Syntactic structures and the general Markov models

Sitanshu Gakkhar
Caltech / USA
University of Toronto / Canada
sgakkhar@caltech.edu

Matilde Marcolli
Caltech / USA
University of Toronto / Canada
matilde@caltech.edu

Abstract

We further the theme of studying syntactic structures data from Longobardi (2017b), Collins (2010), Ceolin et al. (2020) and Koopman (2011) using general Markov models initiated in Shu et al. (2017), exploring the question of how consistent the data is with the idea that general Markov models. The ideas explored in the present paper are more generally applicable than to the setting of syntactic structures, and can be used when analyzing consistency of data with general Markov models. Additionally, we give an interpretation of the methods of Ceolin et al. (2020) as an infinite sites evolutionary model and compare it to the Markov model and explore each in the context of evolutionary processes acting on human language syntax.

1 Introduction

The question of to what extent Markov processes can provide viable models in linguistics has been extensively investigated. Indeed, Markov processes were originally introduced precisely for the purpose of linguistic investigations, to create statistical models of sequences in a text, Link (2006). On the other hand, Markov processes are generally inadequate for the modelling of syntax: early results in generative linguistics showed that regular grammars, generated by finite state automata, do not capture the recursive structures of natural languages, Chomsky (1956), while further limitations on the effectiveness of Markov models can be seen in the fact that pushdown stack automata are exponentially more succint than finite state automata, Berwick (2015). Hidden Markov models play a role in probabilistic context-free grammars, Bridle and Dodd (1987). More interesting for our perspective, models of language change within a Principles and Parameters setting, based on a Markov process in a space of syntactic parameters were developed in Niyogi and Berwick (1997).

In the present paper, we are interested in investigating the adequacy of Markov processes in modelling phylogenetic trees of language families, based on data of syntactic parameters. This complements the recent work towards understanding the extend to which linguistic relationships can be reconstructed bases on syntax along Ceolin et al. (2021, 2020). Typically, mathematical methods for phylogenetic trees reconstruction are based on an underlying assumption that the stochastic variables involved (in our case syntactic parameters) would

be ruled by a dynamics described by a Markov model on a tree. While this hypothesis has generally been justified in the main applications to mathematical biology, the question of its direct applicability to syntactic features is more subtle. Using available data of syntactic structures of world languages, we test the validity of this Markov model hypothesis. Our main conclusion is that working with Markov models on trees provides reasonably accurate results for collections of languages within a given language family, while when the size of the tree grows to include the simultaneous presence of different language families the tree reconstruction becomes more unreliable. This is consistent with what observed using different, coding theoretic methods in Shu and Marcolli (2017), where it is shown that, when one includes different language families and considers the associated lists of syntactic parameters as a binary code, one obtains codes whose position in the space of code parameters is not compatible with a random process of code generation of the type implicit in the usual phylogenetic models. The theoretical problem of identifying a better dynamical model, beyond the Markov processes on trees, to describe evolution of syntactic features remains to be investigated, as well as the relation to the models proposed in Niyogi and Berwick (1997).

The general Markov model on a tree is a widely used setting for modelling and constructing phylogenetic trees. The fact that it allows for describing processes which are not necessarily memoryless is of particular interest in the syntactic structures setting, because one in general expects the length and type of interaction has an effect on the syntax for coevolving langugages. Now real world data coming from a not completely understood process may violate the assumptions that the model makes to various degrees: the data may not be independent and identically distributed, confounding the statistical inference of parameters for the model, or the underlying process may not be well described by a tree. This stymies the attempts to understand the process. Given the data, quantifying how well they fit the general Markov model becomes an important consideration that informs how much of an investment should be made in fine tuning the model to fit the data, versus exploring using more general statistical models to gain insight into the process. We consider the same data of syntactic structures from Longobardi (2017b); Collins (2010) that were analyzed in Shu et al. (2017), using general Markov model derived phylogenetic algebraic geometry. Our goal is to investigate how well the model describes the evolutionary processes on natural language syntax, at the same time comparing the phylogenetic signal we obtain to that of Ceolin et al. (2020). We also explore the question of metricizing the space of syntactic structures that is relevant to the persistent homology machinery used by Port et al. (2019) towards questions on the phylogenetics of language families.

Additionally, we note that the methods of Ceolin et al. (2020) can be reinterpret as describing an alternate model, the infinite sites model. We consider this as an alternative, and point out the similarity between the results from general Markov model the methods of Ceolin et al. (2020).

2 Preliminaries

A phylogenetic tree for species X is a tree T with an identification, $\phi : X \to \text{leaves}[T]$, of the tree leaves with the species set X. The root ρ of the tree is a choice of the a vertex of T. For a rooted tree T on vertex set V and edge set E, with a partial ordering on the vertices given by distance from root, a Markov process on T with state set C is a family of random variables $\{\zeta_v : v \in V\}$, such that if $(u, v) \in E$ then

$$\mathbb{P}(\zeta_v = \alpha | \wedge_{w<v} \zeta_w) = \mathbb{P}(\zeta_v = \alpha | \zeta_u)$$

Such a Markov process can be thought of as obtained by assigning a Markov transition matrix to each edge, governing the dynamics across it. More formally, the κ-state general Markov model on a phylogenetic tree consists of a probability distribution over the state set assigned to the root vertex, together with an assignment of a $\kappa \times \kappa$ transition matrix to each edge. The κ-state random variable χ, called a *character*, evolves from the root to each leaf based on the transition matrices on the path downwards. The probability distribution of the κ states at the leaves can be thought of as a tensor, $\mathcal{P}_T$, indexed by the the patterns possible. This means that the component of the tensor are the probabilities $p_{x_1 x_2 \ldots x_n}$ (with $x_i \in [\kappa]$) of the character χ having state x_i at the leaf i for a n-leaf tree. The central problem is inferring the phylogenetic n-leaf tree given n sequences of length t, where by sequences we mean samples of the values the character takes at each of the leaves.

The inference in the Markov model is usually performed assuming that each site is evolving identically and independently. While it can be reasonably assumed that the topology of the tree is identical for the evolution of each site, assuming that tree parameters are identical can be problematic; selection pressures often induce sites to evolve differently, and the location of the site in the sequence may carry meaning, giving no apriori reason why they should evolve identically, or even independently.

Following Allman and Rhodes (2008), for a n-leaf binary tree with $|E| = 2n - 3$ edges, the parameter space S for the κ-state Markov model sits inside $[0, 1]^N$ with $N = \kappa - 1 + |E|\kappa(\kappa - 1)$, and there is a polynomial map $\phi_r : S \to [0, 1]^{\kappa^n}$ which gives the joint distribution of states at the leaves. Allman and Rhodes (2008) show that there exists polynomials, called the *phylogenetic invariants*, dependent only on the tree topology, which vanish on the image of ϕ_r. This implies that, regardless of the exact parameter values, if the data is indeed coming from a Markov model then it is straightforward to test if the tree topology is supported by the data. With real data the invariants do not exactly vanish, as the probabilities are not exact but only estimates; instead the magnitude of the invariants is often used as a heuristic to select between tree topologies. Phylogenetic algebraic geometry studies the map ϕ_r and the phylogenetic variants. In general, statistical models where such polynomial maps from parameters space to observation space are defined can be studied through an algebraic geometry approach. The general Markov model assumes that the sites are evolving identically independently. There are modification that allow other possibilities, but they require a priori knowledge. There are two assumptions that are implicit in the setup of these phylogenetic models:

1. The sites of the sequences (i.e. the samples of character values at leaves) are independent and identically distributed.

2. The interactions between the taxa at the leaves are described by a tree.

The presence of selection pressures and hybridization, which are both common phenomena, are at odds with these two assumptions. The failure of the first is particularly confounding because of the phenomenon of phylogenetic mimicry: it is well known in the phylogenetics literature that a mixture of data from different phylogenetic models can mimic a different model, in the sense that leaf pattern frequencies can match pattern frequencies that are not distinguishable (see Matsen and Steel (2007); Štefankovič and Vigoda (2007a,b)). While Allman et al. (2012) show that this mimicking is unlikely with mixtures of small number models, when the state space is not small this phenomenon is an issue for 2-state models.

We try to quantify the agreement (or the violation) of these two assumptions with the data of syntactic structures. In the next sections we introduce the key ingredients of our approach: the *logdet* transform and the flattenings of phylogenetic tensors.

The discussion is specialized to the setting of binary characters, both because the phylogenetic algebraic geometry of the 2-state model is the most well developed, and because this is the setting that directly applies to the syntactic structures datasets. However, the techniques carry over to characters with finite state sets in obvious ways; the ideas developed are more generally applicable than to the syntactic structures setting, and can be used when working with the such models.

2.1 Logdet transform and neighbor joining

An important class of metrics relevant to our setting are *tree metrics*: these are metrics on the space of leaves for which there exists a tree with edge lengths such that the distances between leaves correspond to lengths of paths on the tree. The main reason why these metrics are especially useful is that tree reconstruction algorithms that use similarity measures between the data at the leaves for constructing the tree are often well behaved for tree metrics. For example, neighbor joining is a commonly used tree construction method that reconstructs the correct tree topology given an n-point distance matrix from a the tree metric on the tree T.

For neighbor joining, the requirement of being a tree metric can be relaxed so that the reconstruction is still correct as long as each entry of the distance matrix is no more than half the minimum edge length of T from the tree metric associated to T (see, for instance, Theorem 5.8 from Warnow (2017), Theorem 7.7.5 from Semple et al. (2003)). This means the accuracy of construction is compromised if the pairwise dissimilarity between the taxa has a large variation, especially if one assumes that larger dissimilarity corresponds to larger evolutionary distances and larger noise.

Gascuel and Steel (2006) note that neighbor joining greedily optimizes a global criterion – the weighted tree length computed using Pauplin's formula. This has the effect that two

most similar taxa are not necessarily guaranteed to be placed together as that may not be optimal on the full tree, and adding or removing taxa to the set being considered can change relationships inferred between the remaining taxa. Additionally, if the dissimilarity estimates between a small number of taxa are noisier than the rest, on adding more taxa we expect that tree construction to improve because now the outlier has less impact on the tree length. We also empirically observe this in the datasets we consider.

In the context of the Markov model, the *paralinear* distance of Lake (1994), also known as logdet transform, gives a natural tree metric. Lake (1994) defines the paralinear distance[1] $d(S_i, S_j)$ for two sequences S_i, S_j over an alphabet $\{a_k : k \in [n]\}$ as

$$d(S_i, S_j) = -\log \frac{\det J^{ij}}{\sqrt{\det D^1 \det D^2}}$$

where J^{ij} is an $n \times n$ matrix, with (p, q) entry given by the number of instances of (a_p, a_q) in the sequence pair (S_i, S_j), and D^1, D^2 are diagonal matrices with sum of rows and columns (of J^{ij}) respectively on the diagonal. Because the normalized frequencies are approach probabilities, under the assumption that each site of the sequence is independent, this measures how far the joint distribution is from being the product of marginals.

Chapter 8, Section 8.12 of Semple et al. (2003) gives a different formulation that is also useful. For a phylogenetic Markov model on a tree T, with a character χ with state set X evolving on it, and for leaves x, y, define the matrix $J^{xy}_{\alpha,\beta} = \mathbb{P}(\chi_x = \alpha \wedge \chi_y = \beta), \alpha, \beta \in C$, and the matrix P^{xy} with $P^{xy}_{\alpha\beta} = \mathbb{P}(\chi_y = \beta | \chi_x = \alpha)$. One has $P^{xy} = (J^{xx})^{-1} J^{xy}$, implying that the paralinear distance

$$d(x, y) = -\log \frac{|\det J^{xy}|}{\sqrt{\det J^{xx} \det J^{yy}}}$$

becomes

$$d(x, y) = -\frac{1}{2} \log \det P^{xy} P^{yx}.$$

From this observation it is easy to see that, if S_1, S_2 are independent sequences, then $d(S_1, S_2) = \infty$, since P^{xy} has rank 1.

Felsenstein (2004) notes (page 212) that the logdet transform fails to be additive when the stationary distributions for the edge transition matrices do not agree. More generally, it fails when the distribution at the root is not uniform. (Recall that the uniform distribution is a stationary distribution for any Markov transition matrix.) To see this, suppose that we have the tree (ab) with an interior vertex r, and edges $e_a = ar, e_b = rb$, with two edge matrices M_a, M_b (directed along ar, rb). Consider the joint distribution at the vertices r and a, $P^{ra}_{\gamma\alpha} = M_a^{\alpha\gamma} P(a = \alpha)$. Similarly $P^{ba}_{\alpha\beta} = [M_a M_b]^{\alpha\beta} P(a = \alpha)$. The problem happens when computing the joint distribution of b and a after factoring through r. If the distribution at a is not the stationary distribution for M_a then entries M_b are not scaled in a compatible way, and the logdet metric may fail to be additive.

[1]The slight issue with negative determinants in Lake's definition can be side stepped using a constant scaling of the metric and moving it inside the logarithm.

The assumption of a uniform distribution at the root is not very realistic for the syntactic data considered here. Indeed, the phylogenetic algebraic geometry analysis of syntactic structures in Shu et al. (2017) does not assume a uniform distribution at the root. In the setting we consider here, the tolerance for neighbor joining allows for noise and for some deviation from the assumption of uniform distribution at root.

With this last caveat, we have a natural tree metric on the space of sequences coming from a general Markov model; combining it with neighbor joining, the tree topology can be recovered. We will work with the tree constructed using logdet transform and neighbor joining throughout the next sections, and we refer to tree constructed like this as the LOGDET+NJ tree.

2.2 Logdet transform and rate matrices

An important special case is when the edge transition matrices have form e^{Qt} for a real parameter t and a matrix Q called the *rate matrix*[2]. Rate matrices characterize the instantaneous rate of transition between the states for the character and the parameter t can be thought of as time. An important consideration is whether every edge transition matrix, M_e can be given by a single choice of rate matrix, $M_e = e^{Qt_e}$; such a process is a *stationary* process. It is *reversible* when Q is symmetric, the reason being that the parameter t_e associated to each edge can now be interpreted to mean time. *Heterotachy* is the phenomenon where more than one rate matrix is involved, and significantly complicates the analysis. Semple et al. (2003), Section 8.5, notes that for a stationary, reversible process, the logdet transform is closely related to the expected number of substitutions along the edges, which is considered as the *evolutionary distance*.

Consider the covariance $\mathrm{Cov}(C, B)$ of character values at the leaves, i.e. Bernoulli random variables B, C that evolved from the root R on the tree, co-evolving till the interior node to state I_{BC}, then evolving independently. Assume $C = M_C I_{BC}, B = M_B I_{BC}$ for Markov transition matrices, M_C, M_B. By the Law of Total Covariance

$$\mathrm{Cov}(C, B) = \mathbb{E}[\mathrm{Cov}(C, B|I_{BC})] + \mathrm{Cov}(\mathbb{E}[C|I_{BC}], \mathbb{E}[B|I_{BC}]),$$

where $\mathbb{E}[\mathrm{Cov}(C, B|I_{BC}))]$ vanishes, since conditioned on I_{BC}, B, C are independent, while $\mathrm{Cov}(M_C I_{BC}, M_B I_{BC})$ becomes proportional to the variance of the internal state I_{BC}, involving the entries of M_C, M_B.

When the state space is large, or otherwise when that the variance is expected to become proportional to the parameter t, the covariance at the leaves encodes the topology of the tree that can be recovered by a simple greedy strategy: compute all pairwise covariances between the n leaves, group the two leaves with largest covariance, and then replace the leaves that were just grouped by the single node. The covariance between this node and the other leaves is the minimum of the covariances against the constituents of the node. The process is iterated till every leaf is absorbed into some node, the covariances between

[2]Rate matrix is any matrix where each row sums to zero, and all entries are positive off diagonal and not non-positive on it; each edge not is thought of as a continuous Markov chain associated to the rate matrix

two nodes consisting of multiple leaves being computed analogously. The correctness of this follows, since if we assume the evolution happens on the tree then, up to estimation error, the minima of the covariances between leaves in different nodes should all be the same.

In particular, when appropriate assumptions (stationarity, reversibility, proportionality to the time parameter) are met, the disagreement between covaaince and LOGDET+NJ tree indicates heterotachy. This suggests that comparison of the LOGDET+NJ tree and the tree based on covariances can be revealing in general.

2.3 Flattenings, splits and phylogenetic invariants

For a tree T with leaves L_T, a *split* is a partition of the set L_T that is induced by deleting an edge of the the tree. A partition into sets A, B of L_T, not necessarily a split of the leaves, associates a $\kappa^{|A|} \times \kappa^{|B|}$ matrix to the partition, called the *flattening* of the probability tensor, $p_{x_1 x_2 \ldots x_n}$, where we are using the sets $A = \{a_i : i \in [|A|]\}, B = \{b_j : j \in [|B|]\}$ to index the tensor.

$$[\mathrm{Flat}_T(A, B)]_{s_1 \ldots s_{|A|}, t_1 \ldots t_{|B|}} = p_{l_1 \ldots l_{|L_T|}}$$

where l_k is either t_j or s_i, depending on whether the leaf k is b_j or a_i.

The rank of the flattening $\mathrm{Flat}_T(A, B)$ is $\kappa^{\min(1+e(A,B),|A|,|B|)}$ where $e(A, B)$ is the number of edges shared by the subtrees of T obtained by restricting to the leaves A and B (see Pachter and Sturmfels (2005), Theorem 19.5, and Allman and Rhodes (2008), Section 4). If A, B is a split, then the rank is κ, and in particular all $(\kappa + 1) \times (\kappa + 1)$ minors have vanishing determinants. Eriksson (2006) gives a simple way of constructing phylogenetic trees from character data for n taxa by iteratively joining pairs of taxa such that the rank of flattening matrices between the pair and the rest of vertices is closest to κ.

Allman and Rhodes (2008), Theorem 4, shows that for the case of binary trees, for the 2-state general Markov model, the phylogenetic ideal is generated by the 3×3 minors of the flattening matrix for splits induced by each of the edges of the tree.

3 Testing consistency with Markov models

The starting point is a $n \times t$ matrix of data, where each of the n sequences, with t sites, are from a single taxa, with a particular site across sequences representing the evolution of a single character on the tree. Specifically, we work with the data of syntactic structures described in next section – each structure corresponding to a character. There are two major checks that are needed: checking if each site represents an independent and identically evolving copy, and if the tree topology is sufficient.

3.1 Maximum likelihood statistics

The sufficiency of the tree topology is explored using the maximum likelihood on the LOGDET+NJ tree along with behaviour of the rank of the flattening matrices. The key idea here is that if the Markov model does indeed describe the data, then the LOGDET+NJ tree

reconstruction will in the limit give the correct tree topology, T. Coupled with a maximum likelihood estimate of the tree parameters, this gives a way to generate an empirical null distribution against which the statistics can be tested. Generating data from the maximum likelihood model, the distribution of distances from the flattening to the nearest matrix of appropriate rank (where we know the behaviour of the ranks of flattening matrices from last section) provides the expectation against which we examine the actual data. Testing identical evolution in the syntactic structures data is theoretically not possible since there is only one sample from each structure. We do have a proxy that is sufficient (though not necessary) as evidence of failure of independent evolution, simply by considering the joint distribution of two fixed sites against the product of individual site distributions. Again the expectation for this statistic can be empirically estimated using a maximum likelihood model. The simulated data comes from i.i.d. simulated trials on the maximum likelihood trees, this makes z-score usable to characterize the discrepancy of the actual versus the simulated.

3.2 The influence of sites in leaf sequences

A secondary question in the syntactic parameter setting is determining if there are parameters that carry higher relevance than others to determining the relationships between languages in families. We examine this using the idea of *influence* from analysis of boolean function (see O'Donnell (2014)). The tree on n leaves can be represented as a partially ordered collection of subsets of leaves with the order induced by the tree structure. Any algorithm Alg_{Tree} for tree construction can be thought of as a map into the collection of all subsets of the leaves, taking a value one if that subset is present in the output tree representation and otherwise zero. Each site in each sequence in the data affects the output of Alg_{Tree}. The influence of the variable is the probability that changing the value of that variable changes the function.

Intuitively, one expects that a few sites in the data should have small effect on the reconstruction process. Moreover, on average over data distribution, assuming i.i.d. evolution of characters, the expectation is that all sites would have similar influences. However, given this particular data sample, and absent any statistical evidence for or against identical evolution, we can hope to get some insight by considering the homogeneity of sites influences. This can also be revealing about syntactic structures themselves, as well as what structures are key in determining relationships within language families.

We flesh out these ideas more concretely after introducing the syntactic structures program in more detail in the next section.

4 Syntactic structures: Background

Chomsky (1981), along with Chomsky and Lasnik (1993), introduced the Principles and Parameters model of syntax, hypothesizing that syntactic structures for natural human languages can be parameterized by a universal set of binary variables: each variable indicating

the validity of a syntactic construction in that language. In subsequent work, including Longobardi (2017b), Collins (2010) and Dryer and Haspelmath (2013), various families of syntactic features that can be formulated in binary form have been identified and data have been collected on the values of these variables over a significant number of world languages. We consider three independent such sets: the dataset produced by the LanGeLin collaboration (Longobardi and Guardiano (2009), Longobardi (2017a)), collecting the values of syntactic parameters based on the Modularized Global Parameterization approach developed by Longobardi, the more recent data from Ceolin et al. (2020) encoding nominal structures, and the database of Syntactic Structures of World Languages (SSWL) by Collins and collaborators.

Technically, the binary variables used in the SSWL database cannot be regarded as genuine "syntactic parameters" in the sense of the Principles and Parameters model, because of conflation of deep and surface structures. For example, instead of a basic word order variable (BWO) as in The World Atlas of Language Structures (WALS, Dryer and Haspelmath (2013)) Feature 81A, the SSWL dataset has several surface word order variables such as SVO, SOV, etc. (For a discussion of deep and surface structure in word order features see Rizzi (2017) and also Murawaki (2018).) However, as demonstrated by previous analysis carried out on this data set (see for instance Port et al. (2019); Ortegaray et al. (2018)), the SSWL data still provide valid information regarding the distribution of syntactic features across world languages, and historical phenomena of syntactic relatedness. The LanGeLin data can be more reliably considered as genuine syntactic parameters. For simplicity of notation, we will loosely refer to all of the syntactic features collected in these databases in the form of binary variables as "syntactic parameters". This is partly justified by the fact that modern syntactic theory has moved toward a generalization of the notion of parameter with respect to universal grammar (UG) specific parameters, by including parameters that are constructed during language acquisition, or "schemata" in the sense of Longobardi (2017b), where general operations are UG-specified rather than individual parameters. For a recent general theoretical discussion of syntactic parameters we refer the reader to Rizzi (2017). For a general introduction to syntactic structures and the parameters model, we refer the reader to Chomsky and Lasnik (1993) and to the papers collected in the recent volume Karimi and Piattelli-Palmarini (2017), which presents an up-to-date overview of the current understanding of syntactic parameters in the linguists community. For a non-technical introduction to syntactic parameters aimed at a general audience of non-linguists, we recommend Baker (2002).

To each language there is an associated vector of syntactic parameter values which gives coordinates in an ambient metric space, with the choice of metric dependent on the context. A main open question in this parametric model of syntax is identifying a good set of independent variables, or equivalently understanding relations between syntactic parameters and constraints on the locus of possible grammars inside the larger ambient space. We refer to this problem as "the geometry of syntax". Considerable work has been done towards understanding linguistic relationships and how syntax is constrained based on this metric space structure. The latter is an interesting question from the perspective of language

acquisition: within this model of syntax it is assumed that the values of the parameters are learned in the process of language acquisition, based on exposure to a set of positive examples; Niyogi (2006) gives an overview of mathematical models of language acquisition within this syntactic parameters model.

Longobardi and the LanGeLin collaboration introduced the use of syntactic parameters to reconstruct phylogenetic trees of language families of interest to historical linguistics, Longobardi and Guardiano (2009). Linguistic phylogenetic trees based on data of syntactic structures were also analyzed using phylogenetic algebraic geometry in Shu et al. (2017). Topological data analysis of syntactic structures was used in Port et al. (2018) and Port et al. (2019) to identify historical linguistic phenomena not captured by tree structures.

However, as discussed in Shu et al. (2018) and in Port et al. (2019), prior work addressing linguistic relationships based on the analysis of syntactic structures shows certain divergences in the structure of phylogenetic trees, with respect to what is known from historical linguistics. In particular, while the phylogenetic algebraic geometry method of Shu et al. (2017) correctly selects the historically accurate tree among a preselected list of candidates, for languages belonging to preselected and sufficiently small families, tree reconstruction methods based on the use of ℓ_p metrics and neighbor joining, or direct application of phylogenetic packages like PHYLIP to the SSWL data, as well as trees derived from persistent components in the persistent homology computations applied to either SSWL or LanGeLin data, show some significant amount of misplacement of languages both within and across language subfamilies.

In the linguistic context one does not reasonably expect that all leaves are at the same distance from the root (this is known as the "clock assumption" in phylogenetics, see e.g. Warnow (2017)); for example, in the family of early European languages we consider, Tocharian and Hittite are not contemporaneous with others like Albanian and Greek. It is known in the literature that metric space methods are susceptible to failure in absence of the "clock assumption". To see intuitively why this failure happens, notice that, when we compute distances between taxa that far apart in time, distances measured by ℓ_p metrics only see where the vectors describing the taxa disagree, and miss the differences that arose and were undone during the intervening time. This under-estimation of evolutionary distances by the frequently used ℓ_p metrics, due to missing unobserved changes in syntactic structures, approximates an unreasaonable model for linguistic evolution: if such metrics are an accurate representation of the metric relationships, then that implies that a syntacic parameter flips at most once in the evolution process. This makes convergent evolution much less likely and is at odds with known historical phenomena of multiple reversals in some syntactic parameters. Further, in language evolution we do see homoplasy phenomena and horizontal transmission in syntax, as discussed for instance in Longobardi (2012) and detected through persistent first homology computation in Port et al. (2019) and Port et al. (2018).

This leads us to positing that the point of failure here lies in the fact that the metrics used are not capturing the evolutionary distance. The reason for the good results from phylogenetic algebraic geometry also becomes clear: the phylogenetic invariants machinery

is agnostic of the metric structure and only utilizes the general Markov model derived invariants. As the logdet metric is the natural metric in the Markov model setting, we move from using invariants to select phylogenetic trees to using logdet metric to construct them. To establish trust in the contructed phylogenies, we first apply the techniques introduced to testing how well the data match the general Markov model.

Specifically, the phylogenetic algebraic geometry methods of Allman and Rhodes (2008), Pachter and Sturmfels (2007) appear very successful when applied, as in Shu et al. (2017), to trees of language families, and that in itself is evidence in favor of Markov models on trees. However, there are reasons why one can expect significant discrepancies from these models when investigating phenomena of language relatedness at the syntactic level. Markov evolution is a restricted class of models for how syntax/languages can evolve and one does not necessarily expect the relationship between any two languages to be well described by such process. For example, languages evolving in close geographic proximity as in the case of the microvariation phenomena studied in Guardiano et al. (2016), represented in the data of Romance and Hellenic Southern Italian dialects in the LanGeLin data, can present more interaction than permitted by tree models. Known historical linguistic phenomena involving influences across different tree subbranches are well known at the lexical level (the Anglo-Norman bridge for example) but more rare at the syntactic level, although such structures are visible in the persistent first homology studied in Port et al. (2019) and Port et al. (2018) (the Gothic-Slavic-Hellenic loop discussed in Port et al. (2019) for example). Such phenomena are beyond what is describable purely in terms of Markov models on trees. Moreover, different syntactic parameters are not independent variables: some relations are explicitly known (as discussed in Longobardi and Guardiano (2009), Longobardi (2017a) for instance), while other relations can be detected through methods of data analysis, as in Ortegaray et al. (2018), Park et al. (2017), or through methods of coding theory Shu and Marcolli (2017), Marcolli (2016). The presence of dependencies between syntactic parameters violates the Markov models on trees hypothesis that these variables can be treated as identically distributed independent random variables. Thus, the effectiveness of the Markov evolution on trees as a model for syntactic relations between languages lingers on how large the effect of such deviations may be. Our purpose here is to show that, despite these possible discrepancies, the statistics of the data of syntactic structures, tested over sufficient diverse language families, are largely consistent with Markov models on trees when restricting to data within given language subfamilies. We will show that the tree reconstructions obtained by this method become significantly less reliable when the size of the tree is enlarged to include different language families, as the effect of deviations from the Markov evolution hypothesis amplifies with the size of the tree. To be more precise, what we see as the size of the language set grows encompassing different language families is that misplacement errors *within* the subfamilies decrease, while significant misplacements *across* different families occur. We see this, for example, in §6.2 with the Greco-Romance tree, where some misplacements within this subtree disappear when instead of considering only this subset of languages, we consider them within the full Indo-European tree (this subset of languages has a large sampling bias, as it contains a large number of closely related Italian

dialects, considered in the microvariations study of Guardiano et al. (2016)). Examples of misplacements across families can be seen, for instance, in the placement of Welsh within the Germanic tree, in the case of the full Indo-European tree of §6.5. This points, on the one hand, to an improved performance of the neighborhood joining within subfamiiles, but at the same time to a more visible discrepancy with respect to the Markov model hypothesis when different subfamilies are simultaneously taken into consideration.

Understanding when the general Markov model applies, servicing the logdet as the natural evolutionary metric, also gives insight into the Geometry of Syntax paradigm of Port et al. (2018) and Ortegaray et al. (2018) which grapple with choice of metric when trying to understand the geometry: we note that when studying evolutionary relationships it is the evolutionary distance that should be considered.

We note that these databases have been updated since the analysis of Port et al. (2019); Ortegaray et al. (2018), with the SSWL dataset especially being subject to frequent additions and updates of parameter values. This results in some minor discrepancies in values of some invariants that we compute with respect to prior results but these do not change the main conclusions.

4.1 LanGeLin dataset

The LanGeLin dataset collects the values of 83 syntactic parameters based on the Modularized Global Parameterization approach developed by Longobardi, for a set of 62 languages, mostly Indo-European. A complete list of the languages and parameters included in this database is reported in §1.2 of Port et al. (2019).

4.2 Entailment in the LanGeLin dataset

The syntactic parameters from Longobardi's LanGeLin collaboration dataset take on values ± 1 as well as 0 with zeroed values indicating dependence on other parameters. To ensure the assumption of independent evolution of parameters, we disregard all parameters that take on a zero value in the language family in consideration. When defining a metric based on these parameters, this leads to a bias towards underestimation because, when computing the dissimilarity, if the parameters underlying the dependent parameters differ then the dependent parameters will also differ. However, we have disregarded them, leading to a dampening of perceived syntactic difference. This can be viewed as a special case of not all parameters contributing uniformly to the syntax. We briefly touch on this in the discussion. This effect is also present on the SSWL dataset, though the dependence there is not explicitly identified.

In the geometry of syntax formalism, the functional dependence of zeroed parameters is exactly what defines the geometry, and are of particular interest from that perspective. Since we expect this functional dependence to be different for different language families, the scheme of disregarding parameters with zeros only across the language family considered is sufficient.

4.3 Syntactic Structures of World Languages (SSWL) dataset

The current version of SSWL dataset contains 252 languages and 115 syntactic binary variables. The list of languages and syntactic features of the SSWL dataset is discussed in detail §1.2 of Port et al. (2019). The set of languages included in the database range across several non-Indo-European language families: the most represented families are, in decreasing number of languages: Indo-European, Niger-Congo, Austronesian, Afro-Asiatic.

An issue with the SSWL data is that the syntactic features are very unevenly mapped across the languages in the database: some languages have 100% of the syntactic features recorded, while others are only 2% mapped. Any subjectivity that may enter analysis in dealing with this incompleteness is removed by following the approach of previous work, where one either sets incomplete parameters to 0 (with ± 1 the binary values of recorded parameters) or one chooses to work only with those parameters that are completely mapped for the language family under consideration (the advantages and disadvantages of these methods are discussed, for instance, in Port et al. (2019), Port et al. (2018), Shu and Marcolli (2017), Shu et al. (2017)). Note that the second method does bias the analysis towards Indo-European languages, which tend to be more extensively mapped in the SSWL database.

4.4 The Ceolin et al. (2020) nominal structures data

The Ceolin et al. (2020) encode the nominal structures in 69 languages across 13 Eurasian languages using 94 binary variables. The dataset is significantly more complete than either the LanGeLin dataset or the SSWL dataset. The parameters show entailment like the LanGeLin dataset, with entailed parameters marked by using a zero value, as opposed to $\pm$ values otherwise. For a more complete description, we refer to the Ceolin et al. (2020) dataset. We note that there are sets of languages that are degenerate in this set in the sense that for all laguages in these subsets all syntactic structures are identical; we only keep one representative from each subset while Ceolin et al. (2020) use all; we do this since keeping multiple representative adds no information but can bias neighbor-joining because of how it minimizes the balanced minimum evolution criterionGascuel and Steel (2006).

4.5 Reinterpreting the Ceolin et al. (2020) metric

Ceolin et al. (2020) use a modified Jaccard similarity value with Unweighted Pair Group with Arithmetic Mean (UPGMA clustering. In their case the parameter only contributes to the syntactic distance between two languages when it's set in at least one them:

$$d_{\text{modified\_jaccard}}(l_1, l_2) := \frac{N_{-+} + N_{+-}}{N_{-+} + N_{+-} + N_{++}}$$

where N_{ab} is the frequency of value a in for parameters from languages l_1 and b for l_2. This can be thought of as modelling an infinite sites evolutionary model in the sense that it is counting how many events happened in the evolution of the sequence on the tree and how

many of them were different between the pair. There is no contribution from unobserved changes, in effect parameters once set are not unset till an evolutionary split happens, meaning along any branch a site, meaning a nominal structure, changes once. Because the number of structures that separate closely related languages is small with respect to the number of structures, this scheme approximates the infinite sites model of evolution**?**. Additionally as all languages in this data are currently extant, the assumption that all languages are all same distance from root, and therefore choice of UPGMA reconstruction made by Ceolin et al. (2020) is reasonable. This approximate model is an alternative to the general Markov model, the key difference being the possibility of multiple updates to same syntactic structure.

As noted byCeolin et al. (2020), there are asymmetries in state transitions, with transitions primarily only observed in one direction (we see also this asymmetry in the maximum likelihood model we obtain). This asymmetry makes unobserved changes across an edge unlikely, so we expect that a highly asymmetric model with approximate the this model. Pushing this a step further, if the evolutionary process is well described by such a model, then $d_{modified_jaccard}$ would be approximately additive. We will use neighbor joining with this metric to get the correct reconstruction gaurantees that it offers, and use that test if there's alternative that betters fits the data compared to the general Markov model.

Ceolin et al. (2020) also reconstruction using Bayesian phylogenetics (built on the Markov model approach with rate matrices) obtaining results that can be considered arguably better than the UPGMA approach, giving weight to the Markov model. The point we want to make is that one does not expect the evolutionary process along any branch to be necessarily memoryless, that is, with an underlying rate matrix: the evolutionary process for syntactic parameters is less like molecular sequence evolution which is provides motivation for Bayesian phyologenetics but the closer analog in biology is stem cell differentiation which has been modelled as a non-Markov (in the sense that the process is not memoryless - we do not get substitutions accumulating as the exponential of a rate for the length of the branch (Stumpf et al. (2017)). Because of this, the Gamma site model they infer may not be the most appropriate to syntactic evolution, leaving us in the more general setting without rate matrices.

5 Markov evolution in language families

The LanGeLin and SSWL datasets are still active projects and only partially complete, with Indo-European languages being most completely defined. So considerable prior analysis has focused on Indo-European languages. To test the ideas put forward in the previous section, we consider the following five families studied in Ortegaray et al. (2018); Port et al. (2019); Shu and Marcolli (2017).

1. *Germanic*: Dutch, German, English, Faroese, Icelandic, Swedish.

2. *Slavic*: Russian, Polish, Slovenian, Serb-Croatian, Bulgarian.

3. *Romance*: Latin, Romanian, Italian, French, Spanish, Portuguese.

4. *Altaic-Uralic-Tungusic*: Finnish, Estonian, Hungarian, Khanty, Udmurt, Yukaghir, Turkish, Buryat, Yakut, Even, Evenki

5. *Early Indo-European*: Hittite, Tocharian, Albanian, Armenian, Greek.

The last family includes some of the early branchings of the Indo-European family tree. There was some debate in recent years in the linguistic community (see Perelysvaig and Lewis (2015)) around computational reconstructions of the structure of the Indo-European tree near the root. In particular, this subset of languages was chosen in Shu et al. (2017) in order to compare the relative positions of the Anatolian and Tocharian branches and the Albanian, Armenian, and Hellenic branches, between two candidate trees, one obtained in Bouckaert et al. (2012) on the basis of lexical data and one, generally regarded by linguists as more reliable, obtained in Ringe et al. (2002) including morphological data. The phylogenetic algebraic geometry method, applied in Shu et al. (2017) to the SSWL syntactic data for this set of languages, slightly favors the tree of Ringe et al. (2002). While some of these ancient languages, like Ancient Greek, are very completely mapped in the SSWL database, others like Hittite and Tocharian are only very coarsely mapped. This implies that there are only 22 variables in the SSWL dataset that are fully mapped for all of these ancient languages. Since the analysis in this case is based on a very small set of syntactic features, it should be regarded as less reliable than the cases of the other families above, for which a larger set of completely mapped parameters is available.

We use the combined set of parameters from SSWL and LanGeLin datasets when the languages in consideration are present in both databases. For Romance and Slavic families which are present in both databases we use the combined parameter values from both, restricting to parameters which are set in all languages only. So this is the analysis that is based on the most complete set of data; although still having to drop partial unset parameters is un-ideal. For the other families too, as in Shu et al. (2017), we only use parameters that are set for all languages.

The results obtained in this way are discussed in the following subsections; we defer the discussion of romance family to the end, as here we find that not only is the LOGDET+NJ tree different from historically correct tree, but also has a lower phylogenetic invariant.

Working in the setting of binary syntactic structures, the general Markov model setup is specialized to the binary characters. The sequences at leaves are also binary and the transition matrices are 2×2. One could consider the unset parameter to be a third symbol, however, the statistics when a parameter is unset in both languages under consideration become ambiguous, so working only with those parameters that are completely mapped for the selected subfamily of languages is favored.

We first construct LOGDET+NJ trees for these families, and evaluate how consistent these are with what is accepted in the linguistics community. Note that the trees constructed are unrooted. This is the case also when one applies the phylogenetic algebraic geometry methods (this issue for linguistic phylogenetic trees is discussed in detail in Shu et al. (2017)). In particular the placement of the root is related to the knowledge about ancient

languages in the database. While for the Indo-European language family, several ancient languages are represented in the data, and this information can be used to correctly root the trees, for language families where only the modern languages are represented in the data one can obtain the information on the tree topology but not as a rooted tree.

By arguments outlined, under the assumption of general Markov model (including the uniform distribution at the root), the LOGDET+NJ tree will recover the correct tree. The reconstructions for Germanic, Slavic and Altaic languages are briefly discussed, before we focus on the two cases which lend themselves to a richer analysis.

5.1 Germanic languages

The LOGDET+NJ tree constructed for the Germanic family using the 89 completely mapped syntactic parameters correctly identifies the separation between West Germanic (Dutch, German), and the East Germanic (Swedish, (Icelandic, Faroese)). The LOGDET+NJ trees are unrooted. The usual method of rooting trees by choosing an outgroup representative is not meaningful here, since the outgroup element may not be evolutionarily related, or the evolutionary distance may be so large that the noise in estimating it will significantly affect the results. Thus, we have simply placed the root in the tree where it is known to be from historical linguistic information, while the LOGDET+NJ tree is simply providing the tree topology.

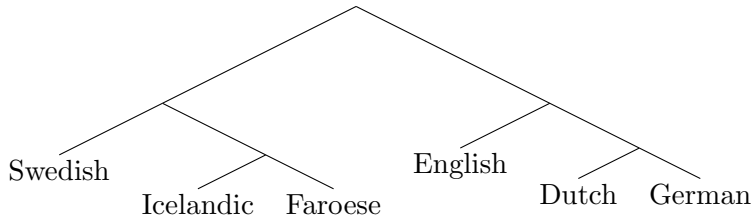

Figure 1: The Germanic LOGDET+NJ tree, showing the East and West Germanic split, with the root placed according to the historically accepted tree.

The reconstruction is very robust and with $\approx 60\%$ of the data, we can recover this topology with probability approximately 0.7.

We note that using various ℓ_p metrics (with UPGMA tree construction) fails to recover the East/West Germanic split, as does in the tree of the persistent connected components of §6.4 of Port et al. (2019) which mixes North and West Germanic languages. The logdet transform is better proxy for evolution distance than ℓ_p metrics. Specifically ℓ_p metrics do not account for unobserved changes in the syntactic structures. This observation underscores using evolutionary distance to explore the geometry of syntax over embedding into other metric spaces.

5.2 Slavic languages

For the Slavic family, there are 68 parameters that are fully mapped between the two datasets: 45 for Longobardi and 23 for SSWL. With the most recent version of the SSWL data, these parameters do not separate Slovenian from Russian, so for this reason we have excluded Slovenian. The version of the data used in the phylogenetic algebraic geometry analysis of Shu et al. (2017) correctly placed Slovenian close to Serb-Croatian, in the South Slavic subbranch, while with the later version of the data used in the persistent components tree of Port et al. (2018), the current ambiguity is resolved by (incorrectly) placing Slovenian next to Russian.

	russian	polish	slovenian	serb-croatian	bulgarian
russian	0.000000	0.127036	0.000000	0.029729	0.092947
polish	0.127036	0.000000	0.127036	0.160433	0.232805
slovenian	0.000000	0.127036	0.000000	0.029729	0.092947
serb-croatian	0.029729	0.160433	0.029729	0.000000	0.126210
bulgarian	0.092947	0.232805	0.092947	0.126210	0.000000

Figure 2: Loget distance matrix for Slavic family including Slovenian

There are still 68 parameters between SSWL and Longobardi datasets that are mapped for the four remaining languages. The SSWL parameters for this family are very poorly mapped, and the distance matrix for the SSWL data alone is highly degenerate. We work with the full collection of 68 parameters spanning the two datasets. Constructing the LOGDET+NJ tree we get

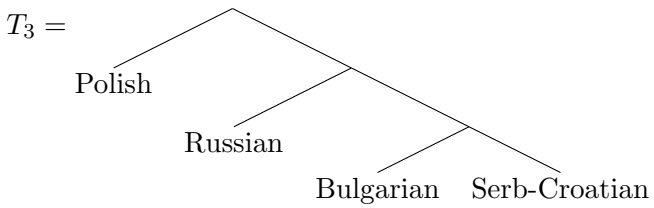

Here again the tree is unrooted, and we have chosen to draw it so that the root is placed consistent with historical linguistic knowledge. Since there are only four branches, the only information contained in the tree topology is the placement of the unique internal edge, namely splitting of the leaves into two pairs of adjacent vertices, $\{\text{Polish}, \text{Russian}\}$ and $\{\text{Bulgarian}, \text{Serb-Croatian}\}$, which here correctly reflects the grouping together of the South Slavic branch.

With Slovenian excluded, all three Longobardi et al. (2013); Shu et al. (2017); Nurbakova et al. (2013) agree on this tree. The reconstruction is robust: with approximately 60% of the parameters sampled uniformly randomly, this topology appears with a probability in

excess of 0.6. An alternate topology which places Russian with Bulgarian also appears in some subsamplings; this proximity of Bulgarian to Russian is also observed in the tree of persistent connected components from Port et al. (2018).

5.3 Altaic-Uralic-Tungusic languages

As in Port et al. (2019), we consider the languages belonging to the Uralic family (Estonian, Finnish, Hungarian, Udmurt, Yukaghir, Khanty) and to the more hypothetical Altaic family (Turkish, Buryat, Yakut, including the Tungusic languages Even and Evenki).

Compared to what was obtained by the topological method of Port et al. (2019), for the Altaic-Uralic-Tungusic languages LOGDET+NJ recovers a tree that appears more consistent with known linguistic relations. We first consider the full set of Altaic-Uralic-Tungusic languages listed above, using 28 parameters from the LanGeLin dataset to obtain the (unrooted) tree:

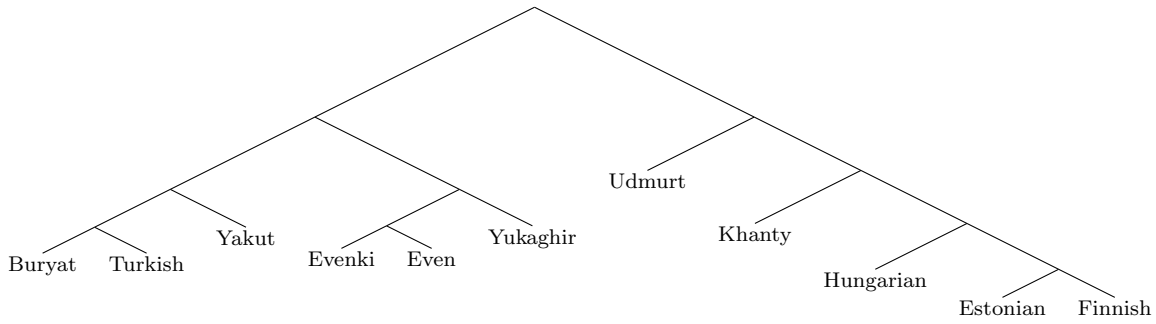

The LOGDET+NJ tree is unrooted: we have drawn it so that the root is placed at the divide between the Uralic and the Altaic-Tungusic languages. Notice that here the two groups of languages are clearly separated, with the only misplacement, with respect to this divide, consisting of the Uralic language Yukaghir that is placed together with the Tungusic languages Even and Evenki.

This tree recovers the large scale structure of the family: Udmurt separates out from rest in the Uralic subtree (Khanty, Hungarian, Finnish, Estonian), which is sensible as it is the lone representative of Permic branch of Uralic languages, while in Port et al. (2019) it was placed into the Altaic subtree. The Altaic subtree is here recovered as (Yakut (Buryat, Turkish)). Regarding the misplacement of Yukaghir, note that this language is considered sufficiently distinct from the Uralic languages to form an Uralo-Yukaghir meta-family and the extent of the relationship between Uralic and Yukaghir is a subject of active investigation (see for instance Piispanen (2013)). Its placement close to the Tungusic Even and Evenki is more consistent with geography (Even, Evenki, Yukaghir all belong to parts of Eastern Russia) than reflecting the Uralo-Yukaghir relation.

The position of Khanty in the Uralic subtree is not as expected: Hungarian and Khanty should separate from the Balto-Finnic subtree (as they do in the topological analysis of

Port et al. (2019)). This can be compared with the historically agreed tree for the Uralic family:

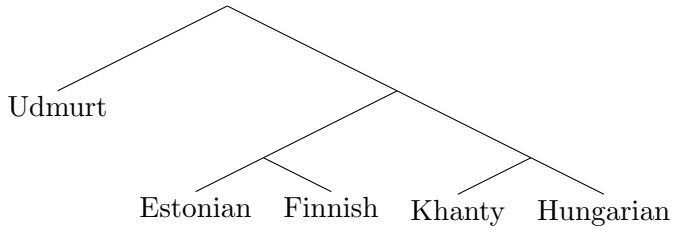

Examining the consistency of this construction, the topology of the LOGDET+NJ tree can be recovered in the case of Uralic with probability a half. There are three topologies that appear in the simulated data, including the correct topology.

5.4 Early Indo-European languages

In this case the sparseness of data – there are only 22 parameters that are completely set for this set of languages – makes this dataset difficult to work with, and conclusions drawn from the analysis should not be regarded as very reliable. There are additional issues: for example, the values of one parameter each for Hittite and Tocharian has been updated [3] since the analysis of Shu et al. (2017), and this renders the Hittite and Tochrian degenerate on the space of the parameters that are completely set. We use this dataset after rolling the update back for comparison with Shu et al. (2017), but this does confound the results.

We find that the LOGDET+NJ tree for Early Indo-European languages rooted appropriately near the Anatolian-Tocharian split is the same as one obtained by Bouckaert et al. (2012). In Shu et al. (2017) phylogenetic invariants are computed for the Ringe et al. (2002) and Bouckaert et al. (2012) tree and it is observed that the tree of Ringe et al. (2002) has a smaller phylogenetic invariant. The disagreement with LOGDET+NJ tree could be interpreted as implying that the evolutionary processes acting on syntax are again not Markov, but this is not confirmed by the phylogenetic invariants computation (also based on the Markov model) that favor the tree of Ringe et al. (2002). It is possible that the discrepancy between the LOGDET+NJ approach and the phylogenetic approach here may reflect the fact that phylogenetic invariants, in the model of Allman and Rhodes (2008), allow for a nonuniform distribution at the root, while as observed earlier the logdet transform fails to be additive when the distribution at the root is non-uniform (Felsenstein (2004), page 212). Since we are looking here at a group of languages that branched out very close to the putative root of the Indo–European tree, this issue may be significant.

On exploring this further we find that there exists a set of 17 parameters [4] from this set

[3] In the updated SSWl Hittite has "11 Adposition Noun Phrase" set to value 0 and Armenian (Western Armenian) has "Neg 01 Standard Negation is Particle that Precedes the Verb" set to value 1.

[4] 01 Subject Verb, 06 Subject Object Verb, 11 Adposition Noun Phrase, 13 Adjective Noun, 15 Numeral Noun, 17 Demonstrative Noun, 19 Possessor Noun, 21 Pronominal Possessor Noun, Neg 03 Standard Negation is Prefix, Neg 08 Standard Negation is Tone plus Other Modification, Neg 10 Standard Negation is Infix, Neg 12 Distinct Negation of identity, Neg 13 Distinct Negation of Existence, Neg 14 Distinct Negation of Location, Order N3 01 Demonstrative Adjective Noun, Neg 04 Standard Negation is Suffix, 12 Noun Phrase

$T_1 =$

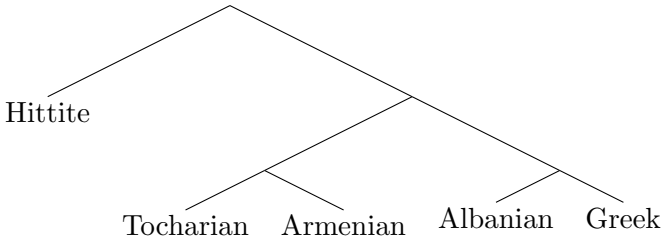

Figure 3: Bouckaert et al. (2012) tree that agrees with LOGDET+NJ tree

$T_2 =$

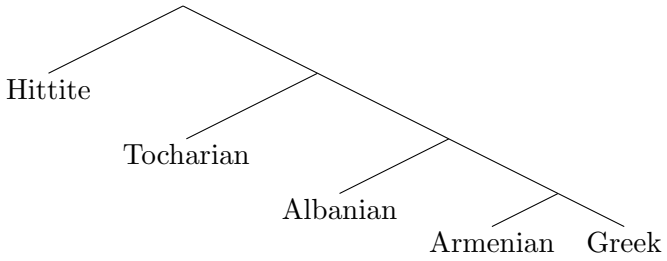

Figure 5: Ringe et al. (2002) tree based on morphological data

of 22 that gives the tree from Rexová et al. (2003):

$T_3 =$

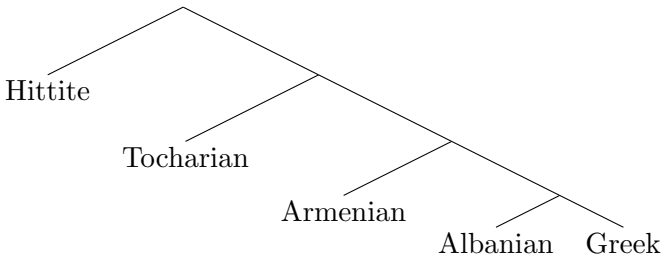

Figure 4: The tree from Rexová et al. (2003) obtained on restricting to 17 out of the 22 parameters that are recorded for this set of languages. Note that the Rexová et al. (2003) is almost the Ringe et al. (2002) tree, except that that Armenian and Albanian are switched.

Adding any other parameter, or removing too many, makes the tree approach the Bouckaert et al tree. This shift suggests that there is an influential (in the sense of being one of the few discriminants of pairs in the family) set of parameters that does not behave like the rest. This indicates that requirement of identical evolution does not hold. Such anomalies are precisely of interest to the linguists studying syntax.

Adposition

The Rexová et al. (2003) construction is based on lexicographic data, like the tree of Bouckaert et al. (2012), while the tree of Ringe et al. (2002) includes both lexicographic and morphological data. None of these previous analyses, with the exception of Shu et al. (2017), are based on syntactic parameters, so the syntactic input can be seen as an independent verification. We find that the tree of Rexova et al. lies between the tree of Bouckaert et al., and that of Ringe et al. when evaluated with phylogenetic invariants: Shu et al. (2017) compares T_1 of Bouckaert et al. (2012) and T_2 Ringe et al. (2002), computing T_3 invariants (on rolled back parameter data) yields

$$\|\Phi_{T_1}(P)\|_{\ell_\infty} = \frac{8}{1331} \approx 0.0060, \|\Phi_{T_1}(P)\|_{\ell_1} = \frac{61}{2662} \approx 0.0229.$$
$$\|\Phi_{T_2}(P)\|_{\ell_\infty} = \frac{8}{1331} \approx 0.0060, \|\Phi_{T_2}(P)\|_{\ell_1} = \frac{18}{1331} \approx 0.0135.$$
$$\|\Phi_{T_3}(P)\|_{\ell_\infty} \approx 0.0060, \|\Phi_{T_3}(P)\|_{\ell_1} \approx 0.0185.$$

Thus, phylogenetic invariants still favor the tree of Ringe et al. (2002). There is an interesting point here, in the discrepancy between the result of phylogenetic invariants computation, which is directly based on a geometric formulation of the Markov model hypothesis, and the Markov model based logdet tree that agrees with the tree of Bouckaert et al. (2012). A possible issue with the uniformity assumption at the root was mentioned above, but there can be other factors involved. This can be compared with the discussion in the next section of the misplacements within the Romance language family.

If the trees appearing were completely random, then one would class that as noise and inherent instability due to sparseness of data, but with support in literature it appears to be more interesting, and suggests that same signals that appear in other data are present here as well; particularly that syntactic structures data is consistent with combined lexicographical and morphological data.

We further try to explore the the presence of highly influential structures that seem to be hinted at as noted previously.

5.4.1 Influence analysis

To explore the effect of linguistic parameters on determining how distinct each language is from the others in the group we do an influence analysis. A tree is thought of as a boolean function by fixing a root and then considering each of the subsets of the leaves that appear under each interior vertex. We are interested in the quantifying how likely the tree is to change on applying *noise* to a parameter: we pick a random set S of parameters, including the parameter of interest, i; given the parameter vector $v_l = (v_{l,x})_x$ for a language l, we flip all coordinates of v_l that are in S to obtain a new vector v_l^S. We bound the size of the noise sets to have intersection of size at most k with the sets $\{x : v_{l,x} = 1\}$ and $\{x : v_{l,x} = 0\}$, so as not to wash out the parameter of interest i. We define the influence of i by

$$\eta_i^k = \sum_{S \in \Gamma_i^k : \textsc{Tree}[v] \neq \textsc{Tree}[v_S]} \frac{1}{|S|}$$

where $\Gamma_i^k = \{S : i \in S, \max(|S \cap \{x : v_{l,x} = 0\}|, |S \cap \{x : v_{l,x} = 1\}) \leq k\}$. The normalization by $|S|$ adjusts the contribution towards the sensitivity of i. We tabulate the parameters which carry largest sensitivity for each member l of the family and for varying k (Figure 6); the influences - we think of highly sensitive parameters as having a higher influence) can also be used to collect parameters to which the family is more sensitive, for e.g 5.4.1.

Because of Hittite and Tocharian are separated by only one parameter, we only consider Hittite in this analysis; the large similarity between the two will not give any meaningful insight.

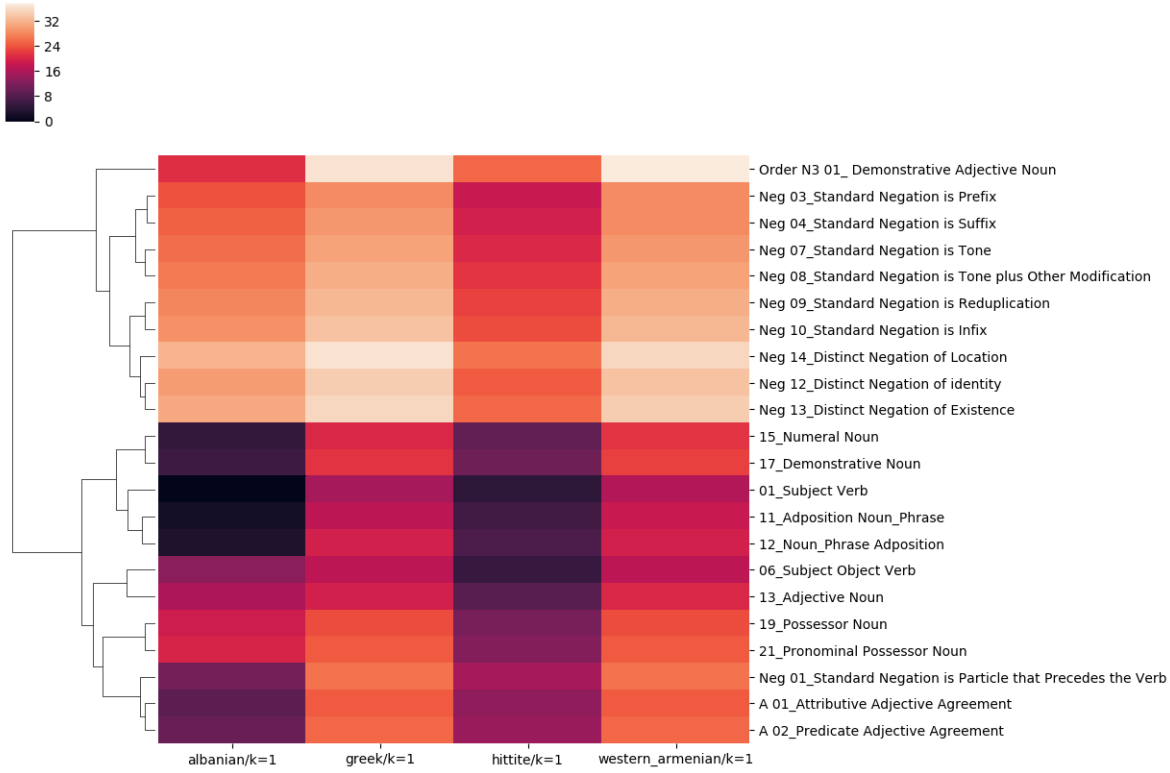

Figure 6: The clustered heatmap showing how parameters group based on their influences: a cluster of parameters with high influences separates out.

	albanian	greek	hittite	wt_armenian
01_Subject Verb	0.0	15.5	4.5	16.5
06_Subject Object Verb	13.0	17.5	5.5	17.5
11_Adposition Noun_Phrase	2.0	17.5	6.5	18.5
12_Noun_Phrase Adposition	3.0	19.5	7.5	19.5
13_Adjective Noun	16.0	19.5	8.5	20.5
15_Numeral Noun	5.0	20.5	9.5	21.5
17_Demonstrative Noun	6.0	21.5	10.5	22.5
19_Possessor Noun	19.0	23.5	11.5	23.5
21_Pronominal Possessor Noun	20.0	24.5	12.5	24.5
A 01_Attributive Adjective Agreement	9.0	24.5	13.5	24.5
A 02_Predicate Adjective Agreement	10.0	25.5	14.5	25.5
Neg 01_Standard Negation is Particle that Prece...	11.0	26.5	15.5	26.5
Neg 03_Standard Negation is Prefix	24.0	28.5	18.5	28.5
Neg 04_Standard Negation is Suffix	25.0	29.5	19.5	28.5
Neg 07_Standard Negation is Tone	26.0	30.5	20.5	29.5
Neg 08_Standard Negation is Tone plus Other Mod...	27.0	31.5	21.5	30.5
Neg 09_Standard Negation is Reduplication	28.0	32.5	22.5	31.5
Neg 10_Standard Negation is Infix	29.0	33.5	23.5	32.5
Neg 12_Distinct Negation of identity	30.0	34.5	24.5	33.5
Neg 13_Distinct Negation of Existence	31.0	35.5	25.5	34.5
Neg 14_Distinct Negation of Location	32.0	36.5	26.5	35.5
Order N3 01_ Demonstrative Adjective Noun	21.0	36.5	25.5	37.5

Figure 7: The table of influences with $k = 1$

	albanian	greek	hittite	wt_armenian
01_Subject Verb	26.0	86.166667	52.500000	209.0
06_Subject Object Verb	127.0	107.500000	53.500000	210.0
11_Adposition Noun_Phrase	28.0	88.166667	54.500000	211.0
12_Noun_Phrase Adposition	29.0	109.500000	55.500000	212.0
13_Adjective Noun	130.0	90.166667	56.500000	213.0
15_Numeral Noun	31.0	91.166667	57.500000	214.0
17_Demonstrative Noun	32.0	92.166667	58.500000	215.0
19_Possessor Noun	133.0	113.500000	59.500000	216.0
21_Pronominal Possessor Noun	134.0	114.500000	60.500000	217.0
A 01_Attributive Adjective Agreement	35.0	95.166667	61.500000	217.0
A 02_Predicate Adjective Agreement	36.0	96.166667	62.500000	218.0
Neg 01_Standard Negation is Particle that Prece...	37.0	97.166667	63.500000	219.0
Neg 03_Standard Negation is Prefix	138.0	118.500000	107.166667	221.0
Neg 04_Standard Negation is Suffix	139.0	119.500000	108.166667	221.0
Neg 07_Standard Negation is Tone	140.0	120.500000	109.166667	222.0
Neg 08_Standard Negation is Tone plus Other Mod...	141.0	121.500000	110.166667	223.0
Neg 09_Standard Negation is Reduplication	142.0	122.500000	111.166667	224.0
Neg 10_Standard Negation is Infix	143.0	123.500000	112.166667	225.0
Neg 12_Distinct Negation of identity	144.0	124.500000	113.166667	226.0
Neg 13_Distinct Negation of Existence	145.0	125.500000	114.166667	227.0
Neg 14_Distinct Negation of Location	146.0	126.500000	115.166667	228.0
Order N3 01_ Demonstrative Adjective Noun	47.0	107.166667	73.500000	230.0

Figure 8: The table of influences with $k \leq 2$

	albanian	greek	hittite	wt_armenian
01_Subject Verb	2255.066667	2386.733333	52.500000	1199.0
06_Subject Object Verb	1793.800000	1922.800000	53.500000	1200.0
11_Adposition Noun_Phrase	2257.066667	2388.733333	54.500000	1201.0
12_Noun_Phrase Adposition	2258.066667	1924.800000	55.500000	1202.0
13_Adjective Noun	1796.800000	2390.733333	56.500000	1203.0
15_Numeral Noun	2260.066667	2391.733333	57.500000	1204.0
17_Demonstrative Noun	2261.066667	2392.733333	58.500000	1205.0
19_Possessor Noun	1799.800000	1928.800000	59.500000	1206.0
21_Pronominal Possessor Noun	1800.800000	1929.800000	60.500000	1207.0
A 01_Attributive Adjective Agreement	2264.066667	2395.733333	61.500000	1252.0
A 02_Predicate Adjective Agreement	2265.066667	2396.733333	62.500000	1253.0
Neg 01_Standard Negation is Particle that Prece...	2266.066667	2397.733333	63.500000	1254.0
Neg 03_Standard Negation is Prefix	1804.800000	1933.800000	135.166667	1211.0
Neg 04_Standard Negation is Suffix	1805.800000	1934.800000	136.166667	1256.0
Neg 07_Standard Negation is Tone	1806.800000	1935.800000	137.166667	1257.0
Neg 08_Standard Negation is Tone plus Other Mod...	1807.800000	1936.800000	138.166667	1258.0
Neg 09_Standard Negation is Reduplication	1808.800000	1937.800000	139.166667	1259.0
Neg 10_Standard Negation is Infix	1809.800000	1938.800000	140.166667	1260.0
Neg 12_Distinct Negation of identity	1810.800000	1939.800000	141.166667	1261.0
Neg 13_Distinct Negation of Existence	1811.800000	1940.800000	142.166667	1262.0
Neg 14_Distinct Negation of Location	1812.800000	1941.800000	143.166667	1263.0
Order N3 01_ Demonstrative Adjective Noun	2276.066667	2407.733333	73.500000	1220.0

Figure 9: The table of influences with $k \leq 3$

Based on the influences that the syntactic structures carry, it's apparent that there are two classes of syntactic structures that are distinct in how much information they encode about the structure of the family.

5.5 Romance languages

For the romance language we obtain the following tree, which we have drawn with the root placed at Latin:

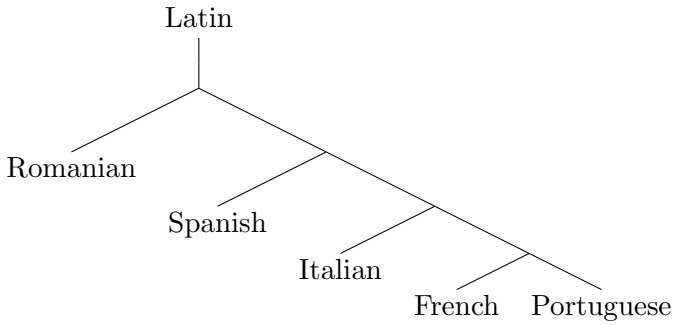

Figure 10: The logdet+nj Romance tree using 85 parameters from across SSWL and Longobardi completely set for the family

This tree misplaces Spanish. On moving Spanish in proximity of Portuguese, we obtain

the historically accepted tree:

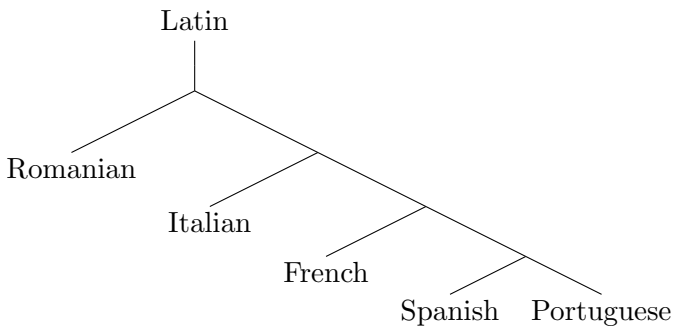

Figure 11: The historically accepted tree for the Romance family

The logdet matrix correctly shows that Spanish is indeed closest to Portugese. However, the raw logdet value is not what neighbor joining optimizes, so the expected (Spanish, Portugese) subtree does not emerge. Additionally, the logdet values involving Spanish and Portugese show high degree of asymmetry.

The similar misplacement of Spanish next to (Romanian, Latin) subtree also appears in Shu et al. (2018). We note that geographic proximity, the related history of Portugese and Spanish, and likely coevolution, makes it a candidate for violating the evolution on tree assumption: we expect these two to have much more similarity and interaction than possible on a tree.

	latin	romanian	italian	french	spanish	portuguese
latin	0.000000	0.418368	0.526778	0.686986	0.513258	0.561157
romanian	0.418368	0.000000	0.180652	0.237485	0.123650	0.152100
italian	0.526778	0.180652	0.000000	0.098833	0.047602	0.023741
french	0.686986	0.237485	0.098833	0.000000	0.100617	0.073974
spanish	0.513258	0.123650	0.047602	0.100617	0.000000	0.023861
portuguese	0.561157	0.152100	0.023741	0.073974	0.023861	0.000000

The LOGDET+NJ tree obtained is also not stable when attempting reconstruction after subsampling down to approximately 60% of the data. One of the topologies that consistently appears is in fact the correct topology. Computing phylogenetic invariants we find that the $\Phi^{\ell_\infty} \approx 0.00156$ invariant does not separate the two topologies, but Φ^{ℓ_1} does in fact separate them and surprisingly selects the LOGDET+NJ tree, taking a value ≈ 0.0094 as versus ≈ 0.0111 for the historically accepted tree.

5.5.1 Maximum likelihood model

Building the tree using the greedy scheme optimizing covariances, we obtain the following tree that misplaces French, Spanish and Italian.

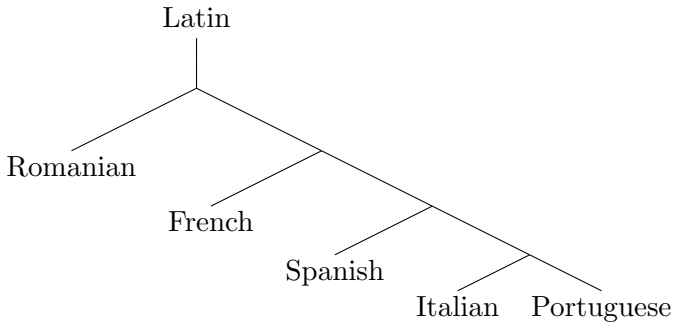

Figure 12: Tree based on covariance

With a comparatively larger set of available parameters, this family lends itself to a richer analysis along the lines we have sketched; we examine this family in more detail. We begin by constructing the maximum likelihood estimate for the tree parameter using the topology of the logdet tree, which would give the correct topology under the assumption that the data come from a general Markov model.

Starting with the matrix of data with one column for each of the $m = 85$ syntactic parameters, and each column containing the values the parameter takes at the leaves of the phylogenetic tree for the Romance family (L:Latin, R:Romanian, S:Spanish, I:Italian, F:French, P:Portuguese). Under the assumption that all parameters are independent, the likelihood of the data matrix, $Z = [z^1, z^2 \dots z^m]$, is the product of individual likelihoods of each column z^i,

$$L(Z) = \prod_{i \in [m]} L(z^i)$$

The likelihood for a column of data for the LOGDET+NJ tree model can be computed by assigning a probability distribution over $\{0, 1\}$ parameter values at the root, $\pi = [\pi_0, 1 - \pi_0]^T$, and 2-dimensional transition matrices,

$$M_k = \begin{bmatrix} \theta_k^0 & 1 - \theta_k^0 \\ 1 - \theta_k^1 & \theta_k^1 \end{bmatrix}$$

to each edge (the pendant edges are labelled by leaves, the interior edges by the left vertex) for $k \in \{L, R, S, I, F, P, i_R, i_S, i_I, i_L\}$ in the L rooted tree:

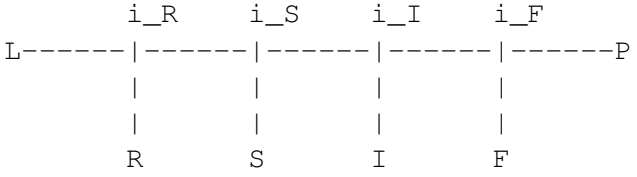

For a given vector $z \in \{0,1\}^5$, we are interested in the probability of seeing z at the leaves as a function of the parameters: $P(z) \equiv P(R = z_1, S = z_2 \dots P = z_5) = f(M_k, \pi)$. The likelihoods can be computed using dynamic programming but here we explicitly sum over the internal states at $i_l, l \in \{R, S, I, F\} \equiv \text{Int}$ and the state x at the root. For simplicity we use the identification $R = 1, S = 2, I = 3, F = 4, P = 5$ i.e. $M_R = M_1$ etc.

$$P(z) = \sum_{x \in \{0,1\}} \sum_{v \in \{0,1\}^4} P(z | i_k = v_k, k \in [4], x) P(v, x)$$

The probabilities at the leaves can be read off the transition matrices once the interior states are fixed (where we use independence). This gives:

$$P(z | i_k = v_k, k \in [4], x) = M_R[v_1, z_1] \cdot M_S[v_2, z_2] \cdot M_I[v_3, z_3] \cdot M_F[v_4, z_4] \cdot M_P[v_4, z_4]$$

All that remains is to estimate is $P(v, x)$. This follows similarly, since

$$P(v, x) = P(v_4|v_3)P(v_3|v_2)P(v_2|v_1)P(v_1, x) = M_{i_I}[v_3, v_4] \cdot M_{i_S}[v_2, v_3] \cdot M_{i_R}[v_1, v_2] \cdot P(v_1, x),$$

where we used that v_i only depends on v_{i-1}. Now $P(v_1, x) = P(v_1|x)P(x) = M_L[x, v_1]\pi[x]$, so that

$$P(z) = \sum_{x,v} \prod_{i \in [4]} M_i[v_i, z_i] \cdot M_P[v_4, z_4] \cdot M_{i_I}[v_3, v_4] \cdot M_{i_S}[v_2, v_3] \cdot M_{i_R}[v_1, v_2] \cdot M_L[x, v_1] \cdot \pi[x] \,.$$

This will give likelihood of the data as a polynomial in the model parameters after taking the product over all parameters. We maximize this using gradient descent to get the maximum likelihood estimate for the parameters, and simulated 10000 evolutions on this tree. Collecting into groups of 85 (which is the number of parameters used for the LOGDET+NJ tree), this gives approximately 110 trials. Since the simulated evolutions are independent, the distribution of distances between any two leafs for all trials are identical; their distribution approaches a Gaussian by the central limit theorem[5]. For the leaves i, j, we used the estimated standard deviation, σ^{ij}, and mean, μ^{ij}, from the simulated data to assign a z-score ($z^{ij} = (d^{ij} - \mu^{ij})/\sigma^{ij}$) to the distances seen in the actual data, d^{ij}. The distance between pairs aren't independent, so coalescing the data into single statistics is not straight forward, and as the the table of z-scores demonstrates, different parts of the tree behave differently; for example, any distances computed between the subgroup of (French, Portuguese) and rest of the family are overestimated by more than two standard deviations in the model, while other distances are within tolerance (one standard deviation of the mean). This gives support to the idea of there being more interaction in this family than tree models permit.

Using the the simulated data across 1000 simulations, we again build the tree based on covariance:

[5]This convergence can be quantified with with the Berry-Esseem theorem Durrett (2019)

	l1	l2	mean	std	z	d
I/P	I	P	2.78336	0.893444	-3.08875	0.0237408
F/I	F	I	2.87366	0.946365	-2.93209	0.0988331
F/S	F	S	2.85387	0.973022	-2.82959	0.100617
F/R	F	R	2.8915	0.997764	-2.65996	0.237485
P/R	P	R	2.90869	1.04035	-2.64968	0.1521
P/S	P	S	2.92124	1.12043	-2.58596	0.0238608
F/L	F	L	2.81107	1.01573	-2.09119	0.686986
L/P	L	P	2.84234	1.22816	-1.8574	0.561157
R/S	R	S	0.131869	0.0597361	-0.137579	0.12365
F/P	F	P	0.0765047	0.0436683	-0.0579418	0.0739745
I/L	I	L	0.520308	0.129903	0.0498074	0.526778
I/R	I	R	0.172632	0.0677208	0.11843	0.180652
L/R	L	R	0.40522	0.104813	0.125442	0.418368
I/S	I	S	0.0420906	0.033077	0.166614	0.0476017
L/S	L	S	0.479633	0.122479	0.274532	0.513258

Figure 13: z-score table for actual logdet distances compared to the maximum likelihood model

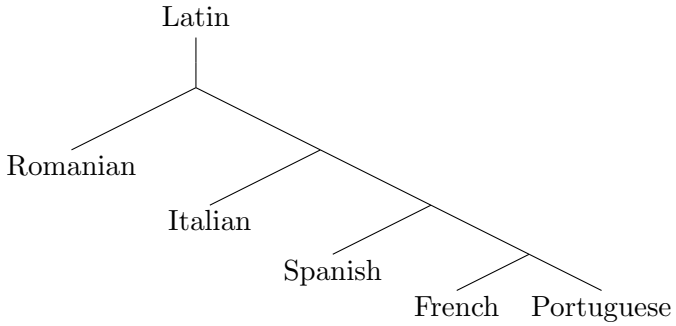

Figure 14: The covariance based tree using simulated data from maximum likelihood model, with the tree drawn with root at Latin.

Again this tree reports incorrectly the relative positions of French, Italian, and Spanish, with French and Spanish interchanged with respect to the tree considered historically correct. The fact that the greedy scheme places French closer to Portuguese is a significant misplacement because it represents that Portuguese and French have higher covariance than Spanish. A pattern showing that the newer Romance languages are often conflated in varied ways has emerged now.

These languages share a history that suggest that the limited interaction possible on

a tree fails to describe them. The linguistic consanguinity also makes it likely that the syntactic structures are not completely independent. This should be compared, for instance, with the analysis in Sections 4.4 and 4.5 of Port et al. (2019), where the dimensionality of the space of syntactic parameters is analyzed over different linguistic families, showing a drop in dimension in certain families, that corresponds to the presence of additional family-specific relations (see Figure 13 of Port et al. (2019) for the case of the Romance languages).

We move to consider both the sufficiency of the tree topology and independence and identical evolution assumption on the syntactic structures.

5.5.2 Independent and identical evolution assumption

French, Spanish and Portuguese because of their geographic proximity are likely to have experienced a higher degree of interaction, including at the syntactic level, than what normally expected in a Markov model on a tree (see our general discussion at the beginning of the paper on the model and its assumptions). It also appears that the syntactic structures within the Romance languages, especially when compared to Latin, are more highly correlated than what expected in general in terms of the dimensionality of the space of syntactic parameters (see Sections 4.4 and 4.5 of Port et al. (2019)). The more recent evolution of the more modern Romance languages also contributes to make their syntactic parameters less likely to behave like independent/identically evolving. Since we do have the simulated data from the maximum likelihood model, we can compare the simulated against the actual, and indeed, we find that this is exactly what we observe.

We start by randomly ordering the syntactic parameters in a given set S. Then for each triple of languages, $L = \{l_1, l_2, l_3\}$, we consider the probability that the syntactic parameter, s_i is set in all three, as well as the probability that the next one in the given ordering, s_{i+1}, is also set. We compute:

$$\frac{\Pr_S[\forall l \in L, s_i = 1]^2}{1/|S| + \Pr_S[\forall l \in L, s_i = 1 \wedge s_{i+1} = 1]}$$

The choice $|L| = 3$ is made because this is the largest size for which the size of the state space, 2^3 is "small" compared the the number of syntactic structures that are available.

We follow this by computing the same for the simulated data. Averaging over 100 random orderings of S, we consider the z-score for each triple of languages against the distribution from the simulated data. Since the simulated data come from iid evolutions on the maximum likelihood model, it gives the expected distribution for what the statistics for iid evolution looks like. This is summarized by the following table:

	z-score
F:S:P	0.222647
I:F:P	0.192141
I:F:S	0.199228
I:S:P	0.180108
L:F:P	0.102739
L:F:S	0.245780
L:I:F	0.223499
L:I:P	0.213800
L:I:S	0.036244
L:R:F	0.167095
L:R:I	-0.004172
L:R:P	0.149725
L:R:S	-0.004571
L:S:P	0.233173
R:F:P	0.159186
R:F:S	0.210917
R:I:F	0.206196
R:I:P	0.183321
R:I:S	0.041099
R:S:P	0.188342

In general, the statistics match what is expected from the simulated data, as it all lies within a quarter standard deviation. This implies that the statistics we will compute next, to test the adequacy of the tree topology, are largely unconfounded by the deviation from iid evolution.

5.5.3 Subfamily splits against maximum likelihood model

The placement of Portuguese and Spanish is one of the confounding factors that repeatedly appears. We consider the some partitions of the family where Portuguese and Spanish are separated, as well as where they form a cherry (while keeping together the consistent Romanian/Latin pair). If these partitions reflect genuine splits in the data, then the rank of the flattening for these partitions must be 2. The distance, d, to the closest rank 2 matrix to the flattening matrix can be computed as the norm of its singular values vector after excluding the top 2. We compute this for the simulated maximum likelihood data, d_{sim} where the logdet tree topology implies the Portuguese and Spanish separate. The z-score for the actual data d value, d_{actual}, against the background from the simulated data is tabulated:

	d_actual	mean[d_sim]	std_dev[d_sim]	z-score
F,P,S; L,R,I	0.012677	0.035968	0.013168	-1.768730
I,F,P,S; L,R	0.012585	0.031150	0.013502	-1.375018
I,F,P; L,S,R	0.012274	0.034641	0.013497	-1.657230
I,F,S; L,P,R	0.017000	0.220083	0.022537	-9.010975
P,S; F,L,R,I	0.011765	0.223664	0.023124	-9.163532

The table shows that the d value implied by the split from the logdet tree, [[L, R, S], [I, F, P]] lies one variance outside what is expected from the maximum likelihood model. And the [S, P] cherry comes with an order of magnitude more extreme z-score. This corresponds to data reflecting that Spanish and Portuguese are extremely likely to form a cherry than what is expected by the iid evolution on the tree, because the distance to the nearest rank 2 approximation is much smaller than what is obtained from the maximum likelihood model. Overall the data is at least one variance outside what is expected. The Portuguese–Spanish cherry statistics strongly suggests that the tree topology is not capturing the full range of interactions.

6 The logdet phylogenetic signal

The richness of the Ceolin et al. (2020) offers a way to test the scales at which the evolutionary models continue to be reliable beyond the setting of small scale examples considered previously. We compare the phylogenetic tree reconstructed using logdet/general Markov and modified Jaccard index/infinite site model, and discuss them in context of the results of Ceolin et al. (2020) using Bayesian phylogenetics (as implemented in BEAST software packageBouckaert et al. (2014)) as well as UPGMA clustering with modified Jaccard index.

Unlike the LanGeLin and SSWL dataset where we only use the parameters that are independent and known in all languages for which the construction is being carried out, for the Ceolin et al. (2020) dataset, we use all parameters, only restricting to independent parameters when computing pairwise distances. This means the distance between different pairs may be based on a different set of parameters. This is done following Ceolin et al. (2020), as otherwise we do not have enough parameters if we discard all that are not independent for any language. Under the assumption that all sites are i.i.d., this does not make a difference. For the reconstructed trees, because the linguistic evolutionary processes are not necessarily memoryless, the branch lengths are not meaningful; and we will reroot the trees with input from what is commonly agreed upon in linguistics literature.

6.1 Languages not included in SSWL and Longobardi datasets

We first consider language families that are not included in the SSWL and Longobardi datasets.

- The Indo-Aryan family represented in the Ceolin dataset by Marathi, Hindi, Telugu, Tamil, Pashto

  - Indo-Aryan reconstructed using Logdet metric:

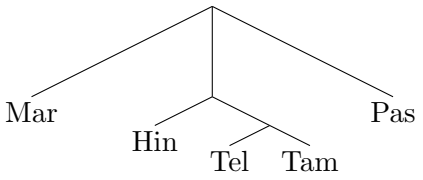

  - Indo-Aryan reconstructed using modified Jaccard metric:

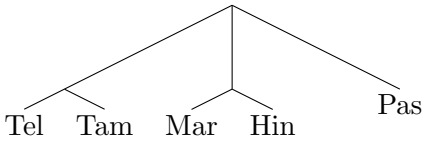

    The modified Jaccard index gives a more accurate reconstruction with the pairs Tamil/Telugu and Hindi/Marathi correctly identified, with the Tamil/Telugu forming the separate Dravidian subtree, identical to the results of Ceolin et al. (2020).

- East Asian languages, Korean, Japanese, Cantonese and Mandarin are correctly reconstructed using both and are in agreement with Ceolin et al. (2020)

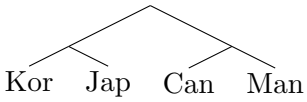

6.2 Greco-Romance Languages

Data for a superset of the Greco-Romance languages considered in section 5 is available in the Ceolin dataset. Restricting to just the Greco-Romance languages in the Ceolin data, both logdet and modified Jaccard similarity give identical trees; this was expected based on the high degree of asymmetry in maximum likelihood model. Portuguese and Spanish are now correctly placed. We do see a misplacement of two Italian dialects: Parma and Casalesco, but we note that this misplacement is also present in the BEAST trees of Ceolin et al. (2020), and that it disappears when we consider the full Indo-European family instead of just the Greco-Romance family. This suggests that this is possibly arising from the large sampling bias in this set which contains a disproportionate number of closely related Italian dialects

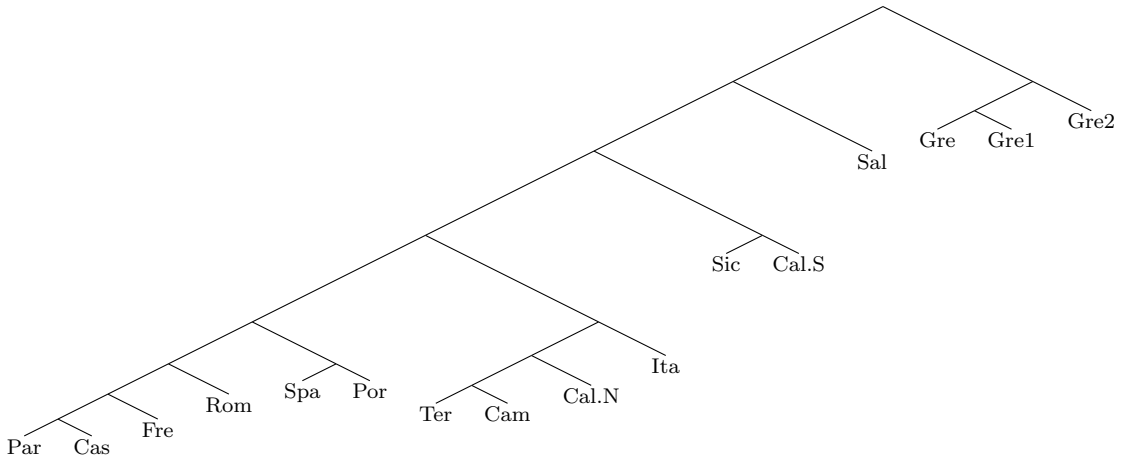

Figure 15: The Greco-Romance tree obtained from LOGDET+NJ construction that is identical to the modified Jaccard index tree when restricted to the Greco-Romance languages

This tree carries some misplacements; in view of these inaccuracies, we are led to conclude that either sampling biases or failure of assumptions underlying the models are significant; for instance one notes that Greek and Romance families are quite distinct and one cannot suppose that an hypothetical root from which these evolve can be recovered from the data. We will reconsider the Greco-Romance languages in context of the full Indo-European family in section 6.5.

As a final check, we revisit the SSWL-LanGeLin data for the Romance family from section 5.5, and apply the modified Jaccard index construction to see if the change in the model resolves the persistent issues there. We find that the reconstruction is identical, with modified Jaccard index values similar to the logdet values.

	french	italian	latin	portuguese	romanian	spanish
french	0.000000	0.098833	0.686986	0.073974	0.237485	0.100617
italian	0.098833	0.000000	0.526778	0.023741	0.180652	0.047602
latin	0.686986	0.526778	0.000000	0.561157	0.418368	0.513258
portuguese	0.073974	0.023741	0.561157	0.000000	0.152100	0.023861
romanian	0.237485	0.180652	0.418368	0.152100	0.000000	0.123650
spanish	0.100617	0.047602	0.513258	0.023861	0.123650	0.000000

Table 1: Table for logdet metric

	french	italian	latin	portuguese	romanian	spanish
french	0.000000	0.100000	0.415094	0.076923	0.209302	0.102564
italian	0.100000	0.000000	0.346154	0.025641	0.162791	0.051282
latin	0.415094	0.346154	0.000000	0.365385	0.294118	0.352941
portuguese	0.076923	0.025641	0.365385	0.000000	0.142857	0.026316
romanian	0.209302	0.162791	0.294118	0.142857	0.000000	0.121951
spanish	0.102564	0.051282	0.352941	0.026316	0.121951	0.000000

Table 2: Table for modified Jaccard index

6.3 Germanic Langauages

For the Germanic family, the Ceolin data contains the additional North Germanic languages Danish and Norwegian, and is missing Swedish. Using Icelandic as the outgroup to root the rest we correctly recover the North/West split with both logdet and modified Jaccard. The interior structures defer, with modified Jaccard placing Faroese with Danish, where the expected would is Norwegian, and lodget fails to assign Dutch and Afrikaans together (same as UPGMA from Ceolin et al. (2020) - which also fails to recover west and north Germanic split, while Ceolin et al's BEAST reconstruction places German next to Afrikaans).

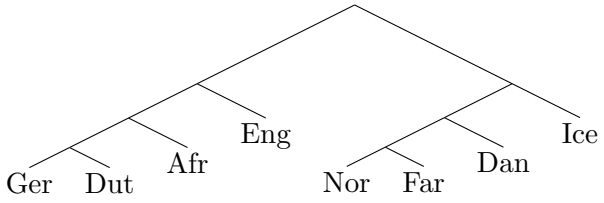

Figure 16: logdet+nj construction

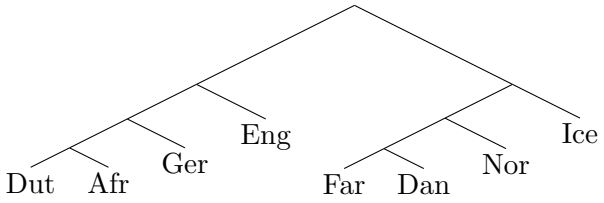

Figure 17: Modified Jaccard construction

6.4 Balto-Finnic, Urgic and Altaic

For the Balto-Finnic, Urgic and Altaic languages the logdet and modified Jaccard both give very similar structures; with some disagreement between the two between the two pairs of closely related dialects of Mari and Udmurt. Similar conflation is observed in the BEAST tree from Ceolin et al. (2020). Compared to the LanGeLin LOGDET+NJ construction, we correctly recover the relationship between Khanty and Hungarian, and that between Turkish and Yakut. We recover the relations between Evenki, Uzbekh, Yakut and Turkish that are the same as Ceolin et al's as well. The modified Jaccard tree is provides a better reconstruction as it separates the three Turkic languages – Uzbek, Yakut and Turkish into their own subtree within the Altaic languages.

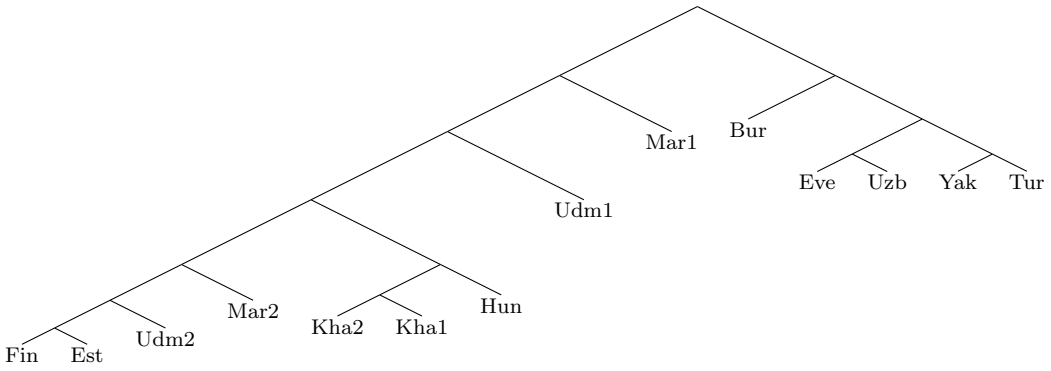

Figure 18: LOGDET+NJ construction

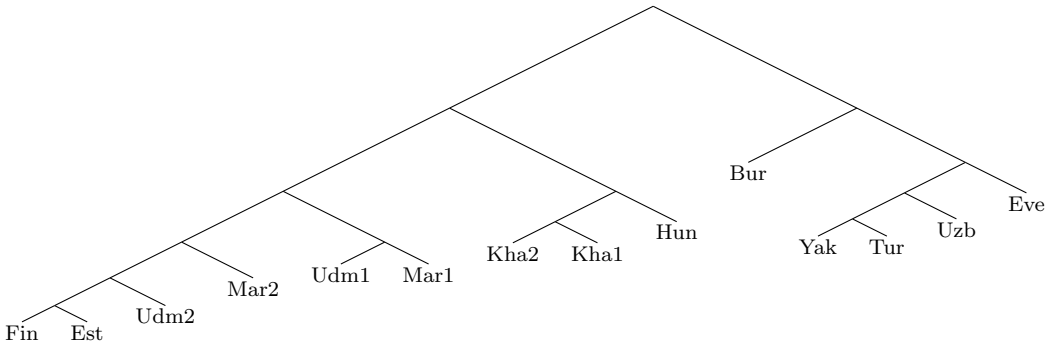

Figure 19: Modified Jaccard construction

6.5 The full Indo-European family

In the analysis of the full Indo-European tree we see that the modified Jaccard index slightly outperforms the LOGDET+NJ construction. Both methods misplace Welsh (note Irish and

Welsh were degenerate in this dataset, so we retained only Welsh; Ceolin et al kept both and so obtained that Welsh and Irish place together), but LOGDET+NJ now loses the west and north Germanic split, while the modified Jaccard index is more stable.

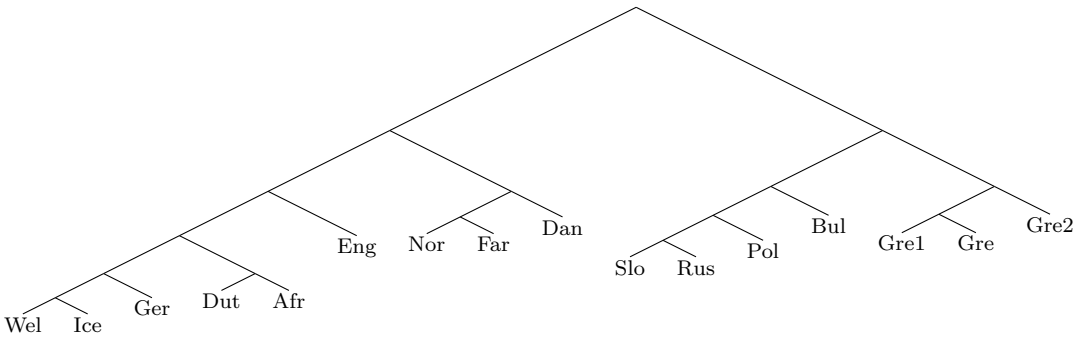

Figure 20: LOGDET+NJ :Germanic, Slavic and Greek languages from the full Indo-European languages set

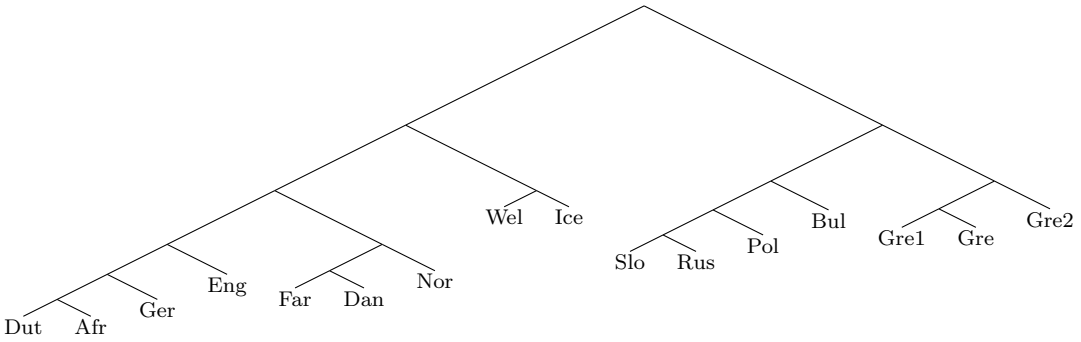

Figure 21: Modified Jaccard:Germanic, Slavic and Greek languages from the full Indo-European languages set

The Romance subtrees now have correct large scale structure with Italian dialects forming their own group. The Italian-Portuguese-Spanish conflation that we saw in SSWL-LanGeLin analysis becomes clearer given how these three separate out.

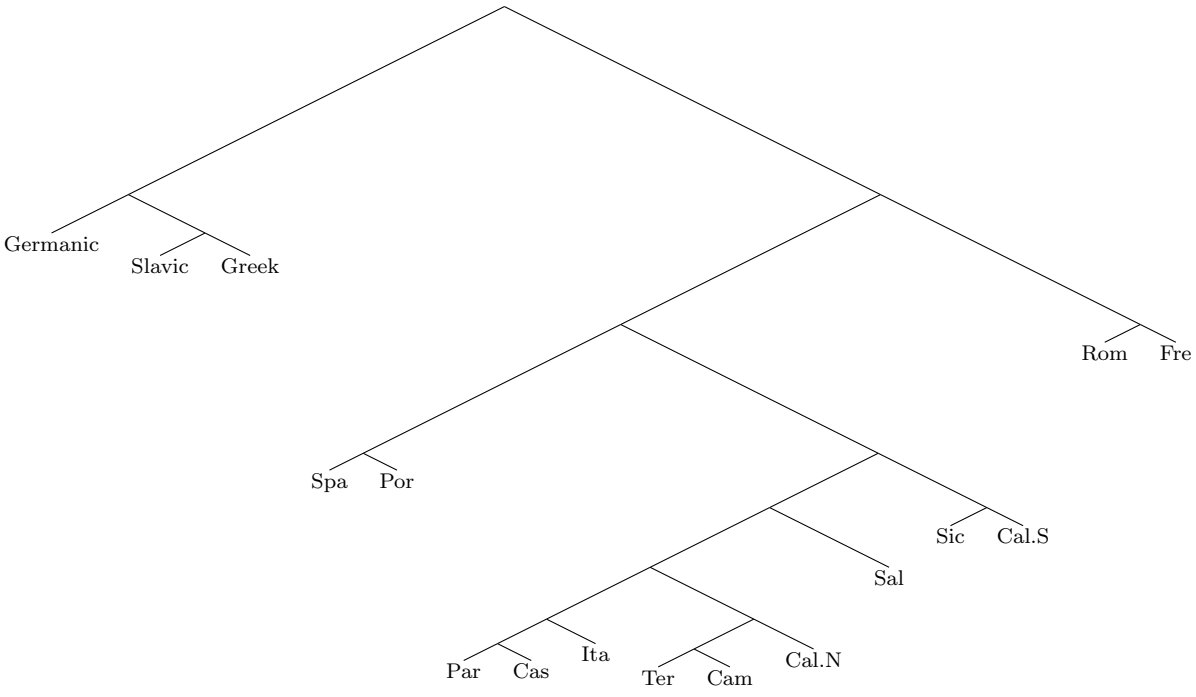

Figure 22: LOGDET+NJ :Germanic, Slavic and Greek languages from the full Indo-European languages set

We see that modified Jaccard and LOGDET+NJ both correctly recover that Indo-Aryan languages split off from the European ones, unlike Ceolin et al. (2020) where they are mixed between the Greek and Romance subtree.

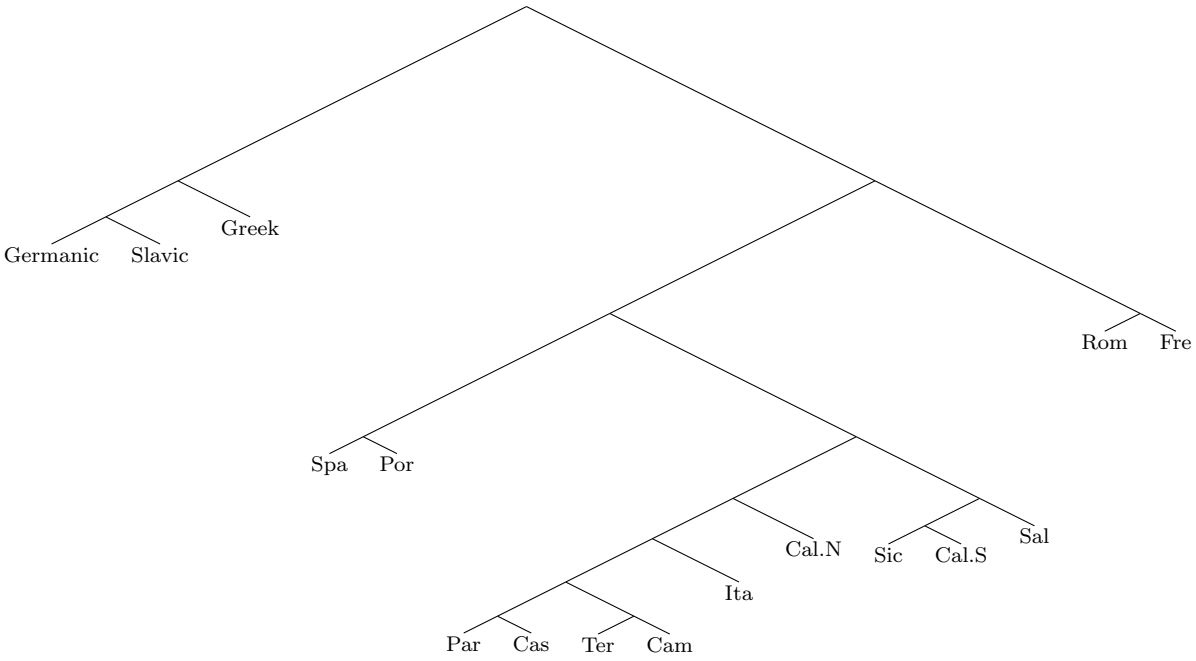

Figure 23: Modified Jaccard:Romance languages from the full Indo-European languages set; the Italian dialects are now correctly forming their own subtree

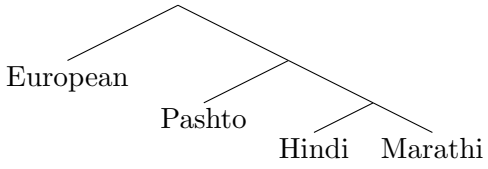

Figure 24: Both modified Jaccard and logdet place the Indo-Aryan languages identically in relation to the European languages

Putting this together we note that both the LOGDET+NJ and modified Jaccard index/infinite site model recover the large scale structure in Indo-European languages. Modified Jaccard index with neighbor joining is more stable than logdet – the asymmetry in how syntactic structure change across language families (and possibly how they evolved) is hardwired into the modified Jaccard metric, while lodget/general Markov model because of its flexibility may overfit. These two methods of reconstruction, reflecting different evolutionary models, are mostly in agreement, disagreeing on finer scale structures where effects of deviations from the assumptions of underlying models affect each differently and to varying significance.

7 Conclusion

This analysis was aimed towards trying to understand how well the general Markov model describes the syntactic structures data to get insight into how human languages change. The point was not to derive a metric, possibly abstract, that yields the expected phylogenetic relationships but to understand how well the phylogenetic relationship can be modelled by a type of process we do understand well: the general Markov process. In doing so we run into expected roadblocks:

- Linguistic relationships across multiple families are often not stably reconstructed.
- Languages with high degree of relatedness are difficult to place.

The first can be addressed by noting that the hypothesis of a single root from which they can be considered to have evolved may be accurate for linguistic subfamilies, but ancestral languages and proto-languages lying behind sufficiently different linguistic families are a highly hypothetical (the contested Ural-Altaic hypothesis being one such example), hence simply trying to fit diverse syntactic data across a large range of language families to such a model with a single root should not be expected to be very meaningful.

The second is indicative of two problems – the tree topology is insufficient to capture how they have influenced each other, but it also represents a type of sampling problem: the representatives of language families are not generated as random samples from the process acting on the family; with a single close relationship in a family, coupled with the small size of the families, the biases become extremely significant. An algorithm, like neighbor joining, that uses both local information (the pairwise distances) and the global information (pairwise distances to the rest of the tree), is likely to be thrown off by these biases. This is evident from the example where while Spanish and Portuguese share the highest similarity in terms of both the modified jaccard and logdet metric, their placement does not reflect this. This is also supported from observation that larger sets from with in the same family tend to give more correct picture of the phylogenetic relationships, even though there is a trade off that the larger collection may be less likely to be described by a simple model.

This leads us to conclude that phylogenetic inference at larges scale across many families in an unsupervised way and without linguistic and historical context is more likely to be an abstract exercise than an approximation of the underlying truth. We suggest that care be taken when working at different scale, for example, inter versus intra linguistic family scale.

We especially think that more significant theoretical work is needed on dynamical models of language change at the syntactic level, that can replace the Markov model with a more accurate model, tailored to linguistic needs, that can be used for better phylogenetic inference of relevance to historical linguistics.

References

E. Allman and J. Rhodes. Phylogenetic ideals and varieties for general markov models. *Adv. Appl. Math.*, 40:127–148, 2008.

Elizabeth S Allman, John A Rhodes, and Seth Sullivant. When do phylogenetic mixture models mimic other phylogenetic models? *Systematic biology*, 61(6):1049–1059, 2012.

Mark C. Baker. *The Atoms of Language*. Basic Books, 2002.

Robert Berwick. Mind the gap. In *50 Years Later: Reflections on Chomsky's Aspects*, pages 1–12. MIT Press, 2015.

Remco Bouckaert, Philippe Lemey, Michael Dunn, Simon J Greenhill, Alexander V Alekseyenko, Alexei J Drummond, Russell D Gray, Marc A Suchard, and Quentin D Atkinson. Mapping the origins and expansion of the indo-european language family. *Science*, 337 (6097):957–960, 2012.

Remco Bouckaert, Joseph Heled, Denise Kühnert, Tim Vaughan, Chieh-Hsi Wu, Dong Xie, Marc A Suchard, Andrew Rambaut, and Alexei J Drummond. Beast 2: a software platform for bayesian evolutionary analysis. *PLoS Comput Biol*, 10(4):e1003537, 2014.

J.S. Bridle and L. Dodd. Formal grammars and Markov models, 1987.

Andrea Ceolin, Cristina Guardiano, Monica Alexandrina Irimia, and Giuseppe Longobardi. Formal syntax and deep history. *Frontiers in psychology*, 11:2384, 2020.

Andrea Ceolin, Cristina Guardiano, Giuseppe Longobardi, Monica Alexandrina Irimia, Luca Bortolussi, and Andrea Sgarro. At the boundaries of syntactic prehistory. *Philosophical Transactions of the Royal Society B*, 376(1824):20200197, 2021.

N. Chomsky. Three models for the description of language. *IRE Transactions on Information Theory*, 2(3):113–124, 1956. doi: 10.1109/TIT.1956.1056813.

Noam Chomsky. *Lectures on Government and Binding*. Walter de Gruyter, 1981.

Noam Chomsky and Howard Lasnik. The theory of principles and parameters. In *Syntax: An international handbook of contemporary research*, pages 506–569. Walter de Gruyter, 1993.

Chris Collins. Syntactic structures of the world's language: A cross-linguistic database. 2010. 27 September 2010, Colloquium: `https://ling.yale.edu/syntactic-structures-worlds-language-cross-linguistic-database`.

Matthew S. Dryer and Martin Haspelmath, editors. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013. URL `https://wals.info/`.

Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.

Nicholas Karl Eriksson. *Algebraic combinatorics for computational biology*. PhD thesis, University of California, Berkeley, 2006.

Joseph Felsenstein. *Inferring phylogenies*, volume 2. Sinauer associates Sunderland, MA, 2004.

Olivier Gascuel and Mike Steel. Neighbor-joining revealed. *Molecular biology and evolution*, 23(11):1997–2000, 2006.

C. Guardiano, D. Michelioudakis, A. Ceolin, M. Irimia, G. Longobardi, N. Radkevich, I. Sitaridou, and G. Silvestri. South by southeast. a syntactic approach to greek and romance microvariation. *L'Italia Dialettale*, 77:95–166, 2016.

S. Karimi and M. Piattelli-Palmarini. Special issue on parameters. *Linguistic Analysis*, 41 (3-4), 2017.

Hilda Koopman. Sswl syntactic structures of the world's languages: An open-ended database for the linguistic community and by the linguistic community `http://sswl.railsplayground.net/` mit 50, 12 2011.

James A Lake. Reconstructing evolutionary trees from dna and protein sequences: paralinear distances. *Proceedings of the National Academy of Sciences*, 91(4):1455–1459, 1994.

David Link. Traces of the mouth: Andrei andreyevich markov's mathematization of writing. *History of Science*, 44(3):321–348, 2006. doi: 10.1177/007327530604400302. URL `https://doi.org/10.1177/007327530604400302`.

Giuseppe Longobardi. Convergence in parametric phylogenies: Homoplasy or principled explanation? In *Parameter Theory and Language Change*, pages 304–319. Oxford University Press, 2012.

Giuseppe Longobardi. Principles, parameters, and schemata: a constructivist UG. *Linguistic Analysis*, 41(3-4):517–556, 2017a.

Giuseppe Longobardi. Principles, parameters, and schemata. a constructivist ug. *Linguistic Analysis*, (41):517–557, 2017b. `https://www.york.ac.uk/language/research/projects/langelin/#tab-1`.

Giuseppe Longobardi and Cristina Guardiano. Evidence for syntax as a signal of historical relatedness. *Lingua*, 119:1679–1706, 2009.

Giuseppe Longobardi, Cristina Guardiano, Giuseppina Silvestri, Alessio Boattini, and Andrea Ceolin. Toward a syntactic phylogeny of modern indo-european languages. *Journal of Historical Linguistics*, 3(1):122–152, 2013.

Matilde Marcolli. Syntactic parameters and a coding theory perspective on entropy and complexity of language families. *Entropy*, 18(4):Paper No. 110, 17, 2016. ISSN 1099-4300. doi: 10.3390/e18040110. URL https://doi-org.clsproxy.library.caltech.edu/10.3390/e18040110.

Frederick A Matsen and Mike Steel. Phylogenetic mixtures on a single tree can mimic a tree of another topology. *Systematic Biology*, 56(5):767–775, 2007.

Yugo Murawaki. Analyzing correlated evolution of multiple features using latent representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4371–4382, 2018.

Partha Niyogi. *The computational nature of language learning and evolution*, volume 43 of *Current Studies in Linguistics*. MIT Press, Cambridge, MA, 2006. ISBN 0-262-14094-2.

Partha Niyogi and Robert C. Berwick. A dynamical systems model for language change. *Complex Systems*, pages 161–204, 1997.

Diana Nurbakova, Sergey Rusakov, and Vassil Alexandrov. Quantifying uncertainty in phylogenetic studies of the slavonic languages. *Procedia Computer Science*, 18:2269–2277, 2013.

Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

Andrew Ortegaray, Robert C. Berwick, and Matilde Marcolli. Heat kernel analysis of syntactic structures. *CoRR*, abs/1803.09832, 2018. URL http://arxiv.org/abs/1803.09832.

Lior Pachter and Bernd Sturmfels. *Algebraic statistics for computational biology*, volume 13. Cambridge university press, 2005.

Lior Pachter and Bernd Sturmfels. The mathematics of phylogenomics. *SIAM Review*, 49 (1):3–31, 2007.

J.J. Park, R. Boettcher, A. Zhao, A. Mun, K. Yuh, V. Kumar, and M. Marcolli. Prevalence and recoverability of syntactic parameters in sparse distributed memories. In *Geometric Structures of Information 2017. Lecture Notes in Computer Science, Vol. 10589*, pages 1–8. Springer, 2017.

A. Perelysvaig and M.W. Lewis. *The Indo-European controversy: facts and fallacies in Historical Linguistics*. Cambridge University Press, 2015.

Peter Piispanen. The uralic-yukaghiric connection revisited: Sound correspondences of geminate clusters. *Suomalais-Ugrilaisen Seuran Aikakauskirja*, 2013(94):165–197, tammi 2013. doi: 10.33340/susa.82515. URL https://journal.fi/susa/article/view/82515.

Alexander Port, Iulia Gheorghita, Daniel Guth, John M. Clark, Crystal Liang, Shival Dasu, and Matilde Marcolli. Persistent topology of syntax. *Mathematics in Computer Science*, 12(1):33–50, Mar 2018. ISSN 1661-8289. doi: 10.1007/s11786-017-0329-x. URL https://doi.org/10.1007/s11786-017-0329-x.

Alexander Port, Taelin Karidi, and Matilde Marcolli. Topological analysis of syntactic structures. *CoRR*, abs/1903.05181, 2019. URL http://arxiv.org/abs/1903.05181.

Kateřina Rexová, Daniel Frynta, and Jan Zrzavỳ. Cladistic analysis of languages: Indo-european classification based on lexicostatistical data. *Cladistics*, 19(2):120–127, 2003.

D. Ringe, T. Warnow, and A. Taylor. Indo-european and computational cladistics. *Transactions of the Philological Society*, 100:59–129, 2002.

Luigi Rizzi. On the format and locus of parameters: the role of morphosyntactic features. *Linguistic Analysis*, 41(3-4):159–191, 2017.

Charles Semple, Mike Steel, et al. *Phylogenetics*, volume 24. Oxford University Press on Demand, 2003.

Kevin Shu and Matilde Marcolli. Syntactic structures and code parameters. *Math. Comput. Sci.*, 11(1):79–90, 2017. ISSN 1661-8270. doi: 10.1007/s11786-017-0298-0. URL https://doi-org.clsproxy.library.caltech.edu/10.1007/s11786-017-0298-0.

Kevin Shu, Andrew Ortegaray, Robert C. Berwick, and Matilde Marcolli. Phylogenetics of indo-european language families via an algebro-geometric analysis of their syntactic structures. *CoRR*, abs/1712.01719, 2017. URL http://arxiv.org/abs/1712.01719.

Kevin Shu, Sharjeel Aziz, Vy-Luan Huynh, David Warrick, and Matilde Marcolli. Syntactic phylogenetic trees. In *Foundations of mathematics and physics one century after Hilbert*, pages 417–441. Springer, Cham, 2018.

Daniel Štefankovič and Eric Vigoda. Phylogeny of mixture models: Robustness of maximum likelihood and non-identifiable distributions. *Journal of Computational Biology*, 14(2):156–189, 2007a.

Daniel Štefankovič and Eric Vigoda. Pitfalls of heterogeneous processes for phylogenetic reconstruction. *Systematic biology*, 56(1):113–124, 2007b.

Patrick S Stumpf, Rosanna CG Smith, Michael Lenz, Andreas Schuppert, Franz-Josef Müller, Ann Babtie, Thalia E Chan, Michael PH Stumpf, Colin P Please, Sam D Howison, et al. Stem cell differentiation as a non-markov stochastic process. *Cell Systems*, 5(3):268–282, 2017.

Tandy Warnow. *Computational Phylogenetics*. Cambridge University Press, 2017.