function words in English do not represent syntactic structures directly; they are instead markers of the types of structures being used. CxG-1 comes next with F1s of 0.80 and 0.76, a significant improvement over the function-word baseline but not approaching unigrams. Note that the experiments using this same grammar in previous work (Dunn, 2018c) were applied to samples of 2k words each. Finally, CxG-2 performs the best, with F1s of 0.96 and 0.92, falling behind unigrams but rivaling bigrams and surpassing trigrams. Because of this, the more detailed experiments below focus only on the CxG-2 grammar.

⁸https://labbcats.canterbury.ac.nz/download/?jonathandunn/VarDial_19

Country	Prec. (CC)	Recall (CC)	F1 (CC)	Prec. (TW)	Recall (TW)	F1 (TW)
South Africa	0.94	0.96	0.95	0.92	0.94	0.93
Nigeria	0.98	0.98	0.98	0.94	0.95	0.94
Canada	0.94	0.94	0.94	0.84	0.79	0.81
United States	0.93	0.95	0.94	0.85	0.89	0.87
India	0.97	0.98	0.97	0.97	0.97	0.97
Pakistan	1.00	0.99	0.99	0.98	0.98	0.98
Malaysia	0.96	0.96	0.96	0.99	0.99	0.99
Philippines	0.98	0.97	0.98	0.98	0.98	0.98
England	0.95	0.95	0.95	0.87	0.90	0.89
Ireland	0.97	0.97	0.97	0.95	0.95	0.95
Portugual	0.99	0.98	0.98	0.93	0.90	0.92
Switzerland	0.97	0.94	0.96	0.98	0.97	0.97
Australia	0.97	0.96	0.97	0.82	0.83	0.83
New Zealand	0.91	0.92	0.91	0.92	0.90	0.91
W. AVG	0.96	0.96	0.96	0.92	0.92	0.92

Table 6: Within-Domain Classification Performance (CxG-2)

A closer look at both datasets by region for CxG-2 is given in Table 6. The two datasets (web-crawled and social media) present some interesting divergences. For example, Australian English is among the better performing varieties on the CC dataset (F1 = 0.97) but among the worst performing varieties on Twitter (F1 = 0.83). This is the case even though the variety we would assume would be most-often confused with Australian English (New Zealand English) has a stable F1 across domains (both are 0.91). An examination of the confusion matrix (not shown), reveals that errors between New Zealand and Australia are similar between datasets but that the performance of Australian English on Twitter data is reduced by confusion between Australian and Canadian English.

In Table 4 we saw that the umbrella-grammar (here, CxG-2) better represents inner-circle varieties, specifically UK English and more closely related varieties. This is probably an indication of the relative representation of the different varieties used to train the umbrella-grammar: grammar induction will implicitly model the variety it is exposed to. It is interesting, then, that less typical varieties like Pakistan English and Philippines English (which had lower feature densities) have higher F1s in the dialect identification task. On the one hand, the syntactic differences between these varieties and inner-circle varieties means that the umbrella-grammar misses some of their unique constructions. On the other hand, their greater syntactic difference makes these varieties easier to

identify: they are more distinct in syntactic terms even though they are less well represented.

Which varieties are the most similar syntactically given this model? One way to quantify similarity is using errors: which varieties are the most frequently confused? American and Canadian English have 221 misclassified samples (CC), while Canadian and UK English are only confused 36 times. This reflects an intuition that Canadian English is much more similar to American English than it is to UK English. New Zealand and Australian English have 101 misclassifications (again, on CC); but New Zealand and South African English have 266. This indicates that New Zealand English is more syntactically similar to South African English than to Australian English. However, more work on dialect similarity is needed to confirm these findings across different datasets.

8 Varieties on the Web and Social Media

How robust are models of syntactic variation across domains: in other words, does web-crawled data provide the same patterns as social media data? We conduct two types of experiments to evaluate this: First, we take dialect as a cross-domain phenomenon and train/test models on both datasets together, ignoring the difference between registers. Second, we evaluate models trained entirely on web-crawled data against testing data from social media (and vice-versa), evaluating a single model across registers. The point is to evaluate the impact of registers on syntactic variation:

Country	Prec. (CC)	Recall (CC)	F1 (CC)	Prec. (TW)	Recall (TW)	F1 (TW)
South Africa	0.88	0.06	0.10	0.68	0.31	0.43
Nigeria	0.43	0.84	0.57	0.73	0.41	0.52
Canada	0.48	0.14	0.22	0.49	0.27	0.35
United States	0.20	0.87	0.32	0.83	0.16	0.27
India	0.65	0.94	0.77	0.38	0.90	0.54
Pakistan	0.96	0.41	0.58	0.88	0.36	0.51
Malaysia	0.45	0.93	0.61	0.98	0.05	0.10
Philippines	0.73	0.61	0.66	0.87	0.22	0.35
England	0.89	0.01	0.03	0.48	0.44	0.46
Ireland	0.94	0.21	0.35	0.78	0.52	0.62
Portugual	0.02	0.00	0.00	0.22	0.17	0.19
Switzerland	0.92	0.04	0.07	0.12	0.80	0.20
Australia	0.89	0.00	0.00	0.33	0.66	0.44
New Zealand	0.27	0.53	0.36	0.64	0.40	0.49
W. AVG	0.62	0.40	0.33	0.62	0.40	0.40

Table 7: Cross-Domain Models, Trained on CC (Left) and Trained on TW (Right), CxG-2

Country	Prec.	Recall	F1
South Africa	0.91	0.92	0.92
Nigeria	0.94	0.95	0.95
Canada	0.87	0.84	0.85
United States	0.85	0.90	0.87
India	0.96	0.97	0.97
Pakistan	0.98	0.98	0.98
Malaysia	0.97	0.96	0.96
Philippines	0.97	0.97	0.97
England	0.87	0.90	0.89
Ireland	0.94	0.95	0.95
Portugual	0.94	0.90	0.92
Switzerland	0.96	0.93	0.95
Australia	0.87	0.86	0.87
New Zealand	0.89	0.87	0.88
W. AVG	0.92	0.92	0.92

Table 8: Single-Set Classification Performance

does Australian English have the same profile on both the web and on Twitter?

Starting with the register-agnostic experiments, Table 8 shows the classification performance if we lump all the samples into a single dataset (however, the same training and testing data division is still maintained). The overall F1 is the same as the Twitter-only results in Table 6. On the other hand, varieties like Australian English that performed poorly in Twitter perform somewhat better under these conditions. Furthermore, the observation made above that outer-circle varieties are more distinct remains true: the highest perform-

ing varieties are the least proto-typical (i.e., Indian English and Philippines English).

But a single model does not perform well across the two datasets, as shown in Table 7. The model trained on Twitter data does perform somewhat better than its counterpart, but in both cases there is a significant drop in performance. On the one hand, this is not surprising given differences in the two registers: we expect some reduction in classification performance across domains like this. For example, the unigram baseline suffers a similar reduction to F1s of 0.49 (trained on CC) and 0.55 (trained on Twitter).

On the other hand, we would have more confidence in this model of syntactic variation if there was a smaller drop in accuracy. How can we better estimate grammars and variations in grammars across these different registers? Is it a problem of sampling different populations or is there a single population that is showing different linguistic behaviours? These are questions for future work.

9 Unmasking Dialects

How robust are classification-based dialect models to a small number of highly predictive features? A high predictive accuracy may disguise a reliance on just a few syntactic variants. Within authorship verification, unmasking has been used as a meta-classification technique to measure the depth of the difference between two text types (Koppel, et al., 2007). The technique uses a linear classifier to distinguish between two texts using chunks of

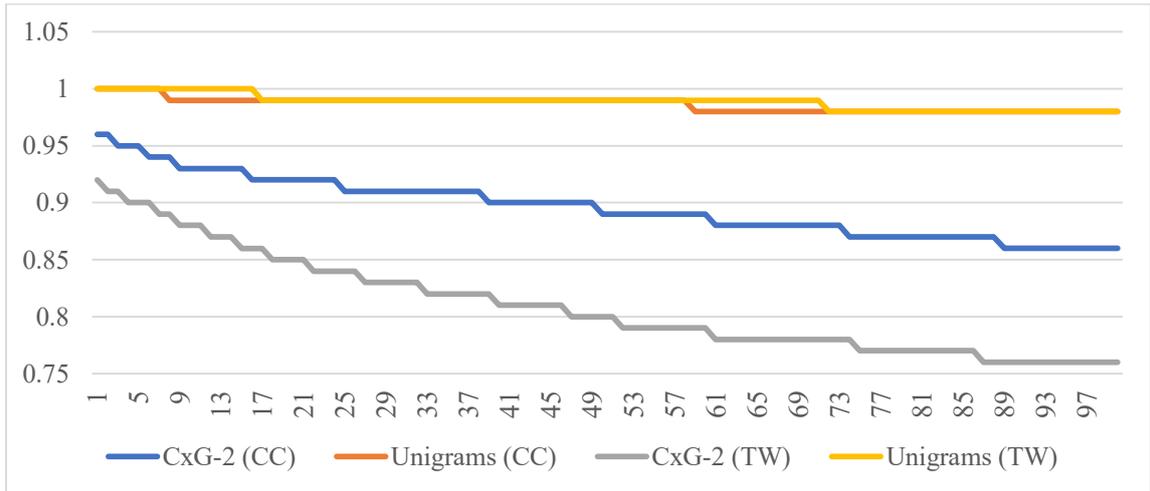


Figure 1: Performance Over 100 Rounds of Unmasking (F1)

the texts as samples. Here we distinguish between dialects with individual samples as chunks. After each round of classification, the most predictive features are removed. In this case, the highest positive and negative features for each regional dialect are removed for the next classification round. Figure 1 shows the unmasking curve over 100 rounds using the F1 score. Given that there are 14 regional dialects in the model, Figure 1 represents the removal of approximately 2,800 features.

For both datasets, the unigram baseline degrades less quickly than the syntactic model. On the one hand, it has significantly more features in total, so that there are more features to support the classification. On the other hand, given that the most predictive features are being removed, this shows that the lexical model has a deeper range of differences available to support classification than the syntactic model. Within the syntactic models, the classifier trained on web-crawled data degrades less quickly than the Twitter model and maintains a higher performance throughout.

This unmasking curve is simply a method for visualizing the robustness of a classification model. The syntactic model is less robust to unmasking than the lexical model. At the same time, we know that the syntactic model does not rely on place-names and place-based content and thus represents a more traditional linguistic approach to variation.

10 Discussion

This paper has used data-driven language mapping to select national dialects of English to be included

in a global dialect identification model. The main experiments have focused on a dynamic syntactic feature set, showing that it is possible to predict dialect membership within-domain with only a small loss of performance against lexical models. This work raises two remaining problems:

First, we know that location-specific content (i.e., place names, place references, national events) can be used for geo-location and text-based models of *place*. To what degree does a lexical approach capture linguistic variation (i.e., “pop” vs. “soda”) and to what degree is it capturing non-linguistic information (i.e., “Melbourne” vs. “London”)? This is an essential problem for dialect identification models. A purely syntactic model does not perform as well as a lexical model, but it does come with more guarantees.

Second, we have seen that inner-circle varieties have higher feature densities given the grammars used here. This implies that there are syntactic constructions in varieties like Philippines English that have not been modeled by the grammar induction component. While dialect identification can be used to optimize regional weights for *known* constructions, how can such *missing* constructions be adapted? This remains a challenge. While the less proto-typical dialects have higher F1s (i.e., Pakistan English), they also have lower feature densities. This indicates that some of their constructions are missing from the grammar. Nevertheless, this paper has shown that a broader syntactic feature space can be used to model the difference between many national varieties of English.

References

- Adams, B. 2015. Finding Similar Places using the Observation-to-Generalization Place Model. *Journal of Geographical Systems*, 17(2): 137-156.
- Adams, B. and McKenzie, G. 2018. Crowdsourcing the character of a place: Character-level convolutional networks for multilingual geographic text classification. *Transactions in GIS*, doi: 10.1111/tgis.12317.
- Ali, M. 2018. Character Level Convolutional Neural Network for Arabic Dialect Identification. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects*, 122-127.
- Argamon, S. and Koppel, M. 2013. A Systemic Functional Approach to Automated Authorship Analysis. *Journal of Law & Policy*, 12:299-315.
- Barbaresi, A. 2018. Computationally efficient discrimination between language varieties with large feature vectors and regularized classifiers. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects*, 164-171.
- Baroni, M., Bernardini, S., Ferraresi, A., and Zanchetta, E. 2009. The WaCky Wide Web: A Collection of Very Large Linguistically Processed Web-crawled Corpora. *Language Resources and Evaluation*, 43: 209-226.
- 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*, 145-152.
- Benko, V. 2014. Aranea: Yet Another Family of (Comparable) Web Corpora. In *Proceedings of Text, Speech and Dialogue. 17th International Conference*. 257-264.
- Calle-Martin, J. and Romero-Barranco, J. 2017. Third person present tense markers in some varieties of English. *English World-Wide*, 38(1): 77-103.
- Clark, L. and Watson, K. 2016. Phonological leveling, diffusion, and divergence: /t/ lenition in Liverpool and its hinterland. *Language Variation and Change*, 28(1): 31-62.
- Collins, P. 2012. Singular agreement in there-existentials: An intervaretial corpus-based study. *English World-Wide*, 33(1): 53-68.
- Cook, P. and Brinton, J. 2017. Building and evaluating web corpora representing national varieties of English. *Language Resources and Evaluation*, 51.
- Davies, M. and Fuchs, R. 2015. Expanding horizons in the study of World Englishes with the 1.9 billion word Global Web-based English Corpus (GloWbE). *English World-Wide*, 36(1): 1-28.
- Donoso, G. and Sanchez, D. 2017. Dialectometric analysis of language variation in twitter. In *Proceedings of the 4th Workshop on NLP for Similar Languages, Varieties and Dialects*. 16-25.
- Dunn, J. 2017. Computational Learning of Construction Grammars. *Language & Cognition*, 9(2): 254-292.
- Dunn, J. 2018a. Modeling the Complexity and Descriptive Adequacy of Construction Grammars. In *Proceedings of the Society for Computation in Linguistics (SCiL 2018)*. Stroudsburg, PA: Association for Computational Linguistics. 81-90.
- Dunn, J. 2018b. Multi-Unit Directional Measures of Association: Moving Beyond Pairs of Words. *International Journal of Corpus Linguistics*, 23(2): 183-215.
- Dunn, J. 2018c. Finding Variants for Construction-Based Dialectometry: A Corpus-Based Approach to Regional CxGs. *Cognitive Linguistics*, 29(2): 275-311.
- Dunn, J. 2019. Frequency vs. Association for Constraint Selection in Usage-Based Construction Grammar. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*. Association for Computational Linguistics.
- Dunn, J.; Argamon, S.; Rasooli, A.; and Kumar, G. 2016. Profile-based authorship analysis. *Literary and Linguistic Computing*, 31(4): 689-710.
- Eisenstein, J.; O'Connor, B.; Smith, N.; and Xing, E. 2010. A latent variable model for geographic lexical variation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 1,227-1,287.
- Eisenstein, J.; O'Connor, B.; Smith, N.; and Xing, E. 2014. Diffusion of lexical change in social media. *PLoS One*, 10.1371.
- Gamallo, P.; Pichel, J.; Algeria, I.; and Agirrezabal, M. 2016. Comparing two Basic Methods for Discriminating Between Similar Languages and Varieties. In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects*, 170-177.
- Ginter, F.; Hajič, J.; Luotolahti, J. 2017. CoNLL 2017 Shared Task - Automatically Annotated Raw Texts and Word Embeddings, LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÁŠFAL), Faculty of Mathematics and Physics, Charles University. <http://hdl.handle.net/11234/1-1989>.
- Goldberg, A. 2006. *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.
- Goldhahn, d.; Eckart, T.; and Quasthoff, U. 2012. Building Large Monolingual Dictionaries at the

- Leipzig Corpora Collection: From 100 to 200 Languages. In *Proceedings of the Conference on Language Resources and Evaluation*.
- Grafmiller, J. & Szmrecsanyi, B. 2018. Mapping out particle placement in Englishes around the world: A study in comparative sociolinguistic analysis. *Language Variation and Change*, 30(3): 385-412.
- Graham, S.; Hale, S.; and Gaffney, D. 2014. Where in the world are you? Geolocation and language identification on twitter. *The Professional Geographer*, 66(4).
- Grieve, J. 2007. Quantitative Authorship attribution: an evaluation of techniques. *Literary and Linguistic Computing*, 22(3): 251–70.
- Hirst, G. and Feiguina, O. 2007. Bigrams of syntactic labels for authorship discrimination of short texts. *Literary and Linguistic Computing*, 22(4): 405–17.
- Hulden, M.; Silfverberg, M.; and Francom, J. 2015. Kernel Density Estimation for Text-Based Geolocation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 145-150.
- Ionescu, R. and Popescu, M. 2016. UnibucKernel: An Approach for Arabic Dialect Identification based on Multiple String Kernels. In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects*, 135-144.
- Joachims, T. 1998. Text categorization with support vector machines: Learning with many relevant features. In C. Nedellec (ed.), *Machine learning: ECML-98: 10th European Conference on Machine Learning*, 137-142. Berlin: Springer.
- Ju, Y.; Adams, B.; Janowicz, K.; Hu, Y.; Yan, B.; and McKenzie, G. 2016. Things and Strings: Improving Place Name Disambiguation from Short Texts by Combining Entity Co-Occurrence with Topic Modeling. In *Proceedings of the 20th International Conference on Knowledge Engineering and Knowledge Management*. LNCS, vol. 10024. Springer, pp. 353-367.
- Kachru, B. 1990 *The Alchemy of English: The spread, functions, and models of non-native Englishes*. Urbana-Champaign: University of Illinois Press.
- Kondor, D.; Csabai, I.; Dobos, L.; Szule, J.; Barankai, N.; Hanyecz, T.; Sebok, T.; Kallus, Z.; and Vattay, G. 2013. Using robust PCA to estimate regional characteristics of language-use from geotagged twitter messages. In *Proceedings of IEEE 4th International Conference on Cognitive Infocommunications (CogInfoCom)*. 393-398.
- Koppel, M., J. Schler, and E. Bonchek-Dokow. 2007. Measuring Differentiability: Unmasking Pseudonymous Authors. *Journal of Machine Learning Research*, 8: 1261-1276.
- Kreutz, T. and Daelemans, W. 2018. Exploring Classifier Combinations for Language Variety Identification. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects*, 191-198.
- Kroon, M.; Medvedeva, M.; and Plank, B. 2018. When Simple n-gram Models Outperform Syntactic Approaches: Discriminating between Dutch and Flemish. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects*, 244-25.
- Kruger, H. and van Rooy, Bertus. 2018. Register variation in written contact varieties of English: A multidimensional analysis. *English World-Wide*, 39(2): 214-242.
- Labov, W.; Fisher, S.; Gylfadottir, D.; and Henderson, A. 2016. Competing systems in Philadelphia phonology. *Language Variation and Change*, 28(3): 273-305.
- Langacker, R. 2008. *Cognitive Grammar: A Basic Introduction*. Oxford: Oxford University Press.
- Lourentzou, I.; Morales, A.; and Zhai, C. 2017. Text-based geolocation prediction of social media users with neural networks. In *Proceedings of 2017 IEEE International Conference on Big Data*, 696-705.
- McNamee, P. 2016. Language and Dialect Discrimination Using Compression-Inspired Language Models. In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects*, 195-203.
- Mocanu, D.; Baronchelli, A.; Perra, N.; Gonçães, B.; Zhang, Q. and Vespignani, A. 2013. The Twitter of Babel: Mapping world languages through microblogging platforms. *PLOSOne*, 10.1371.
- Majliš, M. and Žabokrtský, Z. 2012. Language Richness of the Web. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2012)*. <https://ufal.mff.cuni.cz/w2c>.
- Rangel F., Rosso P., Potthast M., Stein B. 2017. Overview of the 5th Author Profiling Task at PAN 2017: Gender and Language Variety Identification in Twitter. In: Cappellato L., Ferro N., Goeuriot L, Mandl T. (eds.) *CLEF 2017 Labs and Workshops, Notebook Papers. CEUR Workshop Proceedings*. CEUR-WS.org, vol. 1866.
- Roller, S.; Speriosu, M.; Rallapalli, S.; Wing, B.; and Baldrige, J. 2012. Supervised text-based geolocation using language models on an adaptive grid. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. 1,500-1,510.
- Sanders, N. 2007. Measuring syntactic difference in British English. In *Proceedings of the 45th Annual Meeting of the ACL: Student Research Workshop*, 1-6. Association for Computational Linguistics.

- Sanders, N. 2010. *A statistical method for syntactic dialectometry*. Bloomington: Indiana. University dissertation.
- Schilk, M. and Schaub, S. 2016. Noun phrase complexity across varieties of English: Focus on syntactic function and text type. *English World-Wide*, 37(1): 58-85.
- Schreier, D. 2016. Super-leveling, fraying-out, internal restructuring: A century of present be concord in Tristan da Cunha English. *Language Variation and Change*, 28(2): 203-224.
- Strelluf, C. 2016. Overlap among back vowels before /l/ in Kansas City. *Language Variation and Change*, 28(3): 379-407.
- Szmcrcsanyi, B.; Grafmiller, J.; Heller, B.; Rothlisberger, M. 2016. Around the world in three alternations: Modeling syntactic variation in varieties of English. *English World-Wide*, 37(2): 109-137.
- Tamaredo, I. 2018. Pronoun omission in high-contact varieties of English: Complexity versus efficiency. *English World-Wide*, 39(1): 85-110.
- Wing, B. and Baldrige, J. 2014. Hierarchical Discriminative Classification for Text-Based Geolocation. In *Proceedings of the Conference on Empirical Methods in NLP*, 336-348.
- Zeman, D., et al. 2017. CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *Proceedings of the Conference on Natural Language Learning*, 1-19.
- Zeman, D., et al. 2018. CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *Proceedings of the Conference on Natural Language Learning*, 1-21.