



Exploring language relations through syntactic distances and geographic proximity

Juan De Gregorio^{1*} , Raúl Toral¹ and David Sánchez¹

*Correspondence:

juan@ifisc.uib-csic.es

¹Institute for Cross-Disciplinary Physics and Complex Systems IFISC (UIB-CSIC), Campus Universitat de les Illes Balears, Palma de Mallorca, E-07122, Spain

Abstract

Languages are grouped into families that share common linguistic traits. While this approach has been successful in understanding genetic relations between diverse languages, more analyses are needed to accurately quantify their relatedness, especially in less studied linguistic levels such as syntax. Here, we explore linguistic distances using series of parts of speech (POS) extracted from the Universal Dependencies dataset. Within an information-theoretic framework, we show that employing POS trigrams maximizes the possibility of capturing syntactic variations while being at the same time compatible with the amount of available data. Linguistic connections are then established by assessing pairwise distances based on the POS distributions. Intriguingly, our analysis reveals definite clusters that correspond to well known language families and groups, with exceptions explained by distinct morphological typologies. Furthermore, we obtain a significant correlation between language similarity and geographic distance, which underscores the influence of spatial proximity on language kinships.

Keywords: Language distances; Syntactic variation; POS dataset; POS n-grams; High-order Markov processes; Linguistic typology

1 Introduction

The number of languages in the world is estimated to be around 7,000 [1]. This leads to a broad diversity at all linguistic levels: phonetic, morphosyntactic, semantic and pragmatic. The task to comprehend this immense variation is overwhelming. However, researchers have managed to pinpoint linguistic relationships that allow them to cluster languages in groups and families. Such classification was first based on comparative studies [2, 3] and now has increasingly been supported with quantitative approaches [4, 5].

A fertile approximation in historical linguistics describes languages as species in a phylogenetic tree that shows the evolutionary history from proto-languages to today's descendants [6]. In this diachronic view, languages change through linguistic innovations that cause two languages with the same ancestor to become mutually unintelligible. There is also the complementary view, followed in this work, that searches for relations among languages in a given moment of time, in a synchronic manner [7], with the goal of quantifying the distance between languages using an appropriate metric [8, 9].

© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Language distances are encoded in matrices whose entries measure the similarity among certain linguistic features. The same method has been successfully applied in dialectometry, which aims at quantifying the regional differences among varieties of a given language [10]. Quantitative measures of linguistic distances are useful not only for fundamental reasons but also in applied linguistics with the aim of analyzing the learning difficulties of minorities and immigrants [11]. Another application of language distances is for languages in contact since languages that are more congruent to each other are more likely to coexist [12]. An assessment of linguistic similarities is therefore helpful for evidence-based policies and planning that seek to revitalize endangered languages [13].

Now, whilst most of the studies that compute distances either in synchronic linguistics or in dialectometry focus on orthographic, phonetic or lexical variations [14–19], less attention has been paid to morphosyntactic features, other than a few exceptions [20–22]. The latter are interesting because syntax is more robust to change than phonetics or semantics. Therefore, the resulting picture would show a larger time depth, and unique cases of accelerated change would hence stand out.

A particularly simple but elucidating approach to analyze syntactic variations is by means of parts of speech (POS). These denote word classes with well defined grammatical functions that share common morphological properties [23]. For instance, almost all languages distinguish between verbs and nouns, i.e., roughly between actions and entities. As a consequence, one can categorize words or lexical items as members of any of the proposed POS, typically around 15. This classification has its own limitations (e.g., certain languages do not distinguish between verbs and adjectives) but it has the advantage of simplicity while capturing at the same time a large amount of morphosyntactic information. The POS approach has been proven especially useful in natural language processing tasks. The reason is that POS reveal much not only about the syntactic category of a word but also about that of its neighboring words, due to semantic restrictions. For example, if a word is a noun it will most likely be surrounded by determiners and adjectives, forming a noun phrase. Therefore, we can gain insight about the phrasal structure of a language by examining POS sequences. This is precisely the main objective of this work: to model these sequences as stochastic processes, analyze their correlations and compute syntactic distances between languages using POS sequence distributions taken from a multilingual corpus.

Statistics of POS sequences, specifically the analysis of POS r -grams, defined as sequences of r successive POS, has served for diverse research goals. One may hypothesize that genres are characterized by different syntactic structures, which would then allow for reliable genre classification. As demonstrated in Ref. [24], a careful study of POS trigram histograms provides a high-performance genre classifier. Strikingly, series of POS trigrams can be employed for building phylogenetic language trees solely from translations [25]. The premise here is that syntactic features are retained in the translation process. Furthermore, assuming that POS tags can be predicted for historically close languages it is possible to train a language model to measure proximity among languages [26]. However, these previous works set the POS block length in a somewhat heuristic manner. Below, we demonstrate using information-theoretic methods that trigrams suffice to account for the correlations present in POS sequences. In other words, it is not necessary to consider r -grams with $r \geq 4$ to gain more information, a result that considerably simplifies analyses that involve POS series.

The mapping between a corpus and its corresponding POS series can be performed with human or automatic parsers. For the purposes of our investigation, we consider the Universal Dependencies library dataset [27], which includes manually annotated tree-banks across many languages [28]. This dataset consists of texts and speech transcripts originated from various sources: news, online social media, legal documents, parliament speeches, literature, etc. The 17 universal parts of speech used in this particular dataset are grouped into three classes: open, which includes the tags adjective (ADJ), adverb (ADV), interjection (INTJ), noun (NOUN), proper noun (PROPN) and verb (VERB); closed, with the tags adposition (ADP), auxiliar (AUX), coordinating conjunction (CCONJ), determiner (DET), numeral (NUM), particle (PART), pronoun (PRON) and subordinating conjunction (SCONJ); and others, which comprise punctuation (PUNCT), symbol (SYM) and other (X).

We then build a corpus of 67 contemporary languages expressed by means of these tags. Since the number of possible POS r -grams grows exponentially with r , it is thus natural to ask what value of r conveys the maximum information about a language. As aforementioned, we find that $r = 3$ suffices to correctly characterize any of the studied languages. Then, we depict the connections between languages assessing the pairwise distance between POS trigram distributions. Interestingly, our found clusters can be identified with well known families and groups. Exceptions can be understood due to distinct linguistic typologies. This is natural since morphology constrains the possible POS combinations that can form and, consequently, this is reflected in the POS distributions and the distances calculated thereof.

Interestingly, we find a correlation between the obtained linguistic distance and the geographic distance spanned between locations assigned to each language. These are WALS locations, which generally correspond to the geographic coordinates associated to the centre of the region where the analyzed languages are spoken. However, for some languages, the regions where they are spoken are discontinuous. In such cases, the locations are placed within the larger region in which the language is spoken [29]. Despite the fact that the centres are obviously only approximate and that our calculation of the linguistic distances has its own limitations, we clearly find that most of the syntactically close languages are also geographically close. This is expected since similar language usually lie in a continuum but there exist conspicuous exceptions, as we shall discuss below.

2 Methods

2.1 Data

The data utilized in this work is taken directly from the Universal Dependencies library [27], where each language is depicted with one or several corpora that are manually tagged. Thus, each word is classified into one of the possible POS categories as stated above. These tags are deemed universal because they provide a common and consistent way to represent the grammatical categories of words across different languages. For example, the English sentence

Launching this way will hopefully avoid future disasters,
giving more support towards NASA revisiting the stars. (1)

is converted into the POS sequence

$$S = \text{VERB, DET, NOUN, AUX, ADV, VERB, ADJ, NOUN, PUNCT, VERB, ADJ,} \\ \text{NOUN, CONJ, PROP, VERB, DET, NOUN, PUNCT} \quad (2)$$

Since we are interested in lexical classes that are either open or closed, we combine the 3 categories in the others class in a single tag. Hence, we will only consider $L = 15$ distinct categories. As a consequence, in alphabetical order, the possible tags are $\{z_i\}_{i=0}^{14} = \{\text{ADJ, ADP, ADV, AUX, CCONJ, DET, INTJ, NOUN, NUM, PART, PRON, PROP, PUNCT, SCONJ, VERB}\}$, where the categories in the others class are included in the tag PUNCT. We can then count the occurrences of each tag in S , excluding the period at the end of the sentence. For example, the number of times the tag $z_0 = \text{ADJ}$ occurs in Eq. (2) is 2. This way we gain access to unigram statistics. Similarly, we can group S in overlapping blocks of 2 consecutive tags and count the number of times each block is observed. For example, the block $(z_5, z_7) = (\text{DET, NOUN})$ occurs in Eq. (2) 2 times, thus opening the path to bigram statistics, etc.

In general, the set of all L^r possible r -grams, or blocks of size $r \geq 1$, is given by $\{b_j^{(r)} : b_{(i_0, \dots, i_{r-1})_L}^{(r)} = (z_{i_0}, \dots, z_{i_{r-1}}), \quad i_0, \dots, i_{r-1} = 0, \dots, L-1\}_{j=0}^{L^r-1}$, where $(i_0, \dots, i_{r-1})_L$ is a base L number. Given the dataset of language \mathcal{L} , formed by R tagged sentences, we count the number of appearances for each block of size r , with $1 \leq r \leq r_{\max}$, with r_{\max} to be specified below. First, we arrange each of the R sequences in overlapping blocks of size r and calculate the occurrences $\hat{n}_j^{(r)}$ of block $b_j^{(r)}$. With this information we build for each value of r the set of observations $\{\hat{n}_j^{(r)}\}_{j=0}^{L^r-1}$ for every language \mathcal{L} .

\mathcal{L} can be any of the 67 contemporary European and Asian languages that fulfill the criterion of having datasets of at least 10 thousand tokens in the Universal Dependencies library. The comprehensive list is included in App. A, along with their respective language family and group, as well as their corresponding morphological type (agglutinative, fusional, isolating). The latter is a useful information about word formation. We emphasize that these three types are just approximate categories since most of the languages have morphological traits of the three types with variable relevance [30].

Importantly, we select languages within a contiguous region (except for Afrikaans which is included since it belongs to the Indo-European family) because we shall later explore possible correlations between linguistic and geographic distances. Specifically, we focus on Eurasia since this is a single continent with both a rich linguistic diversity and abundant data availability.

Subsequently, we apply the previously discussed procedure of counting POS r -grams occurrences for every language within our dataset. Following an information-theoretic approach, we now argue in the next section that the r -gram probability distribution of each language for $r = 3$ is sufficient to capture the correlations observed in POS sequences and, consequently, it can serve as a reliable basis for calculating distances between languages.

2.2 Predictability gain and memory

The dynamics of many stochastic processes can be described by considering that the transition probabilities to future outcomes depend on previous states. Consider a random variable X with L possible outcomes z_0, \dots, z_{L-1} and probability distribution $P(X) = \{P(X = z_i), \quad i = 0, \dots, L-1\}$, where $P(X = z_i)$ is the probability that X takes the value z_i . In our

case, X are the POS tags taking values specified above, and $P(X = z_i)$ is the probability for occurrence of tag z_i in language \mathcal{L} . Given $k + 1$ repetitions of X , X_0, \dots, X_k , the accuracy of predicting the next outcome of the process, X_{k+1} , generally grows as the number of past states considered increases. For example, the information contained in the stationary probability (or zeroth-order transition probability) $P(X_{k+1} = z_j)$ for outcome z_j is always less or equally accurate than the information provided by the first-order transition probability $P(X_{k+1} = z_j | X_k = z_l)$.

The predictability gain \mathcal{G}_u [31] quantifies the amount of information that one gains when performing predictions taking into account $u + 1$ past states instead of u . For completeness, we prove in App. B that, for homogeneous systems in which the transition probabilities are independent of time, the predictability gain takes the form

$$\mathcal{G}_u = -(H_{u+2} - 2H_{u+1} + H_u), \quad u \geq 0, \quad (3)$$

where H_r is the block Shannon entropy [32] of size $r \geq 1$ ($H_0 \equiv 0$). H_r for $r \geq 1$ can be straightforwardly calculated from the joint probability distribution $P_r(X)$ of r consecutive repetitions of the variable X ,

$$P_r(X) \equiv \{P(X_0 = z_{i_0}, \dots, X_{r-1} = z_{i_{r-1}}), \quad i_k = 0, \dots, L - 1\}, \quad (4)$$

as

$$H_r = - \sum_{j=0}^{L^r-1} p(b_j^{(r)}) \log(p(b_j^{(r)})), \quad (5)$$

where

$$p(b_{(i_0, \dots, i_{r-1})L}^{(r)}) = P(X_0 = z_{i_0}, \dots, X_{r-1} = z_{i_{r-1}}), \quad 0 \leq i_0, \dots, i_{r-1} \leq L - 1, \quad (6)$$

is the probability for occurrence of block $b_{(i_0, \dots, i_{r-1})L}^{(r)} = (z_{i_0}, \dots, z_{i_{r-1}})$ with size r , and \log is hereafter understood as \log_2 .

A stochastic process generated from consecutive repetitions of the variable X has order or memory $m \geq 1$ if the transition probabilities satisfy

$$P(X_{k+1} = z_j | X_0 = z_{i_0}, \dots, X_k = z_{i_k}) = P(X_{k+1} = z_j | X_{k-(m-1)} = z_{i_0}, \dots, X_k = z_{i_k}). \quad (7)$$

These are usually referred to as m -th order Markovian processes [33]. For the case $m = 0$ the probabilities $P(X_s = z_j)$ are independent for all s .

In Ref. [34] it was shown that a system has memory m if and only if H_u is a linear function of u for $u \geq m$. This result amounts to stating that a system has memory m if and only if $\mathcal{G}_u = 0$ for all $u \geq m$. This can be proven directly from Eq. (3). Hence, an analogous definition for the memory m of a stochastic system follows:

$$m = \min(\{\mu : \mathcal{G}_u = 0, \text{ for all } u \geq \mu\}). \quad (8)$$

Therefore, m is the minimum number of past states that we need to consider in order to achieve maximum predictability of the random process. We can use Eq. (8) as a benchmark

to determine the memory of a system. Moreover, if a stochastic process described with the random variable X has memory m , knowing the joint probability distribution $P_{m+1}(X)$ of $m + 1$ consecutive repetitions of X is sufficient to capture all relevant information about the process since P_{m+1} can be used to compute the probabilities of larger and smaller blocks. Hence, the $m + 1$ block size is optimal for accurate predictions and understanding the process dynamics, without adding redundant information.

2.3 Estimation from finite data

Given the counts $\{\hat{n}^{(r)}\} \equiv \{\hat{n}_j^{(r)}\}_{j=0}^{L^r-1}$ of language \mathcal{L} obtained following the procedure outlined in Sect. 2.1 for a value of r , we estimate the probability to observe block $b_j^{(r)}$ of r consecutive POS tags in language \mathcal{L} as

$$\hat{p}_{\mathcal{L}}(b_j^{(r)}) = \frac{\hat{n}_j^{(r)}}{N^{(r)}}, \quad (9)$$

where $N^{(r)} = \sum_{j=0}^{L^r-1} \hat{n}_j^{(r)}$. It is well known that Eq. (9) is an unbiased estimator for the probability $p_{\mathcal{L}}(b_j^{(r)})$. Similarly, from our observations we can estimate the $(r - 1)$ th-order transition probabilities as

$$\hat{p}_{\mathcal{L}}(z_{i_r}|z_{i_1}, \dots, z_{i_{r-1}}) \equiv \hat{P}(X_r = z_{i_r} | X_1 = z_{i_1}, \dots, X_{r-1} = z_{i_{r-1}}) = \frac{\hat{n}_{(i_1, \dots, i_{r-1}, i_r)_L}^{(r)}}{\sum_{\nu=0}^{L-1} \hat{n}_{(i_1, \dots, i_{r-1}, \nu)_L}^{(r)}}, \quad (10)$$

with $i_k = 0, \dots, L - 1$.

Our next step involves calculating the predictability gain in order to determine the memory of the POS sequences by means of Eq. (8). Motivated by Eq. (3) we can perform this task by estimating the block entropies.

Entropy estimation is a largely analyzed problem and, even though there is not known unbiased estimator for the entropy [35], numerous useful estimators can be found in the literature [36]. Usually, all entropy estimators fail when the number of possible outcomes is larger than the available data. Considering that the number of possible blocks grows exponentially with increasing block size, there exists a certain r_{\max} for which entropy estimation is unreliable for $r > r_{\max}$. The value of r_{\max} depends on the chosen entropy estimator, but it is usual to take $r_{\max} \simeq \log(N^{(1)})/\log(L)$ [34].

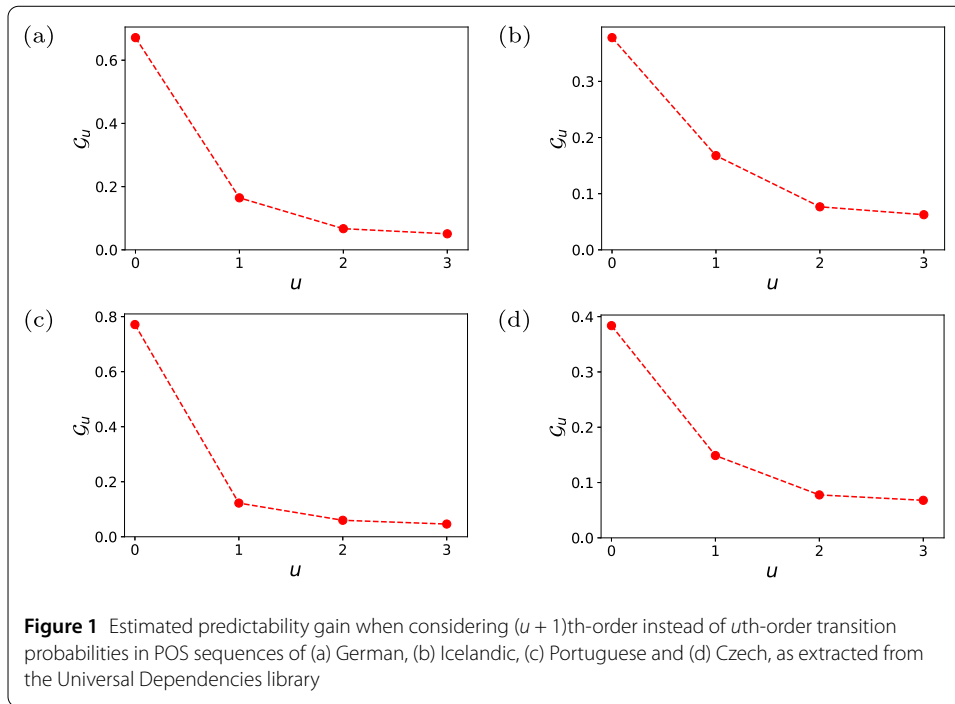
Hereafter we use the NSB entropy estimator [37, 38], which has shown to give good results for correlated sequences [39]. Estimation of the block entropy of size r using this method only requires knowledge of the data observations $\{\hat{n}^{(r)}\}$. The result of applying the estimator \hat{H} to this set is denoted with $\hat{H}[\{\hat{n}^{(r)}\}]$, for $1 \leq r \leq r_{\max}$. Therefore, for finite data Eq. (3) becomes

$$\hat{G}_u[\{\hat{n}\}] = - \left(\hat{H}[\{\hat{n}^{(u+2)}\}] - 2\hat{H}[\{\hat{n}^{(u+1)}\}] + \hat{H}[\{\hat{n}^{(u)}\}] \right), \quad 1 \leq u \leq r_{\max} - 2, \quad (11)$$

and

$$\hat{G}_0[\{\hat{n}\}] = - \left(\hat{H}[\{\hat{n}^{(2)}\}] - 2\hat{H}[\{\hat{n}^{(1)}\}] \right). \quad (12)$$

Due to the limitations of entropy estimation, the condition imposed in Eq. (8) to determine the value of m is too strict. Rather than attempting to ascertain the minimum



block size at which the predictability gain is 0, a more sensible approach is to compare the values of $\hat{\mathcal{G}}_u[\{\hat{n}\}]$ with those obtained from a scenario where we know that the estimated predictability gain should be 0.

2.4 POS trigrams

In this section we provide evidence supporting our claim that in order to capture the correlations in POS sequences it is enough to consider their trigram probability distribution. We show this in two ways: in Sect. 2.4.1, calculating the predictability gain, and in Sect. 2.4.2, analyzing the accuracy in language detection as the order of Markovian models increases.

2.4.1 Predictability gain

Our first objective is to quantify the information gained in POS sequences of various languages when performing predictions considering $(u + 1)$ th-order transition probabilities rather than u th-order.

In Fig. 1 we show $\hat{\mathcal{G}}_u$ for two representative languages of the same group, namely, German (panel a) and Icelandic (panel b) from the Germanic group, and two representative languages of different groups that also differ from the previous one, namely, Portuguese (panel c) from the Romance group and Czech (panel d) from the Slavic group. This way we can test intragroup and intergroup variations, if they exist. Further, these three groups constitute the most extensive clusters, characterized by a wealth of available data, which make the results more trustworthy. We set $r_{\max} = 5$ for the four languages; hence we ensure that the estimation of their predictabilities is reliable up to $u = 3$.

The values of $\hat{\mathcal{G}}_u$ at $u = 0$ for the four languages demonstrate a significant predictability gain when transitioning from zeroth-order predictions to first-order predictions, suggesting that the POS state of the sequence at a given step k is highly informative about the next POS outcome at step $k + 1$. We observe that the values of $\hat{\mathcal{G}}_0$ significantly vary for

each language, even for German and Icelandic which belong to the same group. Then, the $\hat{\mathcal{G}}_u$ values at $u = 1$ indicate that the information provided by the preceding POS outcome at step $k - 1$ is also substantial in terms of predictability. In contrast, the predictability gain at $u = 2$ drops to $\hat{\mathcal{G}}_2 \simeq \frac{\hat{\mathcal{G}}_1}{2}$ and remains relatively constant at $u = 3$. We obtain this decreasing pattern in the curves of $\hat{\mathcal{G}}_u$ as a function of u for all languages considered in this work. Even though this does not prove that POS sequences have memory 2 (see App. C for further details), our results indeed show that considering transition probabilities beyond order 2 does not yield substantially more information about the correlations present in our POS sequences.

2.4.2 Markov models

We now provide more evidence that POS trigrams (or equivalently, POS sequences of memory 2) describe most of the statistical information contained in the represented languages.

From the corpus of language \mathcal{L} we extract a tagged sentence $S = x_1, \dots, x_N$ of length N and compute the probability $P(S|\mathcal{L}')$ of observing S in a certain language \mathcal{L}' . This probability can be computed either by considering the estimated stationary distribution of language \mathcal{L}' ,

$$\hat{P}^{(0)}(S|\mathcal{L}') = \prod_{k=1}^N \hat{p}_{\mathcal{L}'}(x_k), \quad (13)$$

or by incorporating u th-order transition probabilities for $u \geq 1$ (u -order Markov model). For example, the probability of observing the sequence S given the language \mathcal{L}' considering first-order transition probabilities reads


$$\hat{P}^{(1)}(S|\mathcal{L}') = \hat{p}_{\mathcal{L}'}(x_1) \prod_{k=1}^{N-1} \hat{p}_{\mathcal{L}'}(x_{k+1}|x_k). \quad (14)$$

Similarly, we can compute $\hat{P}^{(u)}(S|\mathcal{L}')$ for $u > 1$. This procedure is repeated for various languages, and the language that yields the highest probability is assigned to the one from which S is generated. Ideally, the obtained language would correspond to \mathcal{L} for all values of u .

For each of the four previously considered languages (German, Icelandic, Portuguese, and Czech) we randomly select from their respective corpora K sentences S_1, \dots, S_K , each comprising 5 to 20 word tokens. Subsequently, for every sentence we compute $\hat{P}^{(u)}(S_l|\mathcal{L}')$ for the four languages and for values of u ranging from 0 to 3. The accuracy A_u of correctly identifying the language associated with each case is then assessed as the fraction of correct sentence classifications:

$$A_u(\mathcal{L}) = \frac{\#\{l : \hat{P}^{(u)}(S_l|\mathcal{L}) > \hat{P}^{(u)}(S_l|\mathcal{L}') \text{ for all } \mathcal{L}' \neq \mathcal{L}\}_{l=1}^K}{K}. \quad (15)$$

We compute the estimated stationary and transition probabilities, as specified in Eq. (9) and (10) respectively, without considering the influence of the K sentences utilized for testing purposes. We then repeat this procedure 10 times to find the mean and standard



deviation of the accuracy for each language, as a function of u . We consider $K = 1000$. The results are displayed in Fig. 2.

We observe that the accuracy for all languages significantly increases when first-order transition probabilities are considered ($u = 1$). Subsequently, there is a slight increase at $u = 2$, and afterwards the accuracy remains relatively constant for $u = 3$. These findings agree with the analysis performed for the predictability gain, overall indicating that considering memory values higher than 2 does not provide significant additional information. Conversely, choosing a low memory value proves advantageous as it enhances the probability estimation accuracy. Consequently, this enables us to incorporate a broader set of languages into our analysis.

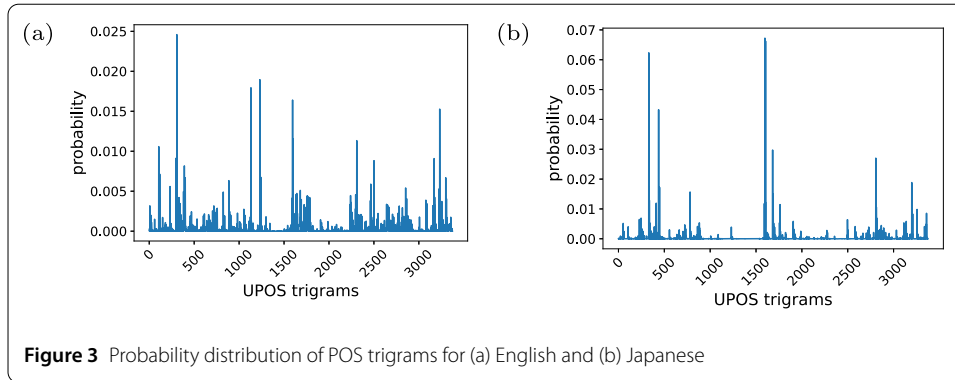
2.5 Language distances

As explained in Sect. 2.2, if a stochastic process generated from repetitions of the variable X has memory m , the probability distribution $P_{m+1}(X)$ is adequate for capturing all relevant information about the process. We previously showed that sequences of POS tags can be modeled as processes with memory 2 with high accuracy. Therefore, we can define a distance metric between languages \mathcal{L} and \mathcal{L}' from the statistical distance between their corresponding trigram distributions. To this end, we consider the Jensen-Shannon (JS) distance [40]. Defining $\{b_j\}_{j=0}^{L^3-1}$ as the set of all possible POS trigrams, the JS distance between languages \mathcal{L} and \mathcal{L}' is determined by

$$d_{JS}(\mathcal{L}, \mathcal{L}') = \sqrt{\frac{1}{2} \sum_{j=0}^{L^3-1} \left(p_{\mathcal{L}}(b_j) \log \left(\frac{2p_{\mathcal{L}}(b_j)}{p_{\mathcal{L}}(b_j) + p_{\mathcal{L}'}(b_j)} \right) + p_{\mathcal{L}'}(b_j) \log \left(\frac{2p_{\mathcal{L}'}(b_j)}{p_{\mathcal{L}}(b_j) + p_{\mathcal{L}'}(b_j)} \right) \right)}. \quad (16)$$

It is worth noting that the d_{JS} measure satisfies all the essential properties expected for a metric and that d_{JS} ranges between 0 and 1.

We estimate the JS distance by replacing the exact probabilities $p(b_j)$ in Eq. (16) with the maximum likelihood estimators given by Eq. (9). As an illustration, we present in Fig. 3 the trigram probability distributions for English (panel a) and Japanese (panel b). The numbers that designate the trigrams correspond to their index. For example, trigram 0 refers to the



block $b_0 = (z_0, z_0, z_0) = (\text{ADJ}, \text{ADJ}, \text{ADJ})$; trigram 1 is $b_1 = (z_0, z_0, z_1) = (\text{ADJ}, \text{ADJ}, \text{ADP})$; and so on until the last trigram 3374, which corresponds to the block $b_{3374} = (z_{14}, z_{14}, z_{14}) = (\text{VERB}, \text{VERB}, \text{VERB})$.

From a simple inspection of Fig. 3 it is clear that the trigram probability distributions of English (\mathcal{E}) and Japanese (\mathcal{J}) show substantial differences. For example, in English the three most probable trigrams are $b_{307} = (\text{ADP}, \text{DET}, \text{NOUN})$, $b_{1231} = (\text{DET}, \text{NOUN}, \text{ADP})$ and $b_{1132} = (\text{DET}, \text{ADJ}, \text{NOUN})$ whereas for Japanese we have $b_{1597} = (\text{NOUN}, \text{ADP}, \text{NOUN})$, $b_{1604} = (\text{NOUN}, \text{ADP}, \text{VERB})$ and $b_{331} = (\text{ADP}, \text{NOUN}, \text{ADP})$. These trigrams differ because determiners are generally absent from Japanese, unlike English, and adpositions follow the Japanese nouns whereas in English adpositions can also appear before the noun. These differences between the two languages can be quantified using Eq. (16). We find $d_{JS}(\mathcal{E}, \mathcal{J}) = 0.79$, which is a high value due to the strong morphosyntactic differences between Japanese and English.

Following a similar methodology, for each language pair within our dataset we compute their corresponding JS distances using Eqs. (16), yielding a distance matrix of dimension 67×67 , on which our subsequent clustering analysis is based. This methodology considers all available data for each language. An alternative approach is presented in App. D, where we calculate distances among text samples extracted from the same language group. We now present the results obtained with the JS distance. Importantly, we also obtain similar results employing a different metric (the Hellinger distance, see App. E), which reinforces the validity of our findings.


3 Results

We first discuss the general results obtained from calculating distances between languages based on their POS distributions. Then, we inquire a possible correlation between linguistic and physical (i. e., geographical) distances.

3.1 Language distances and cluster analysis

3.1.1 Distance matrix

The distance matrix generated from the data with the aid of Eq. (16) can be better visualized through a clustermap (see Fig. 4). This representation employs hierarchical clustering [41] and heatmap visualization. We use the complete linkage method [42] for clustering, organizing rows and columns based on similarity, depicted as dendrograms in the same figure. The heatmap, which represents distance values by a color spectrum as indi-



cated in Fig. 4, enables a comprehensive exploration of our data relationships and structure.

First, since darker colours indicate shorter distance (i.e., greater morphosyntactic similarity) one can clearly distinguish large-scale cluster formation. In the vicinity of the left upper corner we observe the most extensive cluster, corresponding to Slavic languages. Within this language group, we discern smaller clusters, the most clear being among Belarusian, Russian and Ukrainian (the East branch of the Slavic family), as well as between Serbian and Croatian (the South branch). Interestingly, our results point to a close relationship between this group and two members of the Baltic group (Latvian and Lithuanian). This suggests that spatial proximity might be correlated with our POS distances, as we discuss below.

Germanic languages exhibit a more dispersed pattern, manifested itself in two distinct clusters. The first cluster encompasses Afrikaans, Dutch, and German (the West branch), while the second is primarily composed of North Germanic languages along with English,

which highlights its mixed character via a close connection with Romance languages. The latter, in turn, exhibit the most compact clustering, as evidenced by their tight proximity to one another in the heatmap. An exception is Romanian, which falls outside this group, probably because it is the only major Romance language with noun declension and article enclitics. Also, its spatial distance from the West may play a role. This interpretation is supported by the fact that Romanian is clustered with Pomak, a Bulgarian variety spoken in the geographically close region of Thrace.

We next find a small cluster that encompasses Maltese and Hebrew. Surprisingly, Arabic, classified as a Semitic language as well, appears closer to Austronesian languages, such as Indonesian and Javanese, as well as Persian. The reason lies in the different typology of Arabic and its strong dialect diversity. Further visible clusters correspond to Celtic languages, Hindi and Urdu (both belonging to the Indic language group), the Turkic family (Turkish, Kazakh and Uyghur), adjacent to Buryat (a Mongolic language), followed by a larger cluster primarily comprising Uralic, Tungusic, and Sino-Tibetan families. This is the most diverse cluster partly because the Universal Dependencies library contains data of a few languages only as compared with the previous families. However, there exist interesting connections that we explore in the next sections.

3.1.2 Language tree and *k*-medoids clustering

In order to gain further insight on the relations among the distinct languages, we perform a *k*-medoids clustering upon the distance matrix. Hence, we utilize the Partition Around Medoids (PAM) algorithm [43], for which the optimal number of clusters is determined through a silhouette analysis [44]. The corresponding figure and a list of languages constituting each cluster are presented in App. F. For the moment, we will make use of these results to build a single picture of the main interlinguistic connections.

First, we construct a network whose nodes represent languages whereas the edges' weight are given by their pairwise distances. Only the edges that generate a minimum spanning tree are considered as part of our analysis. The minimum spanning tree is a structure that connects all the data points in our dataset results with the minimum possible total edge weight [45]. We depict the tree in Fig. 5, employing the Kamada-Kawai layout [46]. This is a force-directed layout, implying that, for weighted graphs as in our case, the length of the edges tends to be larger for heavier weights. This visualization provides an intuitive representation of the relationships encoded in the minimum spanning tree, offering insights into the overall structure and connectivity patterns, not only among closely related languages, but also between distinct language families. To enhance our interpretability, we assign consistent colors to items within the same cluster, obtained from the *k*-medoids analysis mentioned earlier. Nodes from the same language group are linked with a full line; nodes from the same family but different group are linked with a dashed line; finally, nodes belonging to distinct families are connected with dotted lines. The shape of the nodes represents the language type: circles are assigned to fusional languages, squares to the agglutinative type and diamonds to isolating languages. The languages plotted with half nodes are both isolating and another type. For example, English is plotted with a half diamond node because it is considered to be a fusional-isolating language [47].

Interestingly, the same cluster formations observed with our previous analysis of Fig. 4 naturally emerge in the constructed tree, as well as connections between language families hidden within the distance matrix, which are easier to identify in this simpler visualization. We can find most of the Celtic languages (clusters 6, 28) close to a few Semitic

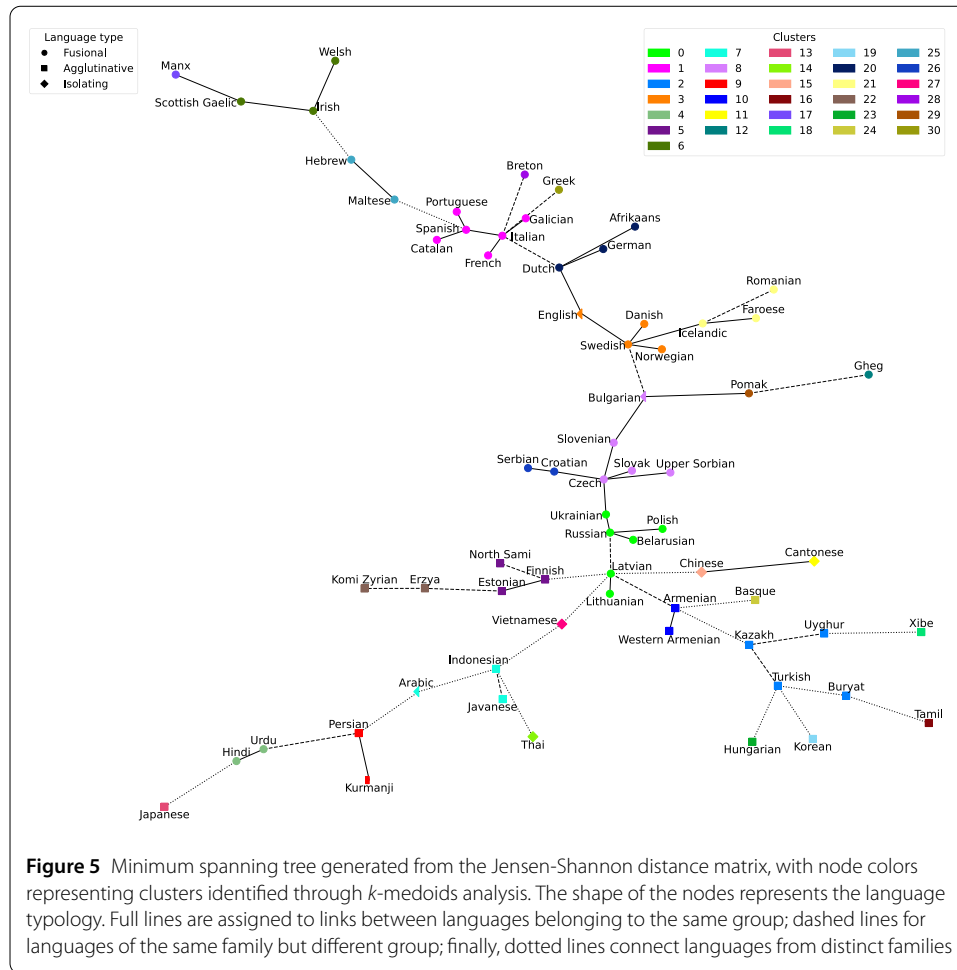


Figure 5 Minimum spanning tree generated from the Jensen-Shannon distance matrix, with node colors representing clusters identified through *k*-medoids analysis. The shape of the nodes represents the language typology. Full lines are assigned to links between languages belonging to the same group; dashed lines for languages of the same family but different group; finally, dotted lines connect languages from distinct families

languages (cluster 25), specially Hebrew, whereas Maltese can be found in close proximity to Romance languages, which are once again forming a compact cluster (1), with the exception of Romanian which is clustered with Icelandic and Faroese (cluster 21). The Germanic group is subdivided into three groups (clusters 3, 20 and 21) that are relatively close to one another, and similarly with the Slavic group (clusters 0, 8, 26, 30). The Baltic languages are clustered together with the Slavic sub-group mostly consisting in East Slavic languages (cluster 0). The proximity of these two groups was also observed in Fig. 4.

The Baltic languages are connected with the Uralic languages (clusters 5, 22) and with the Armenian family (cluster 10). An exception is Hungarian, considered an Uralic language, but here appearing separately (cluster 23) near the Turkic languages, which are grouped together as expected, alongside Buryat (cluster 2). Around this cluster several isolated languages such as Korean (cluster 19), Tamil (cluster 16) and Xibe (cluster 18) can be found.

Arabic is clustered with Austronesian languages (cluster 7), in close proximity to the Iranian cluster (cluster 9). The Indic languages (cluster 4) are noticeably distant from the rest, and even further is Japanese (cluster 13). The remaining clusters are formed by only one language, namely: Ghag (12), Thai (14), Basque (24), Greek (30), Cantonese (11) and Chinese (15). With the exception on the latter two, which belong to the Sino-Tibetan family,

all others are the sole representatives of their respective language groups being considered in this analysis.

It is important to note that, even though this form of visualization can be helpful to recognize clusters and connection among languages, some of these connections are spurious and mainly arise from the inherent nature of the minimum spanning tree, which is designed to avoid leaving isolated nodes. Nevertheless, it is remarkable that most of the connections observed in Fig. 5, represented both by the cluster identification of the nodes and the edges linking them, are between languages belonging to the same group. Moreover, most of the languages displayed in the upper part of the figure belong to the Indo-European family, and they all have fusional features. On the other hand, the lower part of the chart mostly consists of agglutinative and isolating languages.

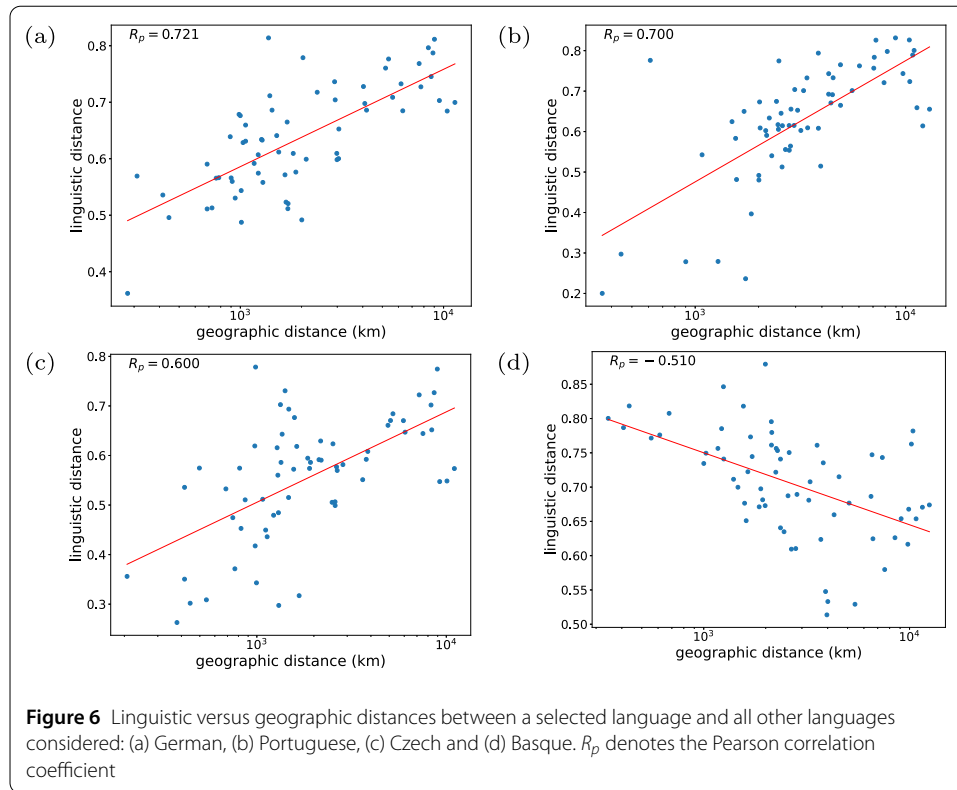
Less known connections, such as among Hebrew and Celtic languages, or between Basque and Armenian are also displayed in the tree. These similarities have been already pointed out in Refs. [48] and [49] respectively, with the proper observation that this does not imply a common origin of the languages. A similar case is that of Altaic languages [50], a putative family relationship between Turkic, Mongolic, Tungusic, Korean and Japanese languages, a subject characterized by ongoing debate within linguistic research [51]. In Fig. 5 this group is represented in the lower part of the tree by the connections among Turkish, Kazakh, Uyghur, Xibe, Buryat and Korean. Despite the controversy surrounding the existence of a genetic relationship, it is generally acknowledged that there exist linguistic similarities among these languages, as evidenced by our analysis.

Romanian language, which is considered to be a Romance language, can be found in Figs. 4 and 5 closer to Germanic languages than to other Romance languages. Interestingly, this fact has been observed in Ref. [26], also using Universal Dependencies data but considering a machine learning approach to calculate linguistic distances. In Ref. [25], a similar cluster analysis based on the analysis of POS trigrams of translations locates Romanian language away from the Romance group.

Another interesting case is that of Arabic, a Semitic language which, in the dendrogram in Fig. 4 is clustered alongside Persian and Kurmanji, both Iranian languages. In the tree presented in Fig. 5, Arabic and Persian are linked but the former is clustered by the k -medoids algorithm with Indonesian and Javanese, which belong to the Austronesian family. In this figure, Arabic appears in the lowest part of the tree, very distant from the other Semitic languages, Hebrew and Maltese. However, this is misleading because their linguistic distance in Fig. 4 is not high, specially between Arabic and Hebrew. Further, the linguistic proximity observed between Arabic and Persian is influenced by their geographic closeness.

Overall, our cluster results are consistent with well established families and linguistic groups. We indeed observe departures that could be attributed to methodology inaccuracies. However, another interpretation is possible. Very recently, a family tree based on the phylogenetic signal of syntactic data has been inferred [52], pointing to salient deviations with respect to the trees derived from the comparative method, which typically did not take into account syntactic data. Therefore, language groups formed upon syntactic analyses need not fully agree with those groups that emerge from phonetic or lexical similarities.

In App. G, we present the results of a similar analysis to the one conducted in this section using tetragrams instead of trigrams. The findings from both methods are consistent,



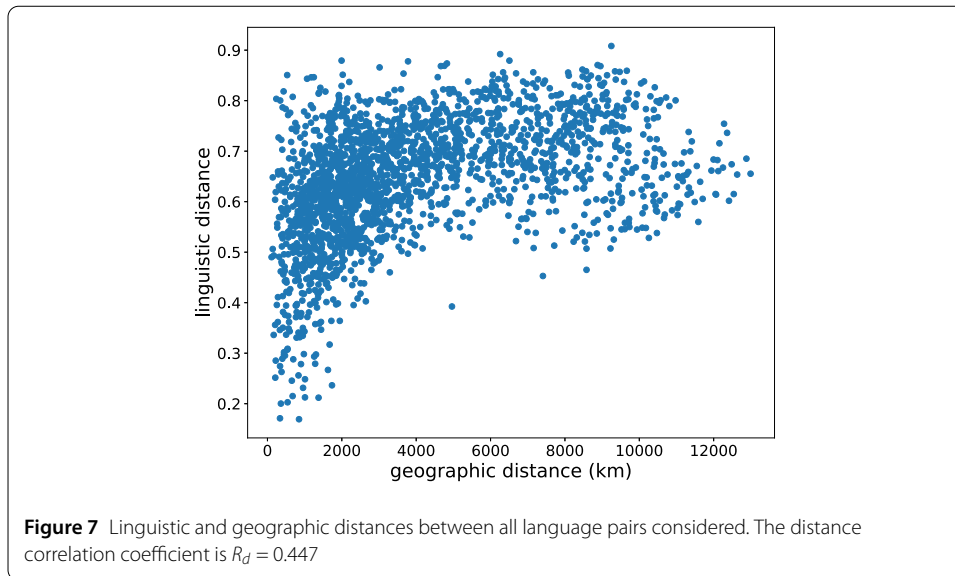
providing strong evidence for our claim that the probability distribution of POS trigrams is sufficient to capture syntactic information in the analyzed languages.

Strikingly, our results also suggest correlations between syntactic distances and spatial proximity. In the following section, we explore the connection between these two variables, which can account for the relationship observed in Fig. 5 between, e.g., Semitic and Romance languages, Indonesian with Vietnamese or Uralic with Slavic languages.

3.2 Relation between linguistic and geographic distances

We calculate the geographic (geodesic) distances between all language pairs by assigning a spatial coordinate to each language considered. This geolocation information is obtained from the World Atlas of Language Structures Online (WALS) [29]. For this analysis we exclude Afrikaans since it is geographically isolated from the rest of the languages considered.

First, for a single language we compute its linguistic and geographic distances with the rest of considered languages. The plots obtained for German (panel a), Portuguese (panel b), Czech (panel c) and Basque (panel d) are presented in Fig. 6, applying logarithmic scale to the geographic distances. The obtained Pearson correlation coefficients ($R_p = 0.721, 0.700, 0.600$, and -0.510 , respectively) point to a logarithmic relation between the two distances (all p-values are < 0.001). Surprisingly, out of the 66 languages analyzed, only Basque [Fig. 6(d)] shows a significant negative correlation between linguistic and geographic distances. This is because Basque is categorized as an agglutinative language and while the majority of Western European languages lean towards fusional characteristics, agglutinative languages are predominantly found in Eastern Europe and



Asia. Consequently, Basque shares more linguistic similarities with languages spoken at greater distances than with those geographically closer.

Finally, we compute the linguistic and geographic distances for all language pairs in order to globally explore the relation between these two variables. We show the resulting plot in Fig. 7. Quite generally, there is a positive correlation between the syntactic and spatial variables. We can quantify this dependence with the distance correlation coefficient R_d [53]. This coefficient is more general than Pearson's R_p since it does not only measures linear dependence between variables. We find $R_d = 0.447$, with a p-value < 0.001 calculated from a permutation test. Even though R_d is not high due to the uncertainties associated to the noisy data and the geographical locations, the value is significantly greater than 0 which indicates that geographical and linguistic distances are indeed correlated.

4 Conclusions

In this work we have collected and analyzed parts of speech tagged sentences available in the Universal Dependencies library for 67 contemporary languages located in a geographically contiguous region. Following an information-theoretic approach, we have provided evidence showing the effectiveness of utilizing the trigram probability distribution of parts-of-speech for characterizing the syntax statistics of languages. Through this method, we have computed distances between languages by calculating statistical divergences between trigram distributions, revealing both well established language groupings and less familiar but already documented linguistic connections. Whereas most analyses of language families are conducted at the phonetic level, our syntactic approach, while yielding similar results, points to a robustness of linguistic family classifications across different linguistic levels. This opens up new avenues for linguistic research where syntactic data can complement phonetic data, providing a more comprehensive understanding of language evolution, classification and geographic relationships. We also stress that our approach is synchronic and thus complements the diachronic views more often encountered in historical linguistics.

A potential impact of our work is in the field of language documentation and revitalization. By quantifying syntactic similarities, our method can help identify which lan-

languages are most similar to endangered languages, therefore guiding efforts to develop educational materials and resources. For instance, if an endangered language has not been extensively documented, educational content from syntactically similar languages can be adapted more efficiently, preserving linguistic heritage more effectively. Furthermore, our approach can be instrumental in computational linguistics, particularly in the development of multilingual natural language processing systems. Understanding syntactic similarities between languages can improve the performance of machine translation systems, especially for low-resource languages. Our results suggest that POS trigrams capture essential syntactic structures, which can be easily integrated into algorithms to enhance cross-linguistic transfer learning. As another potential application, we mention that knowing the language distances across different levels can be useful in language teaching, as it may help focus on the aspects that differ most between the student's native language and the target language.

Furthermore, our analysis has delved into the correlation between linguistic and geographic distances. We have found that spatially proximate languages tend to exhibit more similar morphosyntactic characteristics compared to those located farther apart. Our results suggest a logarithmic relation between these distances. This finding is in fact in good agreement with results reported in Refs. [54, 55], where different measures of linguistic distances are defined. Due to the limitations of selecting a single point to represent the coordinates of an entire language, the results obtained can potentially be improved by considering more accurate methods to assign these locations, such as systematically selecting the regions where the languages are predominantly spoken, or by considering regional varieties that take into account the spatial variation of languages.

We emphasize that the methodology detailed in Sect. 2 for analyzing correlations within discrete sequences is versatile and applicable across numerous disciplines. These techniques hold relevance beyond the realm of linguistics and can be effectively employed in various fields, including but not limited to statistical physics, biology and data science, especially for systems that can be modeled as short-memory stochastic processes with few states.

A limitation of our research is that the UD dataset predominantly covers well documented languages. Therefore, less studied languages and dialects are underrepresented, which may limit the generalizability of our findings to the entire spectrum of global languages, since our results might not fully capture the syntactic diversity present in linguistic minorities. However, the UD library is frequently updated, regularly expanding existing corpora and incorporating new languages. It would be straightforward to extend our analysis to include more languages in future works.

Additionally, the use of POS tags, while effective for our analysis, might oversimplify the complexities of syntactic structures across different languages. POS tagging neglects many linguistic nuances and syntactic intricacies, which may lead to a loss of detailed information. This simplification can potentially affect the accuracy of our distance measurements, as some syntactic phenomena unique to specific languages may not be adequately represented. For example, languages with rich morphological systems or unique syntactic constructions may have their complexities reduced to basic categories that do not fully reflect their syntactic richness. Future work should consider integrating more detailed morphosyntactic annotations, such as dependency and constituency parses, which are already included in the UD library, to capture a broader range of syntactic features.

Appendix A: List of languages

We present in Table A.1 a list of the 67 languages included in our study, categorized by language family, group, and morphological type. The latter is a simplified approximation that is included here for completeness. An exhaustive morphological analysis of the explored languages is beyond the scope of the present work.

Table A.1 List of languages included in this work, classified by language family, group and type

Family	Group	Language	Type	
Afro-Asiatic	Semitic	Arabic	Isolating-Fusional	
		Hebrew	Fusional	
		Maltese	Fusional	
Austro-Asiatic	Viet-Muong	Vietnamese	Isolating	
Austronesian	Javanese	Javanese	Agglutinative	
	Malayo-Sumbawan	Indonesian	Agglutinative	
Basque		Basque	Agglutinative	
Dravidian		Tamil	Agglutinative	
Indo-European	Albanian	Gheg	Fusional	
		Armenian	Agglutinative	
	Armenian	Western Armenian	Agglutinative	
		Baltic	Latvian	Fusional
			Lithuanian	Fusional
	Celtic	Breton	Fusional	
		Irish	Fusional	
		Manx	Fusional	
		Scottish Gaelic	Fusional	
		Welsh	Fusional	
		Germanic	Afrikaans	Isolating-Fusional
		Danish	Fusional	
		Dutch	Fusional	
		English	Isolating-Fusional	
		Faroese	Fusional	
		German	Fusional	
		Icelandic	Fusional	
		Norwegian	Fusional	
		Swedish	Fusional	
	Greek	Greek	Fusional	
	Indic	Hindi	Fusional	
		Urdu	Fusional	
	Iranian	Kurmanji	Fusional-Agglutinative	
		Persian	Agglutinative	
	Romance	Catalan	Fusional	
		French	Fusional	
		Galician	Fusional	
		Italian	Fusional	
		Portuguese	Fusional	
		Romanian	Fusional	
Spanish		Fusional		
Belarusian		Fusional		
Bulgarian		Isolating-Fusional		
Croatian		Fusional		
Czech		Fusional		
Polish		Fusional		
Pomak		Fusional		
Slavic		Russian	Fusional	
	Serbian	Fusional		
	Slovak	Fusional		
	Slovenian	Fusional		
	Ukrainian	Fusional		
	Upper Sorbian	Fusional		

Table A.1 (Continued)

Family	Group	Language	Type
Japanese		Japanese	Agglutinative
Korean		Korean	Agglutinative
Mongolic		Buryat	Agglutinative
Sino-Tibetan		Cantonese	Isolating
		Chinese	Isolating
Tai-Kadai		Thai	Isolating
Tungusic		Xibe	Agglutinative
Turkic	Northwestern	Kazakh	Agglutinative
	Southeastern	Uyghur	Agglutinative
	Southwestern	Turkish	Agglutinative
Uralic	Finnic	Estonian	Agglutinative
		Finnish	Agglutinative
	Mordvin	Erzya	Agglutinative
	Permic	Komi Zyrian	Agglutinative
	Sami	North Sami	Agglutinative
	Ugric	Hungarian	Agglutinative

Appendix B: Predictability gain and block entropy

The predictability gain of going from u th-order to $(u + 1)$ th-order transition probabilities, for $u \geq 1$, can be defined in terms of a conditional relative entropy [56] as

$$\mathcal{G}_u = \sum_{x_0, \dots, x_{u+1}} p(x_0, x_1, \dots, x_u, x_{u+1}) \log \left(\frac{p(x_{u+1} | x_0, \dots, x_u)}{p(x_{u+1} | x_1, \dots, x_u)} \right), \quad (\text{B1})$$

where we adopt the notation that $p(x_0, \dots, x_s) \equiv P(X_0 = x_0, \dots, X_s = x_s)$, with $x_i \in \{z_j\}_{j=0}^{L-1}$, and similarly for the transition probabilities.

We will use the following general result for the difference between the block entropies $H_{r+1} - H_r$, demonstrated in the appendix of Ref. [34]:

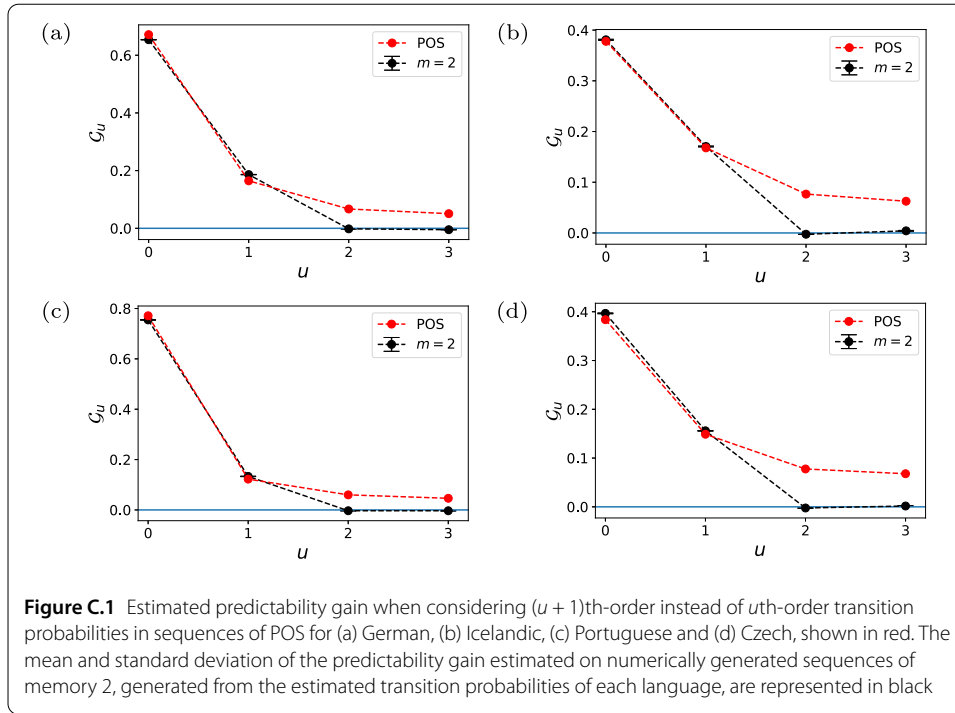
$$H_{r+1} - H_r = - \sum_{x_1, \dots, x_{r+1}} p(x_1, \dots, x_{r+1}) \log(p(x_{r+1} | x_1, \dots, x_r)). \quad (\text{B2})$$

By the properties of the logarithm function, we can write Eq. (B1) as

$$\begin{aligned} \mathcal{G}_u &= \sum_{x_0, \dots, x_{u+1}} p(x_0, \dots, x_{u+1}) \log(p(x_{u+1} | x_0, \dots, x_u)) \\ &\quad - \sum_{x_0, \dots, x_{u+1}} p(x_0, \dots, x_{u+1}) \log(p(x_{u+1} | x_1, \dots, x_u)). \end{aligned} \quad (\text{B3})$$

Adding 1 to all indices in the first sum in Eq. (B3), which is allowed given that we are considering homogeneous sequences, and using $\sum_{x_0} p(x_0, x_1, \dots, x_{u+1}) = p(x_1, \dots, x_{u+1})$, we write Eq. (B3) as

$$\begin{aligned} \mathcal{G}_u &= \sum_{x_1, \dots, x_{u+2}} p(x_1, \dots, x_{u+2}) \log(p(x_{u+2} | x_1, \dots, x_{u+1})) \\ &\quad - \sum_{x_1, \dots, x_{u+1}} p(x_1, \dots, x_{u+1}) \log(p(x_{u+1} | x_1, \dots, x_u)). \end{aligned} \quad (\text{B4})$$



Using Eq. (B2) we express Eq. (B4) as

$$\mathcal{G}_u = -(H_{u+2} - H_{u+1}) + (H_{u+1} - H_u) = -(H_{u+2} - 2H_{u+1} + H_u), \quad (\text{B5})$$

which is exactly the result we wanted to prove. The case $u = 0$ can be demonstrated similarly, considering that

$$\mathcal{G}_0 = \sum_{x_0, x_1} p(x_0, x_1) \log \left(\frac{p(x_1 | x_0)}{p(x_1)} \right). \quad (\text{B6})$$

Appendix C: Memory effects in POS sequences

In Sect. 2.4, we have shown that POS sequences can be approximated as stochastic processes with memory 2 and that this approximation is good. However, this does not necessarily imply that the process has a memory m exactly equal to 2. To see this, we compare the obtained value of $\hat{\mathcal{G}}_2$ with the ones calculated under the null hypothesis. From the original dataset of language \mathcal{L} , composed of R tagged sentences, we compute its second-order transition probabilities by means of Eq. (10). We then generate R sequences, each of equal length as the original sentences, from which we build the set $\{\hat{n}\}_1$. We repeat this procedure K times. By construction, the sets of observations $\{\hat{n}\}_1, \dots, \{\hat{n}\}_K$, built from the generated groups of sequences, have the same size as the original set $\{\hat{n}\}$ of language \mathcal{L} to ensure equal amount of data for comparison. Moreover, these numerical sequences have memory $m = 2$ and consequently we expect the values $\hat{\mathcal{G}}_2[\{\hat{n}\}_k]$ to be close to 0, for $k = 1, \dots, K$. Therefore, we can estimate the p-value p as

$$\hat{p} = \frac{\#\{k : \hat{\mathcal{G}}_2[\{\hat{n}\}_k] \geq \hat{\mathcal{G}}_2[\{\hat{n}\}]\}}{K}. \quad (\text{C1})$$

We apply this method to each of the four languages considered in Sect. 2.4, setting $K = 1000$. We obtain $\hat{p} < 0.001$ for all cases, which leads us to reject the hypothesis that the POS sequences have memory exactly equal to 2. For comparison, we plot in Fig. C.1 with black color the curves corresponding to the mean $\bar{\mathcal{G}}_u$ and standard deviation s_u of the estimated predictability gain for the generated sequences of memory 2, calculated as follows:

$$\bar{\mathcal{G}}_u = \frac{1}{K} \sum_{k=1}^K \hat{\mathcal{G}}_u[\{\hat{n}\}_k], \quad (\text{C2})$$

and

$$s_u = \sqrt{\frac{1}{K-1} \sum_{k=1}^K (\hat{\mathcal{G}}_u[\{\hat{n}\}_k] - \bar{\mathcal{G}}_u)^2}. \quad (\text{C3})$$

We repeat the analysis with the hypothesis that $m = 3$ and obtain similar results, indicating that the POS sequences for these languages possess a memory of at least 4. This is a possible hint of the observed long-range correlations in texts [57]. However, we stress that for the purposes of our work the $m = 2$ approximation works rather well.

Appendix D: Distance between texts belonging to a single language group

We consider a single linguistic group. For each language in this group we randomly select tagged sentences from our database until we reach approximately 10^4 tokens. We iterate this procedure a maximum of 20 times, without replacement (for a few languages there is not enough data to perform this procedure 20 times). Then, for each text portion of each language inside the given group we calculate its probability distribution of POS trigrams and compute the pairwise distances among texts with the JS metric. In Fig. D.1 we present heatmaps corresponding to the distance matrices obtained for Germanic (panel a) and Slavic (panel b) texts. We can observe that, in general, the distance between texts of the same language is smaller when compared to texts coming from distinct languages. This holds even for languages that are considered to be closely related, such as Croatian and Serbian, and Belarusian, Russian and Ukrainian.

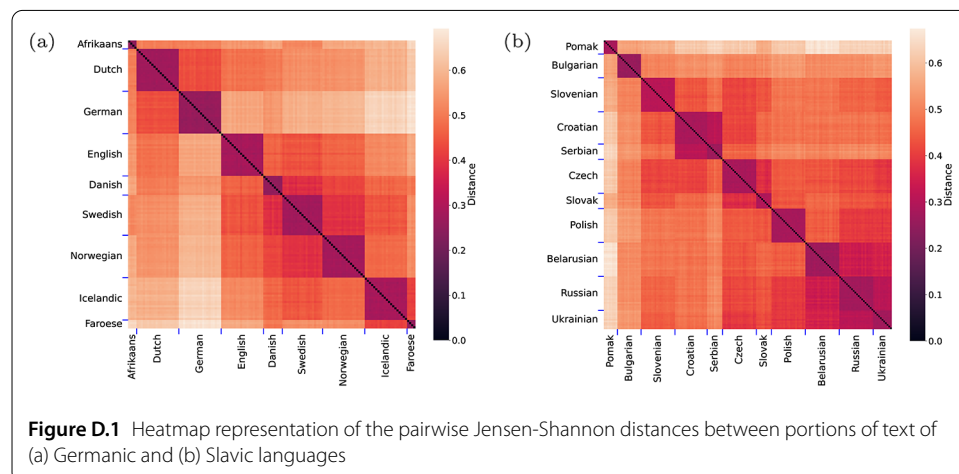


Table F.1 Clusters formed with k -medoids algorithm from the Jensen-Shannon distance matrix

Cluster	Languages
0	Belarusian, Polish, Russia, Ukrainian, Latvian, Lithuanian
1	Catalan, French, Galician, Italian, Portuguese, Spanish
2	Buryat, Kazakh, Turkish, Uyghur
3	Danish, English, Norwegian, Swedish
4	Hindi, Urdu
5	Estonian, Finnish, North Sami
6	Irish, Scottish Gaelic, Welsh
7	Arabic, Indonesian, Javanese
8	Bulgarian, Czech, Slovak, Slovenian, Upper Sorbian
9	Kurmanji, Persian
10	Armenian, Wester Armenian
11	Cantonese
12	Gheg
13	Japanese
14	Thai
15	Chinese
16	Tamil
17	Manx
18	Xibe
19	Korean
20	Afrikaans, Dutch, German
21	Faroese, Icelandic, Romanian
22	Erzya, Komi Zyrian
23	Hungarian
24	Basque
25	Hebrew, Maltese
26	Croatian, Serbian
27	Vietnamese
28	Breton
29	Pomak
30	Greek

This analysis demonstrates the reliability of our proposed metric to discern texts originated from the same language, which is generally a desirable result.

Appendix E: Clustermap obtained from the Hellinger distance matrix

Following the same procedure outlined in Sect. 3.1.1 of the main text, a clustermap was built from the Hellinger (H) distance matrix. The metric in this case is given by

$$d_H(\mathcal{L}, \mathcal{L}') = \sqrt{\frac{1}{2} \sum_{j=0}^{L^3-1} \left(\sqrt{p_{\mathcal{L}}(b_j)} - \sqrt{p_{\mathcal{L}'}(b_j)} \right)^2}. \quad (\text{E1})$$

Similarly to the JS divergence, d_H fulfills the requirements of a distance and lies in the $[0, 1]$ interval.

Figure E.1 shows that the constructed representation resembles the one obtained from the JS matrix. Not only the distributions of the colors match, but also the ordering of the rows and columns obtained with the hierarchical clustering is similar. Only minor differences can be observed in the lower half of the colormaps but, nevertheless, one can identify the same cluster formations.




Figure E.1 Clustermap constructed from the Hellinger distance matrix

Appendix F: Silhouette analysis and clusters obtained with k -medoids algorithm

In order to determine the optimal number of clusters for the k -medoids analysis from the JS distance matrix, we compute the silhouette score [44] for various cluster numbers. The clusterization is then performed using the cluster number that yields the maximum silhouette score. We show the resulting plot in Fig. F.1. As seen, the maximum occurs for 31 clusters.

The clusters obtained using this method are presented in Table F.1.

Appendix G: Minimum spanning tree with tetragrams

In order to validate the results obtained in Sect. 3.1.2 using trigrams, we perform a similar analysis using higher order r -grams for the 67 selected languages.

Specifically, in Fig. G.1 we depict the minimum spanning tree obtained using tetragrams, which can be compared with the tree shown in Fig. 5. For tetragrams, we find that the

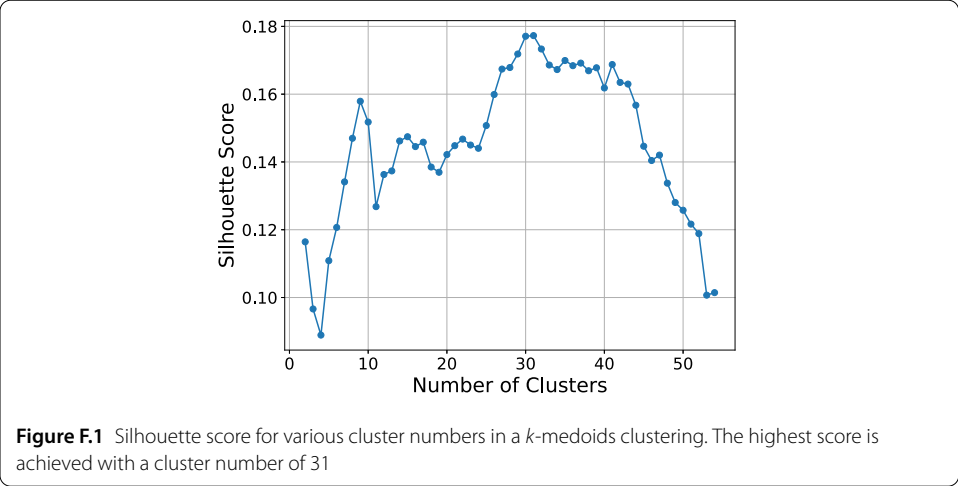


Figure F.1 Silhouette score for various cluster numbers in a k -medoids clustering. The highest score is achieved with a cluster number of 31

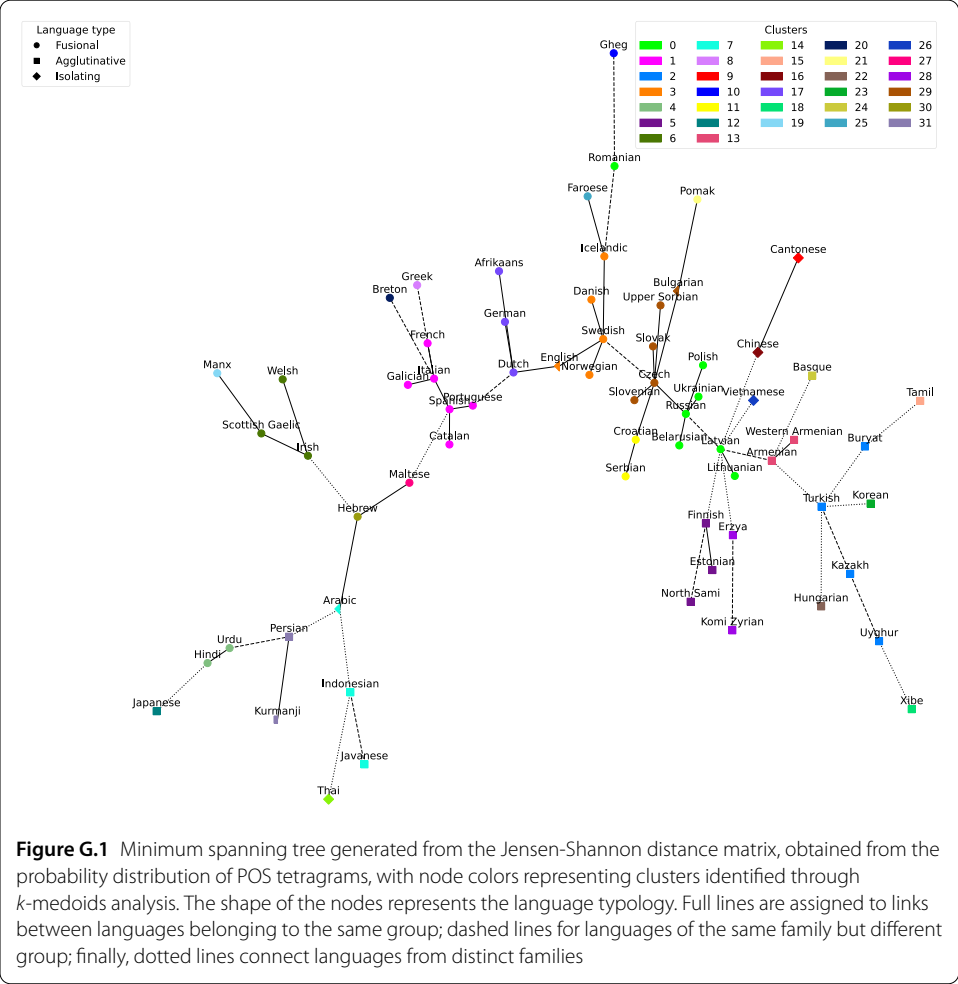


Figure G.1 Minimum spanning tree generated from the Jensen-Shannon distance matrix, obtained from the probability distribution of POS tetragrams, with node colors representing clusters identified through k -medoids analysis. The shape of the nodes represents the language typology. Full lines are assigned to links between languages belonging to the same group; dashed lines for languages of the same family but different group; finally, dotted lines connect languages from distinct families

optimal number of clusters is 32 (instead of 31 for trigrams). We observe that the obtained clusters and connections in the tree are very similar in both cases.

One noticeable difference is that in Fig. G.1 there is a connection between Arabic and Hebrew that was missing in the tree of Fig. 5. This is a nice finding since both languages

are Semitic, although the 3 languages of this family appear in 3 different clusters. This and other discrepancies can be attributed to small data samples, which can contain low values of tokens (of the order of ten thousand for certain languages), whereas the number of possible tetragrams is $15^4 = 50625$.

Choosing the size of the r -grams requires a compromise between the complexity of syntactic structures considered and the accuracy of the parameters estimation. In Sect. 2.4 we have provided evidence that the choice of trigrams offers a good solution to this compromise and the analysis performed with tetragrams supports this claim, since there exist no significant differences with the results based on trigrams.

Abbreviations

POS, Parts of speech; ADJ, adjective; ADV, adverb; INTJ, interjection; PROP, proper noun; ADP, adposition; AUX, auxiliary; CCONJ, coordinating conjunction; DET, determiner; NUM, numeral; PART, particle; PRON, pronoun; SCONJ, subordinating conjunction; PUNCT, punctuation; SYM, symbol; JS, Jensen-Shannon; PAM, Partition Around Medoids; WAL, World Atlas of Language Structures.

Acknowledgements

We thank S. J. Greenhill for useful discussions.

Author contributions

Conceptualization, JDG, DS and RT; methodology, JDG, DS and RT; software, JDG; data collection, JDG; formal analysis, JDG, DS and RT; writing—original draft preparation, JDG, DS; writing—review and editing, JDG, DS and RT; supervision, DS and RT. All authors read and approved the final manuscript.

Funding

This work has been supported by the Agencia Estatal de Investigación (AEI, MICIU, Spain) MICIU/AEI/10.13039/501100011033 and Fondo Europeo de Desarrollo Regional (FEDER, UE) under Project APASOS (PID2021-122256NB-C21), the María de Maeztu Program for units of Excellence in R&D, grant CEX2021-001164-M, and by the Government of the Balearic Islands CAIB fund ITS2017-006 under project CAFECONMIEL (PDR2020/51).

Declarations

Competing interests

The authors declare that they have no competing interests.

Received: 3 April 2024 Accepted: 21 September 2024 Published online: 30 September 2024

References

1. Eberhard D, Simons GFS, Fennig CD (eds) (2023) *Ethnologue*, 23rd edn. SIL International, Dallas
2. Hale M (2007) *Historical linguistics: theory and method*. Backwell Publishing, Hoboken
3. Durie M, Ross M (1996) The comparative method reviewed: regularity and irregularity in language change. Oxford University Press, Oxford
4. Gray RD, Atkinson QD (2003) Language-tree divergence times support the Anatolian theory of Indo-European origin. *Nature* 426(6965):435–439. <https://doi.org/10.1038/nature02029>
5. Gray RD, Drummond AJ, Greenhill SJ (2009) Language phylogenies reveal expansion pulses and pauses in Pacific settlement. *Science* 323(5913):479–483. <https://doi.org/10.1126/science.1166858>
6. Greenhill SJ (2023) Language phylogenies: modelling the evolution of language. Oxford Academic. <https://doi.org/10.1093/oxfordhb/9780198869252.013.61>
7. de Saussure F (2011) *Course in general linguistics*. Columbia University Press, New York
8. Serva M, Petroni F (2008) Indo-European languages tree by Levenshtein distance. *Europhys Lett* 81(6):68005. <https://doi.org/10.1209/0295-5075/81/68005>
9. Holman EW, Brown CH, Wichmann S, Müller A, Velupillai V, Hammarström H, Sauppe S, Jung H, Bakker D, Brown P, et al (2011) Automated dating of the world's language families based on lexical similarity. *Curr Anthropol* 52(6):841–875. <https://doi.org/10.1086/662127>
10. Nerbonne J (2009) Data-driven dialectology. *Lang Linguist Compass* 3(1):175–198. <https://doi.org/10.1111/j.1749-818X.2008.00114.x>
11. Chiswick BR, Miller PW (2005) Linguistic distance: a quantitative measure of the distance between English and other languages. *J Multiling Multicult Dev* 26(1):1–11. <https://doi.org/10.1080/14790710508668395>
12. Mira J, Paredes Á (2005) Interlinguistic similarity and language death dynamics. *Europhys Lett* 69(6):1031. <https://doi.org/10.1209/epl/12004-10438-4>
13. Fernando C, Valijärvi RL, Goldstein RA (2010) A model of the mechanisms of language extinction and revitalization strategies to save endangered languages. *Hum Biol* 82(1):47–75. <https://doi.org/10.3378/027.082.0104>
14. Nerbonne J, Heeringa W (1997) In: Computational phonology: third meeting of the acl special interest group in computational phonology. <https://aclanthology.org/W97-1102>
15. Downey SS, Hallmark B, Cox MP, Norquest P, Lansing JS (2008) Computational feature-sensitive reconstruction of language relationships: developing the aline distance for comparative historical linguistic reconstruction. *J Quant Linguist* 15(4):340–369. <https://doi.org/10.1080/09296170802326681>

16. Heeringa W, Golubovic J, Gooskens C, Schüppert A, Swarte F, Voigt S (2013) Lexical and orthographic distances between Germanic, Romance and Slavic languages and their relationship to geographic distance. *P.I.E. - Peter Lang, Frankfurt am Main*, pp 99–137
17. Donoso G, Sánchez D (2017) In: Nakov P, Zampieri M, Ljubešić N, Tiedemann J, Malmasi S, Ali A (eds) Proceedings of the fourth workshop on NLP for similar languages, varieties and dialects (VarDial). Association for Computational Linguistics, Valencia, pp 16–25. <https://aclanthology.org/W17-1202>
18. Gamallo P, Pichel JR, Alegria I (2017) From language identification to language distance. *Phys A, Stat Mech Appl* 484:152–162. <https://doi.org/10.1016/j.physa.2017.05.011>
19. Eden SE (2018) Measuring phonological distance between languages. Ph.D. thesis, UCL, University College, London
20. Sanders NC (2010) A statistical method for syntactic dialectometry. Indiana University
21. Longobardi G, Guardiano C, Silvestri G, Boattini A, Ceolin A (2013) Toward a syntactic phylogeny of modern Indo-European languages. *J Histor Linguist* 3(1):122–152. <https://doi.org/10.1075/jhl.3.1.07lon>
22. Dunn J (2019) Global syntactic variation in seven languages: toward a computational dialectology. *Front Artif Intell* 2:15. <https://doi.org/10.3389/frai.2019.00015>
23. Manning C, Schütze H (1999) Foundations of statistical natural language processing. MIT press, Cambridge
24. Feldman S, Marin MA, Ostendorf M, Gupta MR (2009) In: 2009 IEEE international conference on acoustics, speech and signal processing. IEEE Press, New York, pp 4781–4784. <https://doi.org/10.1109/ICASSP2009.4960700>
25. Rabinovich E, Ordan N, Wintner S (2017) In: Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: long papers), pp 530–540
26. Samohi A, Mitelman DW, Bar K (2022) In: Proceedings of the 3rd workshop on computational approaches to historical language change, pp 78–88
27. Zeman D, et al. (2023) Universal dependencies 2.13. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University. <http://hdl.handle.net/11234/1-5287>
28. De Marneffe MC, Manning CD, Nivre J, Zeman D (2021) Universal dependencies. *Comput Linguist* 47(2):255–308. https://doi.org/10.1162/coli_a_00402
29. Dryer MS, Haspelmath M (eds) WALS Online (v2020.3). <https://doi.org/10.5281/zenodo.7385533>
30. Comrie B (1989) Language universals and linguistic typology: syntax and morphology. University of Chicago Press, Chicago
31. Crutchfield JP, Feldman DP (2003) Regularities unseen, randomness observed: levels of entropy convergence. *Chaos, Interdiscip J Nonlinear Sci* 13(1):25–54. <https://doi.org/10.1063/1.1530990>
32. Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27(3):379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
33. Raftery AE (1985) A model for high-order Markov chains. *J R Stat Soc B* 47(3):528–539. <https://doi.org/10.1111/j.2517-6161.1985.tb01383.x>
34. De Gregorio J, Sánchez D, Toral R (2022) An improved estimator of Shannon entropy with applications to systems with memory. *Chaos, Solitons & Fractals* 165:112797. <https://doi.org/10.1016/j.chaos.2022.112797>
35. Paninski L (2003) Estimation of entropy and mutual information. *Neural Comput* 15(6):1191–1253. <https://doi.org/10.1162/089976603321780272>
36. Contreras Rodríguez L, Madarro-Capó EJ, Legón-Pérez CM, Rojas O, Sosa-Gómez G (2021) Selecting an effective entropy estimator for short sequences of bits and bytes with maximum entropy. *Entropy* 23(5):561. <https://doi.org/10.3390/e23050561>
37. Nemenman I, Shafee F, Bialek W (2001) In: Dietterich T, Becker S, Ghahramani Z (eds) Advances in neural information processing systems, vol 14. MIT Press, Cambridge
38. Nemenman I, Bialek W, de Ruyter van Steveninck R (2004) Entropy and information in neural spike trains: progress on the sampling problem. *Phys Rev E* 69:056111. <https://doi.org/10.1103/PhysRevE.69.056111>
39. De Gregorio J, Sánchez D, Toral R (2024) Entropy estimators for Markovian sequences: a comparative analysis. *Entropy* 26(1):79. <https://doi.org/10.3390/e26010079>
40. Endres D, Schindelin J (2003) A new metric for probability distributions. *IEEE Trans Inf Theory* 49(7):1858–1860. <https://doi.org/10.1109/TIT.2003.813506>
41. Nielsen F (2016) Hierarchical clustering. Springer, Cham, pp 195–211. https://doi.org/10.1007/978-3-319-21903-5_8
42. Defays D (1977) An efficient algorithm for a complete link method. *Comput J* 20(4):364–366. <https://doi.org/10.1093/comjnl/20.4.364>
43. Kaufman L, Rousseeuw PJ (2009) Finding groups in data: an introduction to cluster analysis. Wiley, New York. <https://doi.org/10.1002/9780470316801>
44. Rousseeuw PJ (1987) Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math* 20:53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7)
45. Gower JC, Ross GJ (1969) Minimum spanning trees and single linkage cluster analysis. *J R Stat Soc, Ser C, Appl Stat* 18(1):54–64. <https://doi.org/10.2307/2346439>
46. Kamada T, Kawai S, et al (1989) An algorithm for drawing general undirected graphs. *Inf Process Lett* 31(1):7–15. [https://doi.org/10.1016/0020-0190\(89\)90102-6](https://doi.org/10.1016/0020-0190(89)90102-6)
47. Haselow A (2011) Typological changes in the lexicon: analytic tendencies in English noun formation, vol 72. de Gruyter, Berlin
48. Gensler O (1993) A typological evaluation of celtic/hamito-semitic syntactic parallels. Ph.D. thesis, University of California
49. Tamrazian A (1994) The syntax of Armenian: chains and the auxiliary. Ph.D. thesis, University of London, University College London, United Kingdom
50. Starostin SA, Dybo AV, Mudrak O, Gruntov I (2003) Etymological dictionary of the Altaic languages, vol 3. Brill, Leiden
51. Janhunen JA (2023) The unity and diversity of Altaic. *Annu Rev Linguist* 9:135–154. <https://doi.org/10.1146/annurev-linguistics-030521-042356>
52. Hartmann F, Walkden G (2024) The strength of the phylogenetic signal in syntactic data. *Glossa: J Gen Linguist* 9(1):1–25. <https://doi.org/10.16995/glossa.10598>

53. Székely GJ, Rizzo ML, Bakirov NK (2007) Measuring and testing dependence by correlation of distances. *Ann Stat* 35(6):2769–2794. <https://doi.org/10.1214/009053607000000505>
54. Nerbonne J (2010) Measuring the diffusion of linguistic change. *Philos Trans R Soc Lond B, Biol Sci* 365(1559):3821–3828. <https://doi.org/10.1098/rstb.2010.0048>
55. Jäger G (2018) Global-scale phylogenetic linguistic inference from lexical resources. *Sci Data* 5(1):1–16. <https://doi.org/10.1038/sdata.2018.189>
56. Cover T, Thomas J (2006) *Elements of information theory*. Wiley, New York
57. Altmann EG, Cristadoro G, Esposti MD (2012) On the origin of long-range correlations in texts. *Proc Natl Acad Sci* 109(29):11582–11587. <https://doi.org/10.1073/pnas.1117723109>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
