

# NEMO: Frequentist Inference Approach to Constrained Linguistic Typology Feature Prediction in SIGTYP 2020 Shared Task

**Alexander Gutkin**

Google, London, United Kingdom  
agutkin@google.com

**Richard Sproat**

Google, Tokyo, Japan  
rws@google.com

## Abstract

This paper describes the NEMO submission to SIGTYP 2020 shared task (Bjerva et al., 2020) which deals with prediction of linguistic typological features for multiple languages using the data derived from World Atlas of Language Structures (WALS). We employ frequentist inference to represent correlations between typological features and use this representation to train simple multi-class estimators that predict individual features. We describe two submitted ridge regression-based configurations which ranked second and third overall in the constrained task. Our best configuration achieved the micro-averaged accuracy score of 0.66 on 149 test languages.

## 1 Introduction

The rapidly developing field of computational linguistic typology (Ramat, 2011) is becoming increasingly popular in natural language processing (NLP) research (Bender, 2016; Ponti et al., 2019), where typological corpora such as the World Atlas of Language Structures (WALS) (Dryer and Haspelmath, 2013) and AUTOTYP (Bickel et al., 2020) are seeing increased use (Naseem et al., 2012; Burdick et al., 2020). Despite their popularity, typological corpora are very sparse. According to Murawaki and Yamauchi (2018), less than 15% of the feature values are present for the 2,679 languages represented in WALS. These databases are human-curated and depend on grammatical descriptions for their development; sparsity is often due to linguistic sources lacking data on particular features for a given language. Developing methodologies for accurately predicting missing typological features on the basis of existing knowledge is therefore crucial for a wider adoption of typological resources in NLP tasks and beyond (Evans and Levinson, 2009).

This paper presents the work done by the “NEMO Team” (Google London and Tokyo) on the *constrained subtask* for the SIGTYP 2020 Shared Task (Bjerva et al., 2020). We experimented with a variety of machine learning models using only the features provided in the training, development and test sets. Our features included genetic features (genus and family), areal features (clusters of languages within a radius of the target language), and derived implicational universals. Originally introduced by Greenberg (1963) and demonstrated to capture the syntactic typology well (Dryer, 1992, 2009; Dunn et al., 2011), similarly to others (Daumé III and Campbell, 2007) we use the framework of universal implications to capture correlations between other types of typological features as well. We describe how these features were derived in detail in Section 3.

As we report below in Section 4, the performance of machine learning algorithms varied across different feature predictions, with some algorithms working better for some features, and less well for others. On balance however we found that ridge regression (Hoerl and Kennard, 1970), also known as Tikhonov regularization (Franklin, 1974), was the most useful approach.

## 2 Related Work

Here we review the approaches corresponding to the constrained sub-task of the shared task, where no external data, such as unlabeled texts, is allowed.

A popular approach to modeling the typological diversity of the world’s languages is based on Bayesian probabilistic inference. Despite recently drawing some criticism on linguistic grounds by Ono (2020), this approach possesses impressive explanatory power. In what possibly represents the earliest model-based typological feature imputer,

Daumé III and Campbell (2007) introduced the probabilistic Bayesian model for uncovering universal implications in WALS data by associating random variables with individual WALS features and discovering the inter-feature correlations from statistical dependence between random variables. The modeling power of the Bayesian approach was further demonstrated by Daumé III (2009) in a non-parametric hierarchical Bayesian model combining linguistic areas and phylogeny.

Murawaki (2015) proposed a deep learning approach to phylogenetic inference by mapping the language vectors to a latent space using an auto-encoder trained using typologically-inspired objective on WALS data with missing values imputed using a regularized iterative variant of Multiple Correspondence Analysis (MCA) (Josse and Husson, 2012; Josse et al., 2012). In later work, Murawaki (2017, 2019) and Murawaki and Yamauchi (2018) abandoned an earlier model in favor of a Bayesian autologistic approach and demonstrated the superiority of Bayesian predictor over MCA. In this approach, the languages are represented as random variables that are explained in terms of other languages related to each other through phylogenetic and spatial neighborhood graphs. Bjerva et al. (2019) introduce a generative model inspired by the Chomskyan principles-and-parameters framework, drawing on the correlations between typological features of languages to tackle the novel task of typological collaborative filtering, a concept borrowed from the area of recommender systems.

While most state-of-the-art missing feature imputation methods are model-based, recently Buis and Hulden (2019) employed the iterative technique based on singular value decomposition (SVD) from the well-studied area of low-rank matrix completion and reported the performance on par with the prior art. Although lacking explanatory power, similarly to MCA, such techniques are attractive due to their simplicity and computational efficiency. Takamura et al. (2016) took a standard machine learning approach by training multinomial logistic regression classifiers for individual WALS features based on other features present in the database under various experimental conditions, hypothesizing that the classifier would capture feature correlations implicit in the data.

The frequentist approach we take is similar in some respects to work of Takamura et al. (2016)

in that we also train vector space classifiers, but there are a few notable differences, we: (i) use typologically-motivated “probabilistic” frequency-based input space by explicitly representing areal and phylogenetic associations and implicational universals, and (ii) explore a wider range of classification approaches.

Although for this task there is a special interest in the geographic aspect of the modeling, for our final submission we limited ourselves to the rather orthodox approach of representing language associations through fixed neighborhoods, in order not to over-complicate our method (described in the next section) that can be difficult to reconcile with the more sophisticated models, such as the very promising model of language evolution from Kauhanen et al. (2019) and the findings emerging from the fields of dialectology and dialectometry (Szmrecsanyi, 2011; Wieling and Nerbonne, 2015; Nerbonne et al., 2020).

### 3 Method

Here we outline the details of our approach used to generate the final submission. The open-source implementation of our training and evaluation pipeline has been released in public domain.<sup>1</sup>

#### 3.1 Precomputation of Features

Typological features were preprocessed to find likely associations between genetic and areal properties of the language. For each typological feature  $f$  and value  $v$  from a set of its values  $V_f$  we computed the following probability estimates:

$$p_{\text{genus}}(v|f) = \frac{\text{count}_{\text{genus}}(v)}{\text{count}_{\text{genus}}(f)} \quad (1)$$

$$p_{\text{family}}(v|f) = \frac{\text{count}_{\text{family}}(v)}{\text{count}_{\text{family}}(f)} \quad (2)$$

$$p_{\text{area}}(v|f) = \frac{\text{count}_{\text{area}}(v)}{\text{count}_{\text{area}}(f)}, \quad (3)$$

where  $\text{count}(f) = \sum_{v_i \in V_f} \text{count}(v_i)$ . Here, “family” and “genus” in equations 1 and 2 were as given in the data, and “area” in equation 3 comprised all languages within a 2,500 kilometer radius around the target language’s latitude and longitude computed using the Haversine formula (Robusto, 1957).<sup>2</sup>

<sup>1</sup>Available at [https://github.com/google-research/google-research/tree/master/constrained\\_language\\_typology](https://github.com/google-research/google-research/tree/master/constrained_language_typology).

<sup>2</sup>A reviewer noted that the hard limit of 2,500 kilometers seemed arbitrary and wondered why we do not weight “neigh-

Type	Example
Family	Niger-Congo Green_and_Blue  3 Black/green/blue 0.667 9
Genus	Bantoid Green_and_Blue  3 Black/green/blue 1.000 2
Area	8.0,4.33333 Green_and_Blue  3 Black/green/blue 0.692 13
Implicational	Red_and_Yellow 2 Red/yellow  Green_and_Blue 3 Black/green/blue  0.583 12 0.148 54

Table 1: Examples of precomputed most likely associations for color features.

In addition we computed a set of *implicational universals* (Greenberg, 1963). For each feature value  $v_j$ , for feature  $f_j$ , we compute the probability of  $v_j$  given  $f_j$ , and each  $f_i, v_i$  pair from the set of known feature-value pairs in the data

$$p(v_j|f_j, f_i, v_i) = \frac{\text{count}(v_j)}{\text{count}(f_j, f_i, v_i)}. \quad (4)$$

For each of the genetic (family and genus), areal (neighborhood) and universal implication types of associations we kept a separate table, where, for all the known features  $f$  we stored

- the feature  $f$ ,
- the total number of samples in the data of the given type with  $f$ , denoted  $c(f)$ ,
- the value  $v$  with the highest estimated probability per the above equations, denoted  $v^{\text{maj}}$ ,
- the prior  $p^{\text{maj}}$  corresponding to  $v^{\text{maj}}$ .

Examples of the most likely associations computed above are shown in the first three rows of Table 1. For the Niger-Congo language family, the most likely value for the feature Green\_and\_Blue (observed 9 times) is 3 Black/green/blue, with the corresponding prior 0.667. For the Bantoid language genus, the feature Green\_and\_Blue was observed twice, both times with the same value 3 Black/green/blue. The areal example corresponds to the neighborhood of Yoruba,

boring languages according to closeness and use weighted clustering instead.” We agree that in principle more sophisticated approaches would be nice, but one should bear in mind that the geographic centroids for languages provided in the data are at best crude, and so doing anything more sophisticated seemed to us to be crude. Also, to do this properly, distance is really not sufficient: one would also need to account for the presence of possible barriers to contact, including impassable mountain ranges, seas, and hostile neighbors, elements that would be hard to model.

Partition	Genus	Family	Area	Implicational
TRAIN	11,369	21,592	213,441	144,719
TRAIN+DEV	12,352	23,098	240,840	147,430
TRAIN+DEV+TEST	12,904	24,130	241,438	148,923

Table 2: Precomputed typological feature associations.

$(\phi, \lambda)=(8.0, 4.3)$  (where  $\phi$  denotes latitude), for which 13 Green\_and\_Blue features were observed, with the most likely value corresponding to 3 Black/green/blue with prior 0.692.

In addition for the implicational features, we stored the prior probability of  $v_i$  given the conditional feature  $f_i$  from equation 4, and the total count for  $f_i$ . As an example of a weak implicational preference consider the example given in the fourth row of Table 1, which means that if a language has 2 Red/yellow for the feature Red\_and\_Yellow, then there is a slight preference ( $p=0.583$ , estimated on the basis of 12 examples) for having 3 Black/green/blue as the value for Green\_and\_Blue. In other words,  $2 \text{ Red/yellow} \supset 3 \text{ Black/green/blue}$ . There were 54 cases of Green\_and\_Blue in the training data, for which the estimated a priori probability for 3 Black/green/blue is 0.148.

Table 2 shows the overall sizes of the association tables for the genetic, areal and implicational types described above for the different partitions of the shared task data.

### 3.2 Sparse Language Vectors

For each typological feature we train a separate feature estimator, resulting in 185 estimators overall. When training an individual estimator, we represent each language in the training and development set as a sparse feature vector. Likewise, the languages whose features need to be predicted at test time are also represented similarly.

The makeup of individual language vector for a given typological feature  $f$  and a language  $l$  is shown in Table 3. The vector consists of dense and sparse subvectors; the components shown in the first four rows of the table are mostly dense. The first subvector consists of language’s latitude  $\phi$  and longitude  $\lambda$  coordinate, represented as two numeric features exactly the same as given in the shared task data. There is no particular rationale for this choice other than we previously found that for a different task (speech synthesis) choosing alternative representations for the language location (e.g., distances to all other languages in the train-

Subvector	Size	Components		Type
Location	2	Latitude	$\phi(l)$	Numeric
		Longitude	$\lambda(l)$	Numeric
Genus	3	Value	$v_{\text{genus}}^{\text{maj}}(f)$	Categorical
		Prior	$p_{\text{genus}}^{\text{maj}}(f)$	Numeric
		Count	$c_{\text{genus}}(f)$	Numeric
Family	3	Value	$v_{\text{family}}^{\text{maj}}(f)$	Categorical
		Prior	$p_{\text{family}}^{\text{maj}}(f)$	Numeric
		Count	$c_{\text{family}}(f)$	Numeric
Area	3	Value	$v_{\text{area}}^{\text{maj}}(f)$	Categorical
		Prior	$p_{\text{area}}^{\text{maj}}(f)$	Numeric
		Count	$c_{\text{area}}(f)$	Numeric
...	...	...	...	...
Implicational <sub><i>i</i></sub>	5	Value	$v^{\text{maj}}(f_i)$	Categorical
		Prob.	$p(v_i f, v, f_i)$	Numeric
		Count	$c(f, v, f_i)$	Numeric
		Prior	$p(v_i f_i)$	Numeric
Count	$c(f_i)$	Numeric		
...	...	...	...	...

Table 3: Language vector for typological feature  $f$  and language  $l$ .

ing set) in the input features did not significantly improve the results (Gutkin and Sproat, 2017).

The next three subvectors representing genus, family and area, are structured similarly using the three components used in association tables described previously: the majority value  $v^{\text{maj}}(f)$  represented as a categorical feature, the prior corresponding to this value  $p^{\text{maj}}(f)$  and the feature frequency  $c(f)$ , both represented as numeric features. For these three subvectors, the missing values are represented by the three-tuple  $(v_\emptyset, 10^{-6}, 0)$ , where  $v_\emptyset$  denotes a global dummy typological feature value.<sup>3</sup>

The first four subvectors described above are followed by multiple subvectors representing individual universal implications, as shown in the fifth row of Table 3. Each implicational, describing the dependence of feature  $f_i$  on  $f$ , is represented as a five-tuple whose elements are stored in the associations table for implicational universals: The most likely value  $v^{\text{maj}}(f_i)$  of  $f_i$  corresponding to the highest conditional probability  $p(v_i|f, v, f_i)$  (interpreted as probability of  $f_i$  taking value  $v_i$  given that  $f$  is  $v$ ), the total count  $c(f_i, f, v)$  of  $f_i$  when  $f$  is  $v$ , the prior  $p(v_i|f_i)$  and the total count  $c(f_i)$  for  $f_i$  when  $f$  is  $v$ . The missing implicational is represented as a five-tuple  $(v_\emptyset, 10^{-6}, 0, 10^{-6}, 0)$ . As mentioned above, the implicational portion of the language vector is very sparse because, for a feature  $f \in F$  the language vector belongs to, one needs to compute all its correlations to other fea-

<sup>3</sup>The choice of  $10^{-6}$  for a missing value prior is arbitrary. Any small non-zero value valid for a log transform is suitable.

	Set	Langs	Families	Genera	Features		
					Types	Unique	Total
TRAIN	1,125	131	283	185	993	42,698	
DEV	83	36	53	182	628	3,246	
TEST	149	35	45	179	577	3,056	
Total	1,357	145	311	185	1,034	49,000	

Table 4: The three partitions of SIGTYP data.

tures  $f_i \in F$ , where  $F$  is the set of all 182 known features. Since the typological database is very sparse, most of the observed correlations between  $f$  and  $f_i$  for a given language  $l$  are poorly instantiated.

The categorical features in the language vector are represented using a one-hot encoding and the numeric features are scaled to zero mean and unit variance. For probability components, prior to numeric feature scaling, the probability features are transformed into log domain.

The overall representation results in the language vectors with a rather high dimensionality. For example, the language vectors for the Order\_of\_Subject,\_Object\_and\_Verb feature have the dimension of 4111. As we shall see below from the shared task details, this representation may already be too specific given the training set which only contains 1,125 data points. This observation also explains our choice of representing features and implicational with the attributes associated with their most likely majority values rather than all the values for that feature observed in the data, as suggested by a reviewer — doing so will dramatically increase the dimension of input feature space even further and render out approach completely intractable.

## 4 Experiments and Discussion

In what follows, we provide a brief overview of the datasets, introduce the baseline systems we evaluated against during the development of our method, provide the evaluation results of miscellaneous machine learning algorithms on the development set that guided our final model selection, describe the two configurations submitted to the constrained subtask and, finally, mention the approaches that did not work well in our settings.



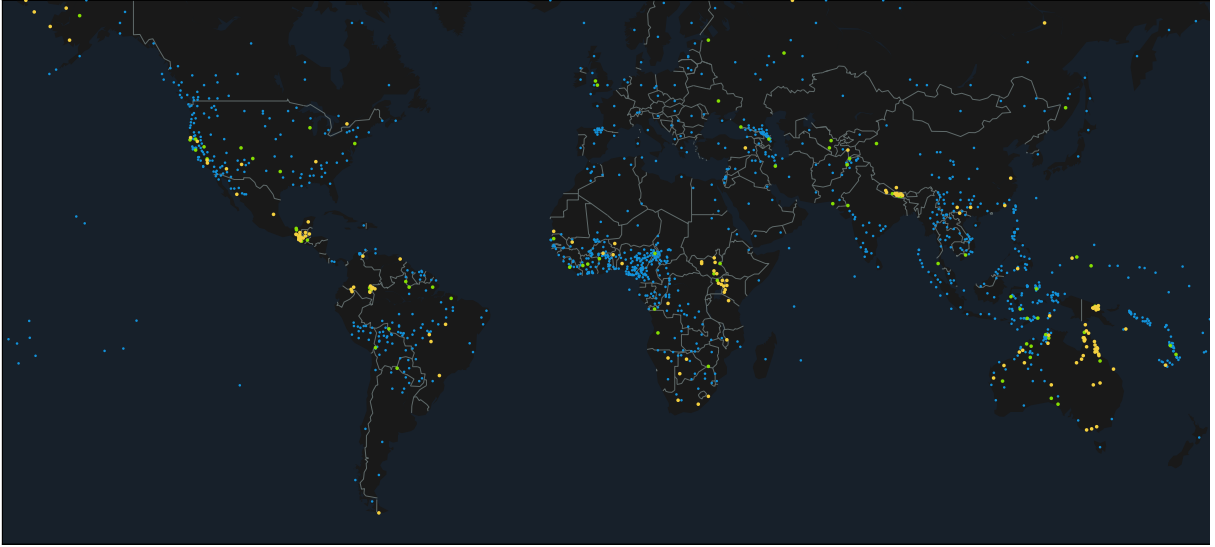


Figure 1: Languages in the SIGTYP training (blue), development (green) and test (yellow) sets.

#### 4.1 Data Overview

Some of the properties of the data used in our experiments<sup>4</sup> are shown in Table 4, where we summarize the counts for three partitions of the data: the training and development sets, and the *blind* test set. For each of the sets, the total number of languages is shown along with the number of unique language families and genera. The count of typological feature types is displayed along with the number of unique feature values and the total number of observed values. As can be seen from the table, the size of the data is very small, which precludes us from using state-of-the-art deep learning methods, especially given our approach to representing each language as a point in high-dimensional space.

The 1,357 longitude and latitude language coordinates provided in the data are displayed in Figure 1 (best viewed in color) using the Mercator projection. The languages in the training set are shown in blue, the development set languages in green and the test set languages in yellow.

#### 4.2 Baselines

For our baselines five simple prediction algorithms based on deterministic search were implemented. Initially we evaluated these systems on the development (DEV) set and later on the test set, after the golden truth data was released by the organizers.

<sup>4</sup>The SIGTYP data is available from <https://github.com/sigtyp/ST2020/tree/master/data>.

The *global* majority class predictor (denoted  $B_1$ ) accumulates the frequencies of all the typological feature values in the training data and predicts the most frequent value for the feature in question irrespective of its phylogenetic or areal attributes. The *clade* majority class predictor (denoted  $B_2$ ) extends the previous predictor by also taking into account language genera and families, producing predictions of the form

$$\hat{v}_{B_2}(f) = \begin{cases} v^{\text{maj}}(f) & \text{if } c_{\text{genus}}(f) = 0 \text{ and} \\ & c_{\text{family}}(f) = 0, \\ v_{\text{family}}^{\text{maj}}(f) & \text{if } c_{\text{genus}}(f) = 0, \\ v_{\text{genus}}^{\text{maj}}(f) & \text{if } c_{\text{genus}}(f) > 0. \end{cases}$$

Predictor  $B_3$  only relies on the distances between languages' geographic coordinates defined in the data. For each typological feature in question, the predictor returns the value belonging to the closest language, according to the Haversine formula (Robusto, 1957), with a known value for that feature. The Haversine distance  $d(l_i, l_j)$  is computed between each pair of languages  $l_i$  and  $l_j$  represented as points on an ideal sphere using their respective latitude and longitude coordinates  $(\phi_i, \lambda_i)$  and  $(\phi_j, \lambda_j)$ .

A naïve approach for combining areal and phylogenetic knowledge is implemented by predictor  $B_4$  which performs its search in two additional clusters of inter-language distances, grouped by genera and families: The algorithm first searches for

Sets			Systems				
Train	Test	# Evals	B <sub>1</sub>	B <sub>2</sub>	B <sub>3</sub>	B <sub>4</sub>	B <sub>5</sub>
TRAIN	DEV	3,246	53.67	62.69	58.87	59.33	<b>63.56</b>
TRAIN+DEV	TEST	5,477	54.66	<b>63.87</b>	58.90	61.68	59.94

Table 5: Micro-averaged accuracies (%) for the five baseline systems.

the known feature within the geographically closest languages of the same genus, followed by a search within the closest languages from the same family, falling back to the closest languages from the global set.

The final baseline configuration B<sub>5</sub> uses an ensemble approach by combining areal and phylogenetic information by using majority voting. The areal estimate is provided by the majority estimate from the candidate’s neighborhood (as described in Section 3.1) or, if that information is not available due to the nearest language being located beyond the neighborhood radius, by the known feature from the closest languages outside the neighborhood. The phylogenetic clade estimates are provided by the majority class values from genera and families (B<sub>2</sub>), respectively. The final estimate is produced by the majority voting, where at least two predictions have to agree, otherwise the predictor falls back to global majority class estimates (B<sub>1</sub>).

The micro-averaged accuracies for five of the baseline systems are shown in Table 5. We performed the evaluation by first obtaining the estimates from a set shown in the first column and then predicting *each* known feature value for the languages from a set shown in the second column (DEV or golden TEST) by pretending that this value is unknown. Best accuracies are highlighted in bold. As can be seen from the table, the best baseline configuration on the DEV set is the B<sub>5</sub> ensemble method, with the pure clade-based B<sub>2</sub> coming second. On the golden TEST set, B<sub>2</sub> is the winning configuration, while the second best configuration is B<sub>4</sub>. It is interesting to note, that the methods that completely ignore phylogeny (B<sub>1</sub> and B<sub>3</sub>) did not perform well in either evaluation.

As one reviewer notes, the apparent predominance of phylogeny in our results may seem surprising given that areal features are well-known to be important in many areas of the world — e.g. India (Emeneau, 1956). This is likely due at least in part to the fact that many of the language families in the sample are small families spoken in a relatively limited geographic area, or else are families

Algorithm	Acc.	Algorithm	Acc.
AdaBoost	41.56	LDA	47.84
Bagging Ensemble	49.04	Linear SVM	55.36
Decision Tree	44.27	Logistic Regression	58.78
DNN	55.55	QDA	34.32
Gaussian Naive Bayes	51.91	Random Forest	63.00
Gaussian Process	23.38	Ridge Regression	<b>75.88</b>

Table 6: Micro-averaged accuracies (%) of various algorithms on the development set.

like Pama-Nyungan, which are more or less isolated from unrelated neighbors. This would tend to confound the influence of phylogeny versus geography, since it is only when *part* of a language family is spoken in an area that is populated by speakers of unrelated families that one will see robust effects of geography.

### 4.3 Model Selection Using Development Set

We approach the feature prediction task by training 185 multi-class classifiers, one for each typological feature. Several standard machine learning algorithms were evaluated on the development set to establish the most optimal algorithm for the task. In particular, we trained decision trees (Breiman et al., 1984), Naive Bayes (Langley et al., 1992) with Gaussian mixtures, Gaussian processes (Rasmussen and Williams, 2006), classifiers based on linear and quadratic discriminant analysis (LDA and QDA, respectively) (Tharwat, 2016), support vector machines (SVM) with linear kernel (Suykens and Vandewalle, 1999), multinomial logistic regression (Böhning, 1992), ridge regression (Hoerl and Kennard, 1970) and simple feed-forward neural networks with a single layer of 200 units (DNN). In addition, three ensemble configurations were also evaluated: multi-class AdaBoost (with 100 estimators) (Zhu et al., 2009), random forests (with 200 estimators, minimum of three samples per leaf and information gain as splitting criterion) (Breiman, 2001) and bagging ensembles (Breiman, 1999). For all the algorithms we used the implementation provided by scikit-learn toolkit.<sup>5</sup> We mostly used default hyperparameters provided performing no special tuning.

The evaluation results for each of the classifiers are shown in Table 6. As can be seen from the table, there are only five algorithms out of twelve which are remotely comparable to our development set baselines summarized in Table 5: DNN, linear

<sup>5</sup><https://scikit-learn.org/>

SVM, multinomial logistic regression, random forest and ridge regression. Out of those, the random forest ensemble method and the ridge regression are the only competitive models, with ridge regression estimation strongly outperforming all of our baselines on the development set achieving the micro-averaged accuracy of 75.88%. On the basis of this result we chose ridge regression algorithm for our experiments on the test set and the final submission.

It is worth noting, however, that it is not difficult to find cases where the ridge regression estimator does not perform as well as other methods. For example, for feature `SVNegO_Order`, the DNN (93% accuracy) and SVM (87% accuracy) both outperform ridge regression (80%). We hypothesized that a better alternative to fixing a certain method for all the features is to employ an ensemble of classifiers which are feature-specific. Although we include this approach in our implementation, so far we only tested it under scenario of performing an iterative stratified  $k$ -fold cross-validation (Witten et al., 2011) over the training, rather than development, set. The optimal feature-specific classifiers obtained by this method did not fare well on the development set, which is probably an indicator that development set was not an optimal reflection of the training data.

Aside from the ridge regression and random forest ensemble methods, the poor average performance of other classification algorithms on the development set is rather unexpected. We hypothesize that this may be due to several confounding factors. The first issue is that our approach of representing training data for each typological feature as sparse language vectors results in small amounts of high-dimensional training data that may not be enough to reliably train most of the classifiers in our particular setting. The second issue is the high feature sparsity which adversely affects multi-class classification for most of the algorithms. Consider the `Order_of_Subject_Object_and_Verb` feature. The frequencies of its seven values (corresponding to class label counts) in the training data represented as an ordered set are  $\{311, 226, 113, 51, 16, 6, 1\}$ . Although we employ balanced class weighting for all the algorithms, using the label values to adjust the weights inversely proportional to class frequencies observed in the data, this may not be enough. Some unsatisfactory

attempts to further remedy this are described below.

We further note that the relatively good performance of ridge regression and random forest classifiers on the development data without any hyper-parameter fine-tuning may possibly be explained by the relative robustness of both methods to sparseness and collinearity effects (which severely affect other types of parametric and non-parametric predictors in our experiment), as previously analyzed in detail by Tomaschek et al. (2018) and observed by others in a typological setting (Burdick et al., 2020) and elsewhere (Josifoski et al., 2019). We briefly reevaluate these assumptions on the released test set in Section 4.6.

#### 4.4 Shared Task Submission

Our submission to shared task consists of two ridge regression classifier configurations denoted `NEMO_system1` and `NEMO_system2`.<sup>6</sup> The difference between the two configurations is in how the training data is generated from the original shared task subsets. The training data for `NEMO_system1` classifier consists of the language vectors generated from the original training (`train.csv`) and development (`dev.csv`) sets only. For `NEMO_system2`, the training data also includes the phylogenetic, areal and implicational relations computed from the *known* features in the blind test set (`test_blinded.csv`).

#### 4.5 Approaches That Disappointed

In addition to the experiments described above, we also tried applying missing feature imputation algorithms to the data provided, as a preprocessing step prior to training the classifiers described above. Neither of the feature imputation algorithms that we tried produced satisfactory results that could improve upon our best baseline. One of the imputation approaches that we evaluated was Multiple Imputation with Denoising Autoencoders (MIDAS) by Lall and Robinson (2020). This approach is particularly attractive because it natively supports categorical features. However, due to reliance on a denoising autoencoder architecture (Vincent et al., 2010), the network requires large amounts of training data to be estimated reliably, which prevented us from getting adequate

---

<sup>6</sup>Submitted as `NEMO_system1_assoc-train-dev_constrained` and `NEMO_system2_assoc-train-dev-test_constrained`, respectively.

Family	# Languages		Acc.
	Train	Test	
Mayan	6	17	0.73
Mahakiranti	4	13	0.72
Tucanoan	3	8	0.67
Nilotic	3	15	0.65
Madang	3	9	0.68
N. Pama-Nyungan	2	24	0.67

Table 7: Numbers of languages from each test family in the test and training data, along with our best system’s per-family accuracy.

performance on the provided typological data.<sup>7</sup>

Similarly disappointing were the attempts to address the class imbalance problem during multi-class classifier training using class resampling techniques. Similar to the feature imputation above, resampling was applied as a preprocessing step prior to training the classifiers. In particular, neither the application of synthetic minority over-sampling technique (SMOTE) (Chawla et al., 2002) nor the adaptive synthetic sampling (He et al., 2008) produced any improvements over our best baseline.

We also extended our method to use country code information provided with the task data for each language. This was achieved by accumulating the per-country typological feature priors, similar to other areal and phylogenetic features, and representing the country codes as categorical features in classification, the process described in Section 3. Ridge regression estimator that included country code information achieved micro-averaged accuracy of 75.66% on the development set – a small deterioration of 0.3% compared to our best result. This may be due to multiple errors in the provided data, e.g., the country code for the Chepang language from Sino-Tibetan family spoken in Nepal (Caughley, 1982) is defined as US in the development set.

#### 4.6 Discussion

We start with some observations about the rather apparent differences between the test data and the training and development data. In Table 7 we list the six language families evidenced in the test data, and the numbers of languages for each represented in the training and development data, versus the test data. As can be seen some of the language families are significantly more represented in the test data than they are in the data that was released

<sup>7</sup>The MIDAS imputer for linguistic typological data is released with the rest of our software.

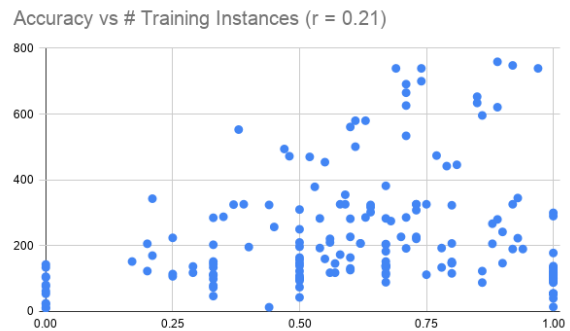


Figure 2: Feature accuracy versus number of instances of the feature in the training and development data ( $r=0.21$ ).

earlier. This is particularly the case for Northern Pama-Nyungan, which comprised two languages in the training-development data, and twenty four languages in the test data. This skew can also be seen in Figure 1 where the yellow dots represent the test data. There are islands of yellow in Australia, Northern South Asia, Central East Africa and in Central America, surrounded by seas of green and blue.<sup>8</sup>

Turning now to our results, we discuss here the performance of our best model NEMO\_system2 (overall micro-averaged accuracy 0.66). First of all, we note that there is very little correlation between the number of exemplars of a feature in training and development, and our system’s performance on that feature in the test set (correlation coefficient  $r=0.21$ ): see Figure 2.<sup>9</sup>

Second, comparing our results against the best baseline provided by the task organizers — frequency-baseline\_constrained (overall micro-averaged accuracy 0.51), we list in Tables 8 and 9 the five features for which our system had the largest win and loss, respectively in terms of absolute accuracy difference. Note that three of the five features for which we showed the largest gains relate to word order. While this may be due to the automatically derived implicational features used in our models — word-order-related features being probably among the most robust of the implicational universals — it is also true that these features are better instantiated in the data.

<sup>8</sup>As one of the reviewers notes, it would be useful for future tasks if the organizers could provide a rationale for the data splits chosen.

<sup>9</sup>A reviewer asks if there is any correlation between accuracy and the number of value settings for each feature in the training and development set. This correlation is somewhat higher, and negative ( $r= -0.34$ ), not surprising since we would expect to do worse if there are more possible values.



Feature	baseline	our best	# in test
Inclusive/Exclusive_Distinction_in_Verbal_Inflexion	0.00	0.75	4
Rel_between_the_Order_of_Obj_and_Vb_and_the_Order_of_Adj_and_Noun	0.14	0.86	35
Order_of_Object_and_Verb	0.26	0.89	38
Order_of_Adposition_and_Noun_Phrase	0.26	0.89	27
Epistemic_Possibility	0.20	0.80	5

Table 8: Five features on which our best system most outperformed the best baseline, with accuracies for each, and counts for each in the test data.

Feature	baseline	our best	# in test
Person_Marking_on_Adpositions	0.50	0.20	10
Genitives,_Adjectives_and_Relative_Clauses	1.00	0.67	3
Noun_Phrase_Conjunction	0.67	0.33	3
Subtypes_of_Asymmetric_Standard_Negation	1.00	0.66	5
Obligatory_Possessive_Inflexion	1.00	0.50	6

Table 9: Five features on which our best system most underperformed the best baseline, with accuracies for each, and counts for each in the test data.

Algorithm	Acc.	Algorithm	Acc.
AdaBoost	0.44	LDA	0.41
Bagging Ensemble	0.50	Linear SVM	0.56
Decision Tree	0.47	Logistic Regression	0.60
DNN	0.58	QDA	0.31
Gaussian Naive Bayes	0.46	Random Forest	0.62
Gaussian Process	0.23	Ridge Regression	<b>0.66</b>

Table 10: Micro-averaged accuracies of various algorithms on the test set.

One reviewer notes that it would be interesting to see the results of ablation studies for some of the high dimensional features discussed in Section 3.2. We agree, but leave this for future work.

Finally, recall from the discussion of development set-based model selection in Section 4.3 that ridge regression was the best performing model, followed by the non-parametric random forest predictor (see Table 6). Once the golden test data was released, we evaluated the same set of twelve classifiers on the test set with the results shown in Table 10. As can be seen from the table, the ranking of three best classifiers on the development set is preserved on the held-out test set as well: ridge regression is the best-performing classifier in our task, followed by random forest and, finally, multinomial logistic regression.

## 5 Conclusion

We have presented the NEMO submission to the 2020 SIGTYP shared task. Our system used genetic, geographical and automatically deduced implicational universals, and a range of classifiers of which ridge regression yielded the best overall performance. Our method achieved 0.66 overall

accuracy on the task, compared to the best baseline provided by the organizers, which had an accuracy of 0.51. Our system tended to do better on test language families that were better represented in the training and development data, though interestingly the correlation between the number of instances of *features* in the training/development data and performance on that feature was not strong.

As discussed above, we include in the training package methods for ensembling classifiers so that one can find per-feature optimal classifiers, but we did not make use of this functionality for our submitted results. We do think, however, that further experimentation with this functionality would be useful.

Finally, as one of the reviewers notes, it would be interesting to ask what use work along the lines presented here could be to the field linguist/typologist who is interested in testing potential relationships between languages based on shared features. Many of the systems reported in Section 2 had this aim in mind. While we do not want to overstress the value of this sort of work compared to good linguistic intuition, we do think that knowing the Bayesian priors of associations, as we discussed in Section 3.1 could at least serve as a reminder that some feature settings may be shared simply because they are very common anyway.

## Acknowledgments

The authors thank the anonymous reviewers for many helpful suggestions on an earlier version of this paper.

## References

- Emily M Bender. 2016. [Linguistic typology in natural language processing](#). *Linguistic Typology*, 20(3):645–660.
- Balthasar Bickel, Johanna Nichols, Taras Zakharko, Alena Witzlack-Makarevich, Kristine Hildebrandt, Michael Rießler, Lennart Bierkandt, Fernando Zúñiga, and John B. Lowe. 2020. The AUTOTYP database. *University of Zurich: Department of Comparative Language Science*. Version 0.1.1: <https://github.com/autotyp/autotyp-data>.
- Johannes Bjerva, Yova Kementchedjiev, Ryan Cotterell, and Isabelle Augenstein. 2019. [A probabilistic generative model of linguistic typology](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1529–1540, Minneapolis, Minnesota. Association for Computational Linguistics.
- Johannes Bjerva, Elizabeth Salesky, Sabrina Mielke, Aditi Chaudhary, Giuseppe G. A. Celano, Edoardo M. Ponti, Ekaterina Vylomova, Ryan Cotterell, and Isabelle Augenstein. 2020. SIGTYP 2020 Shared Task: Prediction of Typological Features. In *Proceedings of the Second Workshop on Computational Research in Linguistic Typology*. Association for Computational Linguistics.
- Dankmar Böhning. 1992. [Multinomial logistic regression algorithm](#). *Annals of the Institute of Statistical Mathematics*, 44(1):197–200.
- Leo Breiman. 1999. Pasting small votes for classification in large databases and on-line. *Machine Learning*, 36(1-2):85–103.
- Leo Breiman. 2001. [Random forests](#). *Machine Learning*, 45(1):5–32.
- Leo Breiman, Jerome Friedman, Charles J. Stone, and Richard A. Olshen. 1984. *Classification and Regression Trees*. Chapman & Hall/CRC Press.
- Annebeth Buis and Mans Hulden. 2019. [Typological feature prediction with matrix completion](#). In *Proceedings of TyP-NLP: The First Workshop on Typology for Polyglot NLP*, pages 13–15, Florence, Italy. Association for Computational Linguistics.
- Laura Burdick, Jonathan K Kummerfeld, and Rada Mihalcea. 2020. [Analyzing the surprising variability in word embedding stability across languages](#). *arXiv preprint arXiv:2004.14876*.
- Ross Charles Caughley. 1982. *The syntax and morphology of the verb in Chepang*. Pacific Linguistics. Department of Linguistics, Research School of Pacific Studies, The Australian National University.
- Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. [SMOTE: Synthetic minority over-sampling technique](#). *Journal of Artificial Intelligence Research*, 16:321–357.
- Hal Daumé III. 2009. [Non-parametric Bayesian areal linguistics](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 593–601, Boulder, Colorado. Association for Computational Linguistics.
- Hal Daumé III and Lyle Campbell. 2007. [A Bayesian model for discovering typological implications](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 65–72, Prague, Czech Republic. Association for Computational Linguistics.
- Matthew S Dryer. 1992. [The Greenbergian word order correlations](#). *Language*, pages 81–138.
- Matthew S Dryer. 2009. [The Branching Direction Theory of word order correlations revisited](#). In Sergio Scalise, Elisabetta Magni, and Antonietta Bisetto, editors, *Universals of Language Today*, volume 76 of *Studies in Natural Language and Linguistic Theory*, pages 185–207. Springer.
- Matthew S. Dryer and Martin Haspelmath. 2013. [WALS online](#). *Leipzig: Max Planck Institute for Evolutionary Anthropology*. <http://wals.info/>.
- Michael Dunn, Simon J Greenhill, Stephen C Levinson, and Russell D Gray. 2011. [Evolved structure of language shows lineage-specific trends in word-order universals](#). *Nature*, 473(7345):79–82.
- Murray Emeneau. 1956. India as a linguistic area. *Language*, 32(1):3–16.
- Nicholas Evans and Stephen C. Levinson. 2009. [The myth of language universals: Language diversity and its importance for cognitive science](#). *Behavioral and Brain Sciences*, 32(5):429–448.
- Joel N Franklin. 1974. [On Tikhonov’s method for ill-posed problems](#). *Mathematics of Computation*, 28(128):889–907.
- Joseph H. Greenberg. 1963. [Some universals of grammar with particular reference to the order of meaningful elements](#). In Joseph H. Greenberg, editor, *Universals of Human Language*, pages 73–113. MIT Press, Cambridge, Mass.
- Alexander Gutkin and Richard Sproat. 2017. [Areal and phylogenetic features for multilingual speech synthesis](#). In *Proc. Interspeech 2017*, pages 2078–2082, Stockholm, Sweden. International Speech and Communication Association (ISCA).
- Haibo He, Yang Bai, Edwardo A. Garcia, and Shutao Li. 2008. [ADASYN: Adaptive synthetic sampling approach for imbalanced learning](#). In *Proc. of International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 1322–1328. IEEE.

- Arthur E Hoerl and Robert W Kennard. 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.
- Martin Josifoski, Ivan S. Paskov, Hristo S. Paskov, Martin Jaggi, and Robert West. 2019. Crosslingual document embedding as reduced-rank ridge regression. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 744–752, Melbourne, Australia.
- Julie Josse, Marie Chavent, Benot Liqueur, and François Husson. 2012. Handling missing values with regularized iterative multiple correspondence analysis. *Journal of Classification*, 29(1):91–116.
- Julie Josse and François Husson. 2012. Handling missing values in exploratory multivariate data analysis methods. *Journal de la Société Française de Statistique*, 153(2):79–99.
- Henri Kauhanen, Deepthi Gopal, Tobias Galla, and Ricardo Bermúdez-Otero. 2019. Geospatial distributions reflect rates of evolution of features of language. In *Proceedings of Workshop on Spatial Patterns of Language Evolution (SPLEV)*, Zurich, Switzerland.
- Ranjit Lall and Thomas Robinson. 2020. Applying the MIDAS touch: How to handle missing values in large and complex data. *American Political Science Association Preprints*. Code available at: <https://github.com/ranjitlall/MIDAS>.
- Pat Langley, Wayne Iba, Kevin Thompson, et al. 1992. An analysis of Bayesian classifiers. In *Proc. 10th Nat. Conf. on Artificial Intelligence*, volume 90, pages 223–228, USA. AAAI Press and MIT Press.
- Yugo Murawaki. 2015. Continuous space representations of linguistic typology and their application to phylogenetic inference. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 324–334, Denver, Colorado. Association for Computational Linguistics.
- Yugo Murawaki. 2017. Diachrony-aware induction of binary latent representations from typological features. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 451–461, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Yugo Murawaki. 2019. Bayesian learning of latent representations of language structures. *Computational Linguistics*, 45(2):199–228.
- Yugo Murawaki and Kenji Yamauchi. 2018. A statistical model for the joint inference of vertical stability and horizontal diffusibility of typological features. *Journal of Language Evolution*, 3(1):13–25.
- Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 629–637, Jeju Island, Korea. Association for Computational Linguistics.
- John Nerbonne, Jelena Prokić, and Martijn Wieling. 2020. Dialectology for computational linguists. In Marcos Zampieri and Preslav Nakov, editors, *Similar Languages, Varieties, and Dialects: A Computational Perspective*. Cambridge University Press. To appear.
- Yohei Ono. 2020. How to handle “missing values” in linguistic typology: A pitfall in the statistical modelling approach. *Northern Language Studies*, 10:61–82.
- Edoardo Maria Ponti, Helen O’Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, Thierry Poibeau, Ekaterina Shutova, and Anna Korhonen. 2019. Modeling language variation and universals: A survey on typological linguistics for natural language processing. *Computational Linguistics*, 45(3):559–601.
- Paolo Ramat. 2011. Linguistic typology. In Georg Bossong and Bernard Comrie, editors, *Empirical Approaches to Language Typology*, volume 1. Walter de Gruyter.
- Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- C. Carl Robusto. 1957. The Cosine-Haversine formula. *The American Mathematical Monthly*, 64(1):38–40.
- Johan AK Suykens and Joos Vandewalle. 1999. Least squares support vector machine classifiers. *Neural Processing Letters*, 9(3):293–300.
- Benedikt Szmrecsanyi. 2011. Corpus-based dialectometry: a methodological sketch. *Corpora*, 6(1):45–76.
- Hiroya Takamura, Ryo Nagata, and Yoshifumi Kawasaki. 2016. Discriminative analysis of linguistic features for typological study. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 69–76, Portorož, Slovenia. European Language Resources Association (ELRA).
- Alaa Tharwat. 2016. Linear vs. quadratic discriminant analysis classifier: A tutorial. *International Journal of Applied Pattern Recognition*, 3(2):145–180.
- Fabian Tomaschek, Peter Hendrix, and R. Harald Baayen. 2018. Strategies for addressing collinearity in multivariate linguistic data. *Journal of Phonetics*, 71:249–267.

Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. 2010. [Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion](#). *Journal of Machine Learning Research*, 11(12).

Martijn Wieling and John Nerbonne. 2015. [Advances in dialectometry](#). *Annual Review of Linguistics*, 1:243–264.

Ian H. Witten, Eibe Frank, and Mark A. Hall. 2011. *Data Mining: Practical Machine Learning Tools and Techniques*, 3rd edition. Morgan Kaufmann, San Francisco.

Ji Zhu, Hui Zou, Saharon Rosset, and Trevor Hastie. 2009. [Multi-class AdaBoost](#). *Statistics and its Interface*, 2(3):349–360.