

Sparse Logistic Regression with High-order Features for Automatic Grammar Rule Extraction from Treebanks

Santiago Herrera, Caio Corro, Sylvain Kahane

▶ To cite this version:

Santiago Herrera, Caio Corro, Sylvain Kahane. Sparse Logistic Regression with High-order Features for Automatic Grammar Rule Extraction from Treebanks. Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024), May 2024, Turin, Italy. 10.48550/arXiv.2403.17534. hal-04701105

HAL Id: hal-04701105 https://hal.science/hal-04701105v1

Submitted on 18 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Sparse Logistic Regression with High-order Features for Automatic Grammar Rule Extraction from Treebanks

Santiago Herrera¹, Caio Corro², Sylvain Kahane^{1,3}

¹Université Paris Nanterre, CNRS, Modyco ²Sorbonne Université, CNRS, ISIR ³Institut Universitaire de France sherrera@parisnanterre.fr. caio.corro@isir.upmc.fr. svlvain@kahane.fr

errera@parishamerre.ir, calo.corro@isir.upinc.ir, sylvain@kanal

Abstract

Descriptive grammars are highly valuable, but writing them is time-consuming and difficult. Furthermore, while linguists typically use corpora to create them, grammar descriptions often lack quantitative data. As for formal grammars, they can be challenging to interpret. In this paper, we propose a new method to extract and explore significant fine-grained grammar patterns and potential syntactic grammar rules from treebanks, in order to create an easy-to-understand corpus-based grammar. More specifically, we extract descriptions and rules across different languages for two linguistic phenomena, agreement and word order, using a large search space and paying special attention to the ranking order of the extracted rules. For that, we use a linear classifier to extract the most salient features that predict the linguistic phenomena under study. We associate statistical information to each rule, and we compare the ranking of the model's results to those of other quantitative and statistical measures. Our method captures both well-known and less well-known significant grammar rules in Spanish, French, and Wolof.

Keywords: grammar extraction, grammar rules, corpus based grammar, quantitative grammar.

1. Introduction

Speakers are obliged to select, pronounce, use, mark, and combine linguistic units in precise ways in order to communicate in a language. A grammar is this set of constraints that a language imposes on its speakers, *i.e.* a set of constraints satisfied by all well-formed utterances. In this setting, a grammar rule describes a specific linguistic pattern enforced in a given context and in a given language, and can be of various interest (e.g. education, theoretical linguistics, linguistic typology). In practice, these rules are of probabilistic nature (in the frequentist interpretation). For example, in English, a simple rule would be that the "subject of a verb is before its governor". However, this rule is not deterministic, a more accurate description would be: "given a syntactic dependency of type subject, the dependent is before its governor in approximately 93% of cases". In practice, it is of interest to understand when this majority rule is not followed.¹ To continue with the same example: "given a syntactic dependency of type subject with an expletive complement in a pre-verbal position, the subject is after its governor in 97% of cases", see Figure 1. As such, rules are not only probabilistic but also more or less fine-grained and potentially overlapping.

With the development of large treebank collec-

tions like Universal Dependencies, there has been a recent interest in automatically extracting grammars from syntactically annotated sentences. Importantly, previous work by Chaudhary et al. (2020, 2022) proposes to rely on machine learning techniques to this end. However, their approach suffers from some limitations, which we overcome in our work: (1) their syntactic rules do not capture fine-grained phenomena due to an under-specified search space and the use of decision trees (e.g. their model capture the fact that the subject inversion can be due to the presence of 'there', but not by the presence of an expletive in general, missing the example of Figure 1); (2) the use of decision trees (Quinlan, 1986) requires tuning of several hyperparameters (entropy choice, depth, splitting constraint, etc.), which is difficult to do on held-out data due to the frequentist nature of predictions; (3) a decision tree does not produce a ranked list of rules by saliency.

In this paper, we propose a different method to extract and explore more fine-grained grammar rules from syntactic treebanks in order to automatically create an easy-to-understand corpusbased grammar. More specifically, we extract descriptions across languages for two linguistic phenomena, agreement and word order, paying special attention to the ranking order of the extracted rules. To that end, we use a linear classifier trained with a sparsity inducing regularizer (Bach et al., 2012). The strength of the regularization, denoted $\lambda \in \mathbb{R}_+$, directly impacts the number of extracted rules for the linguistic phenomena under study. In this setting, the regularization path is the se-

¹The value also depends on the annotation scheme and on the corpus considered. Here the results are given for SUD_English-GUM@2.13, converted from UD in SUD (Gerdes et al., 2018), where the subject depends on the auxiliary in complex verbal forms.



Figure 1: Example from the SUD version of the English GUM treebank. Two clauses with subjects in different positions. In the main clause, the subject follows the verb and the first position is filled by an expletive. In the subordinate clause, the subject occupies the dominant pre-verbal position.

quence of solutions to the training problem when λ varies (Markowitz, 1952; Osborne et al., 2000; Efron et al., 2004). The earlier a rule appears in the path when decreasing λ , the salient it is for the linguistic phenomena under study. Another important benefit of our approach compared to decision trees is that it is (almost) hyperparameter free, which is a crucial feature as hyperparameter selection is non-trivial (and even hopeless), see Section 4.3.

Our contributions can be summarized as follows:

- we propose a novel formalization of grammatical rules with the aim of automatic extraction from treebanks;
- compared to previous work, we broaden the expressiveness of patterns to allow for more precise rules;
- we investigate the use of sparse logistic regression to extract and rank rules;
- we conduct qualitative evaluation on three languages (Spanish, French and Wolof).

All in all, we contribute to the development of descriptive grammars that, while time-consuming to write, are fundamental resources for theoretical linguistic, language-specific or cross-linguistic studies, and an essential resource for language documentation. Our goal is to contribute to the brisdging of the gap between computational linguistics (CL) and theoretical linguistics (Baroni, 2021), as well as between CL and field linguistics.

2. Related Works

Formal grammars (FG). There exists a large amount of works on extracting FGs from treebanks, including CFG (Charniak, 1996), TAG (Chen et al., 2006; Bhatt et al., 2012), LFG (Frank et al., 2003; Burke et al., 2004; Rehbein, 2009) and HPSG (Cahill et al., 2005; Zhang et al., 2012) grammars. FGs are rich and complete descriptions of structures observed in treebanks, usually intended for

parsing. However, they are also redundant and difficult to interpret as they are not descriptive grammars (Zaefferer et al., 1992) and because of their very large size.² Moreover, these FG have limitations when it comes to incorporating quantitative information and highlighting salient features of corpora as they are based on strong symbolic formalisms.

Typologicial features induction. Although linguists still rely on databases like WALS (Dryer and Haspelmath, 2013) and GramBank (Skirgård et al., 2023) for typological features, it is becoming increasingly important to infer these features using corpora, see Ponti et al. (2019) for motivations. Importantly, there has been a recent interest on gradient and quantitative feature descriptions (Levshina, 2022; Baylor et al., 2024) instead of the standard categorical setting. For examples, previous work have considered formal grammars parametrized by categorical typological features, parameters which are inferred from various types of corpora, including interlinear glosses (Bender et al., 2014; Howell and Bender, 2022). Similarly, Choi et al. (2021a,b) used manually crafted rules to extract typological statistics from treebanks, whereas Alves et al. (2023) also relied on information from an automatically extracted CFG. Contrary to these works, our aim is to provide more fine-grained and quantitative grammar rules while requiring less manual pattern crafting.

Grammar rule extraction. Closer to our work, Chaudhary et al. (2020, 2022) demonstrate the ability to extract human-readable grammar rules using machine learning techniques. The authors focus on a few specific linguistic phenomena, including agreement and word order, that are cast as classification tasks. For example, they investigate the subset of features (encoded in governordependent pairs) used by a decision tree to predict if there is a number agreement between two dependent words. Finally, Blache et al. (2016) investigate the use of an automatically extracted CFG

 $^{^{2}\}mbox{Charniak's}$ (1996) English grammar contains $15\,953$ rules.

Treebank	Sentences	Tokens
English-GUM	10 k	187 k
French-GSD	16 k	400 k
Spanish-AnCora	17 k	559 k
Wolof-WTB	2 k	44 k

Table 1: Treebank sizes.

with the same goal. Their approach also suffers from the expressiveness issue due to the locality of CFG's rewriting rules.

3. Grammatical Formalism

3.1. Syntactic Dependencies

Although our method for extracting rules does not depend on a particular formalism, we briefly present here the data to which the method has been applied, that is, the Universal Dependencies project (UD, De Marneffe et al., 2021). It is an open project on cross-linguistic syntactic annotation, which has 259 treebanks for 148 different languages in its last available release (UD 2.13). It is based on a consistent framework for multilingual annotation of parts of speech, lemmas, morphological features and syntactic dependencies, as well as guidelines for lemmatization, tokenization, and other tasks.

UD defines content words as syntactic heads in order to make annotations more seamless between different languages. While this favors comparability between treebanks, some morphosyntactic properties remain "less accessible" and syntactic differences less evident. Given that extracted rules depends on the annotation scheme, we instead use the Surface-Syntactic UD (SUD) framework (Gerdes et al., 2018, 2022), which defines syntactic heads by distributional criteria. This is more suitable to extract word order and agreement rules, amongst others. For example, in SUD, contrary to UD, auxiliaries are analyzed as the head of the clause and the subject depends on it. In English (and more generally in Indo-European languages), the subject is positioned towards the auxiliary, with which it agrees, rather than the verb. Moreover, this allows as to compare our work to previous work (Chaudhary et al., 2020, 2022) that also use SUD.

In order to demonstrate how our method works, we have chosen to analyze the results on languages in which we are experts. This allows us to conduct precise qualitative analysis of results, which is in our opinion more important that large scale evaluation via crowdsourcing native speakers.³ In section 5, we present rules extracted



Figure 2: The search space is defined around a governor node (GOV) and a dependent node (DEP), including the grandparent's governor (GRANPARENT) and the governor's other dependents (SIBLINGS), as well as the dependent's children (GRANDCHILDREN).

from SUD_Spanish-AnCora (Taulé et al., 2008), SUD_French-GSD (Guillaume et al., 2019), and SUD_Wolof-WTB (Dione, 2019) treebanks. The English GUM treebank (Zeldes, 2017) was also used to extract examples to illustrate our approach. Statistics are reported in Table 1.

3.2. Grammatical Rules

As stated, a grammar rule is a regular constraint in a language system occurring in a particular context and involving particular subsets of patterns. Importantly, these rules are probabilistic. In order to build a system capable to extract such rules, we need first to formalize what is a grammatical rule.

In practice, the interpretation of a rule ultimately depends on the theoretical framework adopted. A framework, in a broad sense, includes the grammatical formalism (*e.g.* dependency grammar in our case) and the annotation scheme (*e.g.* UD, SUD, etc.) How the linguistic hypothesis and question are formulated, and which grammatical features are used to answer to the linguistic question also depend on this. Our proposed rule extraction formalization is easy to generalise to any theoretical framework, independent of the rule extraction method as well as it is agnostic to the language treebank used.

We first give an intuitive definition of what we call a syntactic grammar rule in this work, based on the example of the subject position in Figure 1:

- 1. Firstly, we need to define what we are looking at, *e.g.* all dependency relations or a specific subset of them? The scope S of the rule is a given pattern and we consider all dependencies satisfying S. In the example, this is all relations of type subject.
- Secondly, we need to define what the linguistic phenomenon of interest is. In other words, we seek to identify what triggers satisfaction of response pattern Q amongst all relations that satisfies S. In the example, we are interested in subject inversion (subjects that are after their governor in English).

their tasks using large language models (Veselovsky et al., 2023), which make these evaluations less reliable.

³There are evidences that crowd workers automatize

3. Third and finally, the last element of a rule is the pattern P that acts as a trigger of Q in the scope S. In the example, P would be the fact that the subject has an expletive codependent.

Definition 1 (Syntactic grammar rule). Within all dependencies that match a pattern S (the scope), a grammar rule with frequency α identifies a predictor P that triggers linguistic phenomena Q in α % of cases. We denote such a rule by:

 $\mathsf{S} \implies (\mathsf{P} \stackrel{{}_{\scriptstyle \alpha} \, {}_{\scriptstyle \infty}^{\ast}}{\Longrightarrow} \; \mathsf{Q}) \, .$

Our formalization is inspired by correspondence rules in Meaning-text Theory (Mel'čuk, 1988).⁴

4. Rule Extraction Method

In order to extract rules governing grammatical structures, we first need to define what are the pieces of information that can fill the patterns in these rules. In practice, patterns S and Q are manually defined, as they define the linguistic phenomena of interest in a given scope. They can therefore rely on any information available in treebanks, without any kind of restriction.

However, potential patterns P are the ones that the machine learning (ML) model must fill. Therefore, we need to define which are the potential pieces than compose them. We then explain how we use sparse logistic regression to select and rank them.

4.1. Features

Treebanks contain attributes that are used to build features, the input of the ML model. Remember that what we aim to predict is also an attribute related to a dependency relation satisfying S (Is there a number agreement between the governor and its dependent? Is the dependent after the governor?).

The bigger the search space, the more expressive the pattern P can be. We include high-order features, that is, our search space includes all linguistic information related to the grandparent (governor of the governor), the codependents of the governor and the grandchildren (dependent's children). Figure 2 illustrates this search space.

Our features are based on all attributes appearing in the search space, with some additions:

- We add relative position attributes for the governor and the grandparent, *i.e.* one attribute indicating if the governor is before or after the dependent, and one attribute indicating if the grandparent (when existing) is before or after the governor;
- The form attribute is never used, and we use the lemma attribute only for words from a predefined set of POS tags — the main motivation is that the large number of values for these attributes introduces noise into the decision process, but that lemma information can be important for closed classes (*e.g.* AUX or ADP);
- Although the UD part of speech (upos) tagset is considered "coarse", there are distinctions that may not be meaningful for our purpose

 therefore we add additional attributes indicating if the upos of a word belongs to one of several predefined lists of tags (*e.g.* an attribute indicating that upos of the governor is either DET or NUM).

We also apply filters on these attributes, especially to prevent information leaks. For example, when we are interested in rules governing number agreement, we remove the Number attribute.

A feature is an indicator of the value of one or several specific attributes. For example, the feature DEP.upos=VERB is equal to 1 if the upos of the dependent (DEP) is VERB, and 0 otherwise. Features can focus on one attribute (*e.g.* DEP.upos) or two attributes (*e.g.* DEP.upos and the relative position of its governor). To limit noise (annotation errors or attributes that are poorly represented in the corpus), we only add features that appear at least 5 times in the scope of the rule.

4.2. Sparse Logistic Regression

We will now describe the method used to select the features that we will use to define the P patterns of the rules. We formalize the extraction process as *selecting features that allow to identify when a linguistic phenomenon Q is triggered.* For example, for number agreement between a dependent and its governor, the majority case is defined as "no agreement or agreement by chance" and we wish to select features that identify dependencies where number agreement is expected (or forced). To that end, we rely on a sparse logistic regression.⁵

We denote *F* the set of features. Given an input $x \in \{0, 1\}^F$ represented by the boolean feature in-

⁴In the notation of Mel'čuk (1988), the rule would be written "S \implies Q | P" and read "S can correspond to Q in the context P" (the subject is positioned after the verb in presence of an expletive codependent before the verb).

⁵We refer the reader to (Shalev-Shwartz and Ben-David, 2014, Sec. 9 & 25) for a more in-depth introduction to sparse logistic regression.

dicator vector, the decision of a binary linear classifier is defined as follows:

$$f_{\boldsymbol{a},b}(\boldsymbol{x}) = \begin{cases} 1 & \text{if } \boldsymbol{a}^\top \boldsymbol{x} + b > 0 \,, \\ 0 & \text{otherwise,} \end{cases}$$

where $a \in \mathbb{R}^{F}$ (feature weights) and $b \in \mathbb{R}$ (intercept term) are the parameters. However, this decision is "hard" and therefore inappropriate for our purpose. Continuing with our example, number agreement can be observed by chance: what we want to learn is the probability to observe an agreement. To this end, we can instead interpret the predictor as giving the probability that there is a number agreement between the modifier and its governor for input x:

$$\mathsf{P}(ext{``number agreement"}|m{x}) = \sigma \left(m{a}^{ op}m{x} + b
ight) \ , \ \$$
(1)

where $\sigma(w) = \frac{\exp(w)}{1+\exp(w)}$ is the sigmoid function that maps a real value to a value in [0, 1]. As such, we seek to identify features that drive this probability closer to one (forced agreement) compared to the majority case (where an agreement may happened by chance in roughly 50% of observations).

By definition of dot product $a^{\top}x = \sum_{f \in F} a_f x_f$, a feature $f \in F$ contributes to the probability in Equation 1 if and only if its associated weight a_f is not equal to 0. Intuitively, as the sigmoid function is strictly increasing, a positive a_f (resp. negative a_f) means that the presence of this feature will increase the probability of observing a number agreement (resp. a disagreement).⁶

To learn parameters a and b of the model, we rely on the following standard training objective :

$$\min_{\boldsymbol{a} \in \mathbb{R}^{F}} \quad \frac{1}{|D|} \sum_{(\boldsymbol{x}, y) \in D} \ell \left(\boldsymbol{a}^{\top} \boldsymbol{x} + b, y \right) + \lambda r(\boldsymbol{a}),$$
(2)

where *D* is the training dataset (the scope of the rule), ℓ is a loss function, *r* is a regularization term and $\lambda \ge 0$ is the regularization strength. In particular, we choose:

- \$\emp(w, y) = -yw + log(1 + exp w)\$, the negative log-likelihood, which is well-calibrated for probability estimation (Reid and Williamson, 2010);
- $r(a) = ||a||_1 = \sum_{f \in F} |a_f|$, the L1-norm which induces parameter sparsity, *i.e.* depending on the value of λ , more or less values in *a* will be exactly equal to 0 (Bach et al., 2012).

Importantly, the intercept term b is not regularized. Applying regularization to b is equivalent to assuming a Laplacean prior with mean 0 (Murphy, 2012, Sec. 13.3), which may be unsuitable.⁷ For unregularized *b*, when $\lambda \to \infty$, all feature weights will be null, *i.e.* a = 0, and the intercept *b* will capture the mean probability of phenomena Q in the scope:

$$\sigma(b) = \frac{1}{|D|} \sum_{(\boldsymbol{x}, y) \in D} y.$$

Practical consideration. The optimization problem in Equation 2 is challenging as there is a large number of parameters and the regularization term is non-smooth. The SCIKITLEARN library relies on LIBLINEAR for this problem, but does not support unreguralized intercept term. In this work, we therefore use SKGLM (Bertrand et al., 2022).

4.3. Ranking via Regularization Path

The term regularization path refers to the function $\lambda \mapsto \hat{a}(\lambda)$ that maps the regularization parameter λ to the model parameters a minimizing the training problem (Eq. 2), denoted $\hat{a}(\lambda)$. We say that feature $f \in F$ is in (resp. out) the regularization path at λ if $\hat{a}(\lambda)_f \neq 0$ (resp. = 0). Importantly, when $\lambda \rightarrow \infty$, $\hat{a}(\lambda)_f = 0$, *i.e.* no feature is selected by the classifier.

To select and rank features that form pattern P, we proceed as follows:

- 1. We fix $\lambda^{(0)}$ to a reasonable value where $\widehat{a}(\lambda^{(0)}) = \mathbf{0};$
- As we cannot explore all possible λ values, we fix a number k of additional values we will inspect;
- 3. We fix $\lambda^{(k)}$ to a reasonable minimum value for the regularization strength.

Then we train the model for k + 1 regularization strengths $\lambda^{(i)}$, $0 \le i \le k$, evenly distributed between $\lambda^{(0)}$ and $\lambda^{(k)}$. In practice, we fix k = 100, $\lambda^{(0)} = 0.1$ and $\lambda^{(k)} = 0.001$.

We say the feature $f \in F$ enters the regularization path at step i if and only if $\forall j \in \{0, ..., i-1\}$: $\hat{a}(\lambda^{(j)}) = 0$ and $\hat{a}(\lambda^{(i)}) \neq 0$. In other words, $f \in F$ enters the regularization path at step i if it is the first step where its associated weight in the parameter vector a is non-null. The features are ranked by entering order in the regularization path: if features f and f' enter the path at steps i and j, respectively, such that i < j, we consider that feature f is more salient than feature f' as f is more important to predict the frequency of pattern Q in the scope defined by S.

⁶Note that there are limitations to this interpretation that will be discussed in Section 4.4.

⁷For example, in the case of number agreement, regularizing *b* is equivalent to postulating that there is a 50% chance of observing an agreement. Without regularization, we leave full flexibility to the model to estimate this prior to agreement.

On hyperparameters. Importantly, note that the set of considered regularization strengths is the only hyperparameter of our model.⁸ This is a very important feature as it is difficult to tune hyperparameters. Indeed, the standard ML methodology consists of tuning them on a development set. However, predictions of the classifier are frequencies: they should be interpreted, given a set of observed features, as the frequency of the targeted phenomenon. If the model predicts 45% of agreements, this mostly means agreements by chance, and it does not make sense to consider that the model makes an error if for a given data point x, the output label is 1. As product of features can be rare, especially for not-so-frequent linguistic phenomena, the accuracy variance will be high on a small development set, making the hyperparameter tuning procedure meaningless. Our approach does not suffer from such limitation.

4.4. Rule Analysis

It is important to note that, contrary to popular belief, the parameters of a linear model are not directly interpretable as explanatory factors. As explained in previous section, the regularization path only gives us ranking of features (by their saliency), but the particular non-null values in *a* must be ignored due to potential non-linear correlations (Achen, 2005). Therefore, for each considered feature for pattern P, we must: (1) check if it is a trigger of Q or \neg Q and (2) compute other statistical measures to provide more insight.

In the following, we denote #(P) the number of dependencies matching a given pattern *P*.

Statistical test. We denote μ the Bernoulli parameter of the expected distribution, that is, the ratio of pattern Q in the scope:

$$\mu = \frac{\#(\mathsf{S} \land \mathsf{Q})}{\#(\mathsf{S})}$$

Consider a potential rule $S \Rightarrow (P \stackrel{\alpha \%}{\Rightarrow} Q)$ where:

$$\alpha = \frac{\#(\mathsf{S} \land \mathsf{P} \land \mathsf{Q})}{\#(\mathsf{S} \land \mathsf{P})}$$

This grammatical rule is of statistical interest if α is "different enough" from μ . However, given that we approximate α using a limited number of samples n, we need to test if this divergence is statistically significant. The G-test statistic is defined as follows:

$$G = 2 \times n \times \left(\alpha \cdot \ln \frac{\alpha}{\mu} + (1 - \alpha) \cdot \ln \frac{1 - \alpha}{1 - \mu} \right) .$$
 (3)

Note that this is simply a KL-divergence between Bernoullis parameterized by μ and α but reweighted by the number of samples used to estimate α . The p-value of a G-test is calculated by looking at the tail probability of the χ^2 distribution with the right degrees of freedom (*i.e.* 1 in our case). We consider a grammar rule to be statistically significant if p < 0.01, *e.g.* agreement is not by chance if it is statistically significant (see Chaudhary et al. (2020)). Note that the G-test statistic can be adapted for rules triggering \neg Q. Moreover, it can also be used to rank rules: as it is (mostly) a KL-divergence, rules that deviate more from the base distribution μ will have a higher value.

Association measures. We compute the *Cramer's phi* ϕ_c effect size using the G-test statistic. Unlike Chaudhary et al. (2020), we do not use it to filter rules because the threshold considered for doing so vary depending on the data and domains. However, we use effect size as another interpretable measure to rank rules.

Coverage and precision. We hypothesize that reliable grammar rules are those that covers largely the linguistic phenomenon of interest and those that are precise. Coverage/recall and precision of the pattern P are defined as follows:

Rule	Coverage	Precision		
$\overline{S \implies (P \stackrel{\alpha \text{\ensuremath{\sim}}}{\Longrightarrow} Q)}$	$\frac{\#(S \land P \land Q)}{\#(S \land Q)}$	$\frac{\#(S \land P \land Q)}{\#(S \land P)}$		
$S \implies (P \stackrel{\overset{\alpha\%}{\Longrightarrow}}{\Longrightarrow} \neg Q)$	$\frac{\#(S \land P \land \neg Q)}{\#(S \land \neg Q)}$	$\frac{\#(S \land P \land \neg Q)}{\#(S \land P)}$		

Note that these measures must be used with cautions. For example, a number agreement between a dependent and its governor can happen by chance. As such, a rule that trigger forced number agreement will never achieve a 100 % cover (and will probably be way lower).

5. Experimental Results

The extracted rules for word order are based on the relative position of a governor and its dependent. We look at particular cases, such as the SUBJ relation and the relative position of the subject regarding its verb/auxiliary governor. For example, we obtained patterns that favor placing the subject before or after the verb, compared to the general positioning of the subject. We partly follow Chaudhary et al. (2022), which in turn based their work on WALS word order rules (Dryer and Haspelmath, 2013). We extract word order rules for the nondominant position for pairs subject-verb/auxiliary, object-verb and adjective-noun. As for the agreement rules, we explore agreement in number, person and gender between a governor and a dependent having these features, as well as agreement for particular relations, such as SUBJ or COMP:OBJ.

⁸More precisely, the model also have hyperparameters related to the optimization algorithm. However, they can be safely be ignored and set to default values.

	P PATTERN	position	λ	coverage	precision
1	dep.NumType=Ord,dep.rel_synt=mod	before	0.042	21.3	98.2
2	dep.rel_synt=mod,siblings.upos=ADP	before	0.033	40.9	39.9
3	grandchildren.upos=ADP	after	0.022	12.8	98.9
4	gov.upos=NOUN PROPN PRON,grandchildren.upos=ADP	after	0.021	12.8	98.9
5	dep.rel_synt=mod,siblings.lemmas=de	before	0.012	31.8	41.6
6	grandparent.position=before_gov,grandparent.upos=ADP	after	0.011	49.9	75.1
7	dep.Degree=Cmp,dep.rel_synt=mod	before	0.011	5.9	90
8	grandparent.lemma=de	after	0.01	23.8	79.7
9	gov.upos=NOUN PROPN PRON,grandparent.lemma=de	after	0.01	23.8	79.7
10	gov.rel_synt=comp:obj,siblings.Definite=Def	after	0.01	35.1	74

Table 2: The ten most salient noun-adjective word order rules, extracted from the SUD_Spanish-AnCora treebank and ranked based on the order given by the linear classifier. P is the pattern that favors the adjective position indicated in the second column. See Section 4.3 for the interpretation of λ . Coverage and precision are expressed as percentages.

The code to reproduce the experiments is freely available,⁹ together with a tool accessible online to visualize results on different treebanks.¹⁰

In the following, we analyze in detail, as an exploratory example, the ten most salient word-order rules between a noun and its dependent adjective in Spanish (Table 2). We will then provide a study on word order involving the object in Wolof (Table 3). Finally, we comment on the extracted agreement rules. As far as we know, the rules analysed have not been found in comparable previous works, unless otherwise stated. The following qualitative analysis covers only a small portion of the rules that were extracted, but it is intended to be representative of the rules extracted by our method. No filters or pruning techniques were applied. Interesting and precise grammar rules and patterns are found for all the languages studied.

5.1. Noun-adjective order in Spanish

In Spanish, the dominant order is for the adjective to follow its governing noun. In the SUD_Spanish-AnCora treebank, about 28% of dependent adjectives precede their noun. We search for rules that, for all noun-adjective in a dependency relation with a governing noun (S GOV.UPOS=NOUN,DEP.UPOS=ADJ), trigger the adjective anteposition (GOV.POSITION=AFTER_DEP) and more exactly a shift in distribution.

Adjective position is primarily determined lexically and our most salient rule (rule 1 in Table 2) indicates that ordinal adjectives (1) in a modifier dependency relation precede almost always (precison 98%) their noun, while adjectives in general precede their noun only 28% of cases. It's worth noting that rule 1 covers about 21% of the

¹⁰https://autogramm.github.io/

adjectives that precede the noun in the treebank. This is not the only rule that includes a lexical feature: rule 7 shows that comparative adjectives (DEP.DEGREE=CMP) are also much more likely to be before their noun.

(1) Desde un **sexto** piso from a **sixth.ord** floor.noun 'From a sixth floor'

However, contrary to previous works (notably Chaudhary et al. (2022)), we also capture pure syntactic rules (*i.e.* not involving lexical attributes). In the ten most salient rules, the majority of them involve the presence of prepositional phrases (PP) in some way (cf. ADP for adposition). Rule 2, for example, indicate that nouns having a PP tend to have more preceding adjectives (2). Rule 5 is a finer-grained version of this rule indicating that nouns governing a preposition de ('of', 'from') have also more preceding adjectives than the expected. The syntactic interpretation of rules such as these must be done with caution. The anteposition may be explained by the fact that the postnominal position is already occupied by the PP, but also by the fact that the noun with a PP prefers adjectives that are usually placed before it, like ordinals. These two rules can be considered as syntactic tendencies due to their low/medium precision. They appear in the top rules due to their significant coverage.

(2) Rara y triste evidencia de madurez Rare and sad evidence of.ADP maturity 'Rare and sad evidence of maturity.'

Rule 6 (and its finer-grained versions, rules 8 and 9) shows that adjectives follow their noun in almost 75% of cases when the latter is governed by a preposition (GRANDPARENT.UPOS=ADP) that is, *i.e.* when they are in a PP. Again, nouns may prefer adjectives in a post-nominal position than preced-

⁹https://github.com/FilippoC/

grex-lrec-coling-2024

grex-lrec-coling-2024

ing them.

Finally, we capture also really precise rules for the post-position of the adjective. Dependent adjectives governing a preposition (GRANDCHILDREN.UPOS=ADP), that is, which have a PP, are more likely to be postponed (rules 3, 4). We find the same rule in English, where the dominant order is adjective-noun. In these cases, the postposition of an adjective is explained by the weight or heaviness of the syntactic group they govern, see Thuilier (2012) for French.

Our method can capture these rules due to the a large search space and sensitivity to distribution shifts. In addition, even though we proposed a more expressive system than previous works, our rules remain easy to interpret. However, the interpretation of the rules must be lexico-syntactic and must take into account the probabilistic and overlapping nature of the extracted rules.

5.2. Object order in Wolof

Chaudhary et al. (2022) extract three word order rules for the object in Wolof, which are our rules 5, 12, and 54 of Table 3: the object tends to be placed before the verb when it is an accusative pronoun governed by a verb depending on an auxiliary (rule 12), or a pronoun governed by a verb being in a relative clause (rule 5), or bearing the perfective aspect marker (rule 54).¹¹ They do not capture more general tendencies, neither more specific rules.

Our set of rules can be interpreted as a hierarchy of rules, where the three previous rules are particular nodes. Rule 2,3 indicate that pronouns tends to be placed before the verb (contrary to nouns). Rules 1 and 4 shows that this tendency increases if the governor is a finite verb or if it is a verb in the indicative mood. Rule 6 indicates that the tendency also increases if the verb depends on an auxiliary (precision 72%). It increases even more if the auxiliary is di (rule 11, precision 85%) or if the pronoun is marked by the accusative case (rule 12, precision 78%). Even if rule 12 is descriptive by itself, we find more interesting to have it in relation with rules 6 and 11, which are more covering for the first and more precise for the second. Similarly, rule 54 for verbs with a perfective aspect is interesting, but rule 27 adding the constraint for the pronoun to be accusative is more precise (92% vs 85%). It is not clear why Chaudhary et al.'s tool add this additional condition for object pronouns depending on a verb and on an auxiliary (and choose rule 12 rather than rule 6), and not for pronouns depending on a verb with a perfective aspect (rule 54 rather than rule 27).

Having this rich hierarchy of rules (and tendencies) is invaluable, but it also raises the question of the selection of particular rules. We are not sure that a unique and simple solution can be provided and the answer probably depends on the objective. We can be more interested by results with a high precision (they are the cases where we can really talk about "rules") or prefer a compromise between precision and coverage, which is what our method seems to do (see Section 5.4). If we do not keep the whole hierarchy of results, it is also important to select a set of rules/tendencies that covers most of the data and are not too redundant, which is what Chaudhary et al.'s method based on decision trees does well.

5.3. Agreement

Agreement is a phenomenon that is more restrictive than word order, really stable in Indo-European language such as English, French and Spanish. This is the reason why the rules extracted are less diverse and more expected. For number agreement, we capture usual rules like agreement between a noun and a determiner (e.g. rule 17 in Spanish) or the assessments that dependent adjectives agree with their governor and that verbs agree with their subjects (e.g. rule 4 and 41 in French). As regards negative rules, we exclude, for example, agreement between objects and their verbs in French, a crucial test for our method, which represent 3% of mismatch between heads and dependents having the number feature.

The extraction of gender agreement rules suffers from different annotation strategies. Being an inherent property of the noun or the referred entity, it may not always be expressed by morphosyntactic flexion. In any case, principal agreement rules were found in French and Spanish, as the agreement between the noun and its adjectives and determiners.

The most informative way to extract agreement patterns is by checking for agreement between any head and dependent, as its done in Chaudhary et al. (2020). This provides a general overview of the agreement system for a given language. However, due to its under specification, some rules are more general. For example, we extract rules expressing the fact that number and gender agreement co-occur both in French and Spanish (see French rules 28 and 29 for number agreement). Our findings also indicate that in French and Spanish, number agreement generally occurs when the governor follows the dependent. This is particularly evident in the agreement between determiners and nouns, which is the largest set of agreed pairs in both languages. On the contrary, about 85% of dependents located after their head, if they

¹¹Our rule 12 does not indicate that the dependent is a pronoun, but only pronouns have a Case feature in the Wolof-WTB treebank.

	P PATTERN	position	λ	coverage	precision
1	dep.upos=PRON,gov.VerbForm=Fin	before	0.1	92.6	63.2
2	dep.upos=NOUN PROPN PRON,dep.upos=PRON	before	0.077	94.7	58.9
3	dep.upos=PRON	before	0.067	94.7	58.9
4	dep.upos=PRON,gov.Mood=Ind	before	0.064	91.7	63.4
5	dep.upos=PRON,gov.rel_deep=mod@relcl	before	0.041	36.7	96.7
6	dep.upos=PRON,gov.rel_synt=comp:aux	before	0.033	47	72.1
11	dep.upos=PRON,grandparent.lemma=di	before	0.025	30.3	85
12	dep.Case=Acc,gov.rel_synt=comp:aux	before	0.023	35.1	78.8
27	dep.Case=Acc,gov.Aspect=Perf	before	0.009	8.2	92.5
54	dep.upos=PRON,gov.Aspect=Perf	before	0.003	9.1	85.7

Table 3: Word order rules for the pre-verbal position of the object mentioned in 5.2 and extracted from the SUD_Wolof-WTB treebank. The order is given by the linear classifier, the last rule of the table coming at position 54. See Section 4.3 for the interpretation of λ . P is the pattern that favors the object anteposition. Coverage and precision are expressed as percentages.

have a number feature, do not agree in French (see rule 19). Global analysis like the precedent one are useful in quantitative typology (Gerdes et al., 2021).

5.4. Ranks

By expanding the search space and due to the nature of our rules, we extract a greater number of rules compared to previous works. It is why order and ranking the rules is crucial to achieving better interpretability. Beyond the order given by the linear classifier (section 4.3), it is possible to rank the rules using the other statistical measures described (section 4.4), as is done in Corpus Linguistics (e.g. Pecina and Schlesinger, 2006), prioritizing differences of aspects of the same rule, like the precision or coverage of the rule. In addition, these measures could be used to prune subsets in different ways and reducing the final number of rules.

We compared the orders of the rules given by the model and by using the G-test statistic. We computed a Spearman's rank order correlation for word order experiments, and we obtained significant results, going from weak-moderate to high positive correlation (r > 0.40, p value < 0.01) with some exceptions (*e.g* adjective-noun word order). This could mean that while they are in most cases order is correlated, they do not highlight exactly the same type of information. It is worth noting that the order given by the classifier is partial and less granular, with several rules occupying the same place.

6. Conclusion

We proposed a novel method for grammatical rule extraction from treebanks. Our approach is based on (1) a formal definition of what is a syntactic grammatical rule and (2) the regularization path of a sparse logistic classification model. Moreover, we showed in our result analysis that it is important to capture less pronounced shift distributions to include a greater variety of rule patterns. This allows to increase the expressiveness of rules and to capture linguistic phenomena that are out of the scope of previous work (Chaudhary et al., 2020, 2022; Blache et al., 2016).

Alongside our main contribution, we hope that our clear formal description of each component will help to develop future work on the topic.

7. Limitations

The extracted grammatical rules rely on the annotation scheme. Therefore, some patterns extracted may reflect corpus properties and theoretical decisions more than actual grammar rules. Our search for agreement rules, which is consistent with previous researches, is limited to syntactic agreement. As a result, one of the most salient (negative) rule for the subject agreement is to exclude subjects having coordination, which is due to the fact that the verb is linked to the first conjunct, not reflecting the features of the coordinated phrase. While our extraction system does not capture this kind of semantic agreement, its behavior remains consistent and contributes to understandability.

Grammar rules also reflect annotation errors or lack/excess of annotation. The most precise negative rule for number agreement in French captures full dates (*6 janvier 1902* 'January 6, 1902') because the year is labeled as plural while its head, the month, as a singular noun.¹²

Extracted rules for marginally existing linguistic phenomena are much more loose and less significant. This is the case of agreement between ob-

¹²Note that *1902* when used as a proper name as here should be singular, since it triggers a singular agreement: *1902 is ...*

jects and verbs in the languages studied. The observed results show preferences for language use, not grammar rules, such as the fact that singular objects are significantly paired with singular subjects (likewise for plurals), resulting in a tendency for the verb to agree with its object. Conversely, the agreement rule of a past participle with its object when the object precedes it existing in French was not found, probably because the relative pronoun *que* ('that') does not bear any agreement feature.

Acknowledgement

We thank Mathurin Massias and Badr Moufad for the help with the SKGLM library. We also thank Bruno Guillaume, Guillaume Bonfante, and Guy Perrier for their insightful conversations and feedback. We thank the anonymous reviewers for their comments and suggestions. All authors are supported by a public grant overseen by the French ANR (Autogramm, ANR-21-CE38-0017).

8. Bibliographical References

- Christopher H Achen. 2005. Let's put garbage-can regressions and garbage-can probits where they belong. *Conflict Management and Peace Science*, 22(4):327–339.
- Diego Alves, Božo Bekavac, Daniel Zeman, and Marko Tadić. 2023. Analysis of corpus-based word-order typological methods. In *Proceedings* of the Sixth Workshop on Universal Dependencies (UDW, GURT/SyntaxFest 2023), pages 36– 46, Washington, D.C. Association for Computational Linguistics.
- Francis Bach, Rodolphe Jenatton, Julien Mairal, Guillaume Obozinski, et al. 2012. Optimization with sparsity-inducing penalties. *Foundations and Trends*[®] *in Machine Learning*, 4(1):1–106.
- Eric Baroni. 2021. On the gap between computational and theoretical linguistics.
- Emi Baylor, Esther Ploeger, and Johannes Bjerva. 2024. Multilingual gradient word-order typology from Universal Dependencies. In *Proceedings* of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers), pages 42–49, St. Julian's, Malta. Association for Computational Linguistics.
- Emily M. Bender, Joshua Crowgey, Michael Wayne Goodman, and Fei Xia. 2014. Learning grammar specifications from IGT: A

case study of chintang. In *Proceedings of the 2014 Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 43–53, Baltimore, Maryland, USA. Association for Computational Linguistics.

- Q. Bertrand, Q. Klopfenstein, P.-A. Bannier, G. Gidel, and M. Massias. 2022. Beyond I1: Faster and better sparse models with skglm. In *NeurIPS*.
- Rajesh Bhatt, Owen Rambow, and Fei Xia. 2012. Creating a Tree Adjoining Grammar from a multilayer treebank. In *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+11)*, pages 162– 170, Paris, France.
- Philippe Blache, Stéphane Rauzy, and Grégoire Montcheuil. 2016. MarsaGram: an excursion in the forests of parsing trees. In *Proceedings* of the Tenth International Conference on Language Resources and Evaluation (LREC'16), pages 2336–2342, Portorož, Slovenia. European Language Resources Association (ELRA).
- Michael Burke, Olivia Lam, Aoife Cahill, Rowena Chan, Ruth O'Donovan, Adams Bodomo, Josef van Genabith, and Andy Way. 2004. Treebankbased acquisition of a chinese lexical-functional grammar. In *Pacific Asia Conference on Language, Information and Computation*.
- Aoife Cahill, Michael Burke, Martin Forst, Ruth O'donovan, Christian Rohrer, Josef van Genabith, and Andy Way. 2005. Treebank-based acquisition of multilingual unification grammar resources. *Research on Language and Computation*, 3:247–279.
- Eugene Charniak. 1996. Tree-bank grammars. In Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2, AAAI'96, page 1031–1036. AAAI Press.
- Aditi Chaudhary, Antonios Anastasopoulos, Adithya Pratapa, David R. Mortensen, Zaid Sheikh, Yulia Tsvetkov, and Graham Neubig. 2020. Automatic extraction of rules governing morphological agreement. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 5212–5236, Online. Association for Computational Linguistics.
- Aditi Chaudhary, Zaid Sheikh, David R Mortensen, Antonios Anastasopoulos, and Graham Neubig. 2022. Autolex: An automatic framework for linguistic exploration.

- John Chen, Srinivas Bangalore, and K. Vijay-Shanker. 2006. Automated extraction of treeadjoining grammars from treebanks. *Natural Language Engineering*, 12(3):251–299.
- Hee-Soo Choi, Bruno Guillaume, and Karën Fort. 2021a. Corpus-based language universals analysis using Universal Dependencies. In Proceedings of the Second Workshop on Quantitative Syntax (Quasy, SyntaxFest 2021), pages 33– 44, Sofia, Bulgaria. Association for Computational Linguistics.
- Hee-Soo Choi, Bruno Guillaume, Karën Fort, and Guy Perrier. 2021b. Investigating dominant word order on Universal Dependencies with graph rewriting. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 281–290, Held Online. INCOMA Ltd.
- Marie-Catherine De Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal dependencies. *Computational linguistics*, 47(2):255–308.
- Cheikh Bamba Dione. 2019. Developing Universal Dependencies for Wolof. In *Proceedings of the Third Workshop on Universal Dependencies (UDW, SyntaxFest 2019)*, pages 12–23, Paris, France. Association for Computational Linguistics.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online (v2020.3)*. Zenodo.
- Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. 2004. Least angle regression. *Annals of Statistics*.
- Anette Frank, Louisa Sadler, Josef van Genabith, and Andy Way. 2003. From treebank resources to lfg f-structures: Automatic f-structure annotation of treebank trees and cfgs extracted from treebanks. *Treebanks: Building and Using Parsed Corpora*, pages 367–389.
- Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2018. Sud or surface-syntactic universal dependencies: An annotation scheme near-isomorphic to ud. In *Universal Dependencies Workshop (UDW)*.
- Kim Gerdes, Bruno Guillaume, Sylvain Kahane, and Guy Perrier. 2022. Starting a new treebank? Go SUD! Theoretical and practical benefits of the Surface-Syntactic distributional approach. In SyntaxFest Depling 2021 - 6th International Conference on Dependency Linguistics, pages 35–46, Sofia, Bulgaria. Association for Computational Linguistics.

- Kim Gerdes, Sylvain Kahane, and Xinying Chen. 2021. Typometrics: From implicational to quantitative universals in word order typology. *Glossa: a journal of general linguistics*, 6(1).
- Bruno Guillaume, Marie-Catherine de Marneffe, and Guy Perrier. 2019. Conversion et améliorations de corpus du français annotés en Universal Dependencies. *Revue TAL : traitement automatique des langues*, 60(2):71–95.
- Kristen Howell and Emily M. Bender. 2022. Building analyses from syntactic inference in local languages: An HPSG grammar inference system. In Northern European Journal of Language Technology, Volume 8, Copenhagen, Denmark. Northern European Association of Language Technology.
- Natalia Levshina. 2022. Corpus-based typology: applications, challenges and some solutions. *Linguistic Typology*, 26(1):129–160.
- HM Markowitz. 1952. Portfolio selection. *Journal* of *Finance*.
- Igor Mel'čuk. 1988. *Dependency Syntax: Theory and Practice*. State University of New York Press.
- Kevin P Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- Michael R Osborne, Brett Presnell, and Berwin A Turlach. 2000. On the lasso and its dual. *Journal of Computational and Graphical statistics*, 9(2):319–337.
- Pavel Pecina and Pavel Schlesinger. 2006. Combining association measures for collocation extraction. In *Proceedings of the COLING/ACL* 2006 Main Conference Poster Sessions, pages 651–658, Sydney, Australia. Association for Computational Linguistics.
- Edoardo Maria Ponti, Helen O'Horan, Yevgeni Berzak, Ivan Vulić, Roi Reichart, Thierry Poibeau, Ekaterina Shutova, and Anna Korhonen. 2019. Modeling language variation and universals: A survey on typological linguistics for natural language processing. *Computational Linguistics*, 45(3):559–601.
- J. Ross Quinlan. 1986. Induction of decision trees. *Machine learning*, 1:81–106.
- Ines Rehbein. 2009. *Treebank-based grammar acquisition for German*. Ph.D. thesis, Dublin City University.
- Mark D. Reid and Robert C. Williamson. 2010. Composite binary losses. *Journal of Machine Learning Research*, 11(83):2387–2422.

- Shai Shalev-Shwartz and Shai Ben-David. 2014. Understanding machine learning: From theory to algorithms. Cambridge university press.
- Hedvig Skirgård, Hannah J. Haynie, Damián E. Blasi, Harald Hammarström, Jeremy Collins, Jay J. Latarche, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Sam Passmore, Angela Chira, Luke Maurits, Russell Dinnage, Michael Dunn, Ger Reesink, Ruth Singer, Claire Bowern, Patience Epps, Jane Hill, Outi Vesakoski, Martine Robbeets, Noor Karolin Abbas, Daniel Auer, Nancy A. Bakker, Giulia Barbos, Robert D. Borges, Swintha Danielsen, Luise Dorenbusch, Ella Dorn, John Elliott, Giada Falcone, Jana Fischer, Yustinus Ghanggo Ate, Hannah Gibson, Hans-Philipp Göbel, Jemima A. Goodall, Victoria Gruner, Andrew Harvey, Rebekah Hayes, Leonard Heer, Roberto E. Herrera Miranda, Nataliia Hübler, Biu Huntington-Rainey, Jessica K. Ivani, Marilen Johns, Erika Just, Eri Kashima, Carolina Kipf, Janina V. Klingenberg, Nikita König, Aikaterina Koti, Richard G. A. Kowalik, Olga Krasnoukhova, Nora L.M. Lindvall, Mandy Lorenzen, Hannah Lutzenberger, Tônia R.A. Martins, Celia Mata German, Suzanne van der Meer, Jaime Montoya Samamé, Michael Müller, Saliha Muradoglu, Kelsey Neely, Johanna Nickel, Miina Norvik, Cheryl Akinyi Oluoch, Jesse Peacock, India O.C. Pearey, Naomi Peck, Stephanie Petit, Sören Pieper, Mariana Poblete, Daniel Prestipino, Linda Raabe, Amna Raja, Janis Reimringer, Sydney C. Rey, Julia Rizaew, Eloisa Ruppert, Kim K. Salmon, Jill Sammet, Rhiannon Schembri, Lars Schlabbach, Frederick W.P. Schmidt, Amalia Skilton, Wikaliler Daniel Smith, Hilário de Sousa, Kristin Sverredal, Daniel Valle, Javier Vera, Judith Voß, Tim Witte, Henry Wu, Stephanie Yam, Jingting Ye , Maisie Yong, Tessa Yuditha, Roberto Zariquiey, Robert Forkel, Nicholas Evans, Stephen C. Levinson, Martin Haspelmath, Simon J. Greenhill, Quentin D. Atkinson, and Russell D. Gray. 2023. Grambank reveals global patterns in the structural diversity of the world's languages. Science Advances, 9.
- Mariona Taulé, M. Antònia Martí, and Marta Recasens. 2008. AnCora: Multilevel annotated corpora for Catalan and Spanish. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Juliette Thuilier. 2012. *Contraintes préférentielles et ordre des mots en français*. Ph.D. thesis, Université Paris-Diderot-Paris 7.

- Veniamin Veselovsky, Manoel Horta Ribeiro, and Robert West. 2023. Artificial artificial intelligence: Crowd workers widely use large language models for text production tasks.
- Dietmar Zaefferer, Bernard Comrie, William Croft, and Christian Lehmann. 1992. "a framework for descriptive grammars". In *Proceedings of the XVth International Congress of Linguists*, pages 159–170, Sainte-Foy :. Presses de l'Université Laval.
- Amir Zeldes. 2017. The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581– 612.
- Yi Zhang, Rui Wang, and Yu Chen. 2012. Joint grammar and treebank development for Mandarin Chinese with HPSG. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 1868–1873, Istanbul, Turkey. European Language Resources Association (ELRA).