

Mapping ‘when’-clauses in Latin American and Caribbean languages: an experiment in subtoken-based typology

Nilo Pedrazzini

The Alan Turing Institute (London, United Kingdom)
npedrazzini@turing.ac.uk

Abstract

Languages can encode temporal subordination lexically, via subordinating conjunctions, and morphologically, by marking the relation on the predicate. Systematic cross-linguistic variation among the former can be studied using well-established token-based typological approaches to token-aligned parallel corpora. Variation among different morphological means is instead much harder to tackle and therefore more poorly understood, despite being predominant in several language groups. This paper explores variation in the expression of generic temporal subordination (‘when’-clauses) among the languages of Latin America and the Caribbean, where morphological marking is particularly common. It presents probabilistic semantic maps computed on the basis of the languages of the region, thus avoiding bias towards the many world’s languages that exclusively use lexified connectors, incorporating associations between character n -grams and English *when*. The approach allows capturing morphological clause-linkage devices in addition to lexified connectors, paving the way for larger-scale, strategy-agnostic analyses of typological variation in temporal subordination.

1 Introduction

Across the 7000+ world’s languages recorded by the Glottolog database (Nordhoff and Hammarström 2011, Hammarström et al. 2023)¹ there is great variation in how temporal relations between different eventualities can be encoded in a sentence or discourse unit. English has one main generic temporal subordinator, *when*, which is relatively underspecified with respect to the temporal semantic relation between the clause it introduces and its matrix clause, compared to semantically more precise connectors (e.g. *after*, *before*, or *while*). The number and scope of generic temporal subordinators can vary cross-linguistically from one (e.g.

Italian *quando*), to two (e.g. German *wenn/als*) or several more (e.g. Pular *ndel/sil/bay/fewndoltuma*; Evans 2017; Pedrazzini 2023). Crucially, languages can additionally or exclusively encode WHEN-clauses² morphologically on the predicate, rather than using a lexified subordinator (cf. Spanish *viendo* ‘see.GER’³ as opposed to *cuando vio* ‘when saw.3.SG’; Ukrainian *pobačyvšy* ‘see.GER’ as opposed to *koly vyn pobačyv* ‘when he saw’). Because of the very nature of competition, overarching semantic differences between subordination strategies within individual languages cannot be fully captured in terms of discrete, categorical variables, but they should be modeled as a continuum allowing for a degree of overlap, aiming to reveal broader patterns in a probabilistic, rather than a fully deterministic way. Previous studies (Haug and Pedrazzini 2023) have employed a ‘token-based approach’ (Levshina 2019, 2022) to explore the semantic ground covered by English *when* and *induce* cross-linguistically common semantic dimensions from parallel corpora. In Haug and Pedrazzini (2023), probabilistic semantic maps (Croft and Poole 2008; Wälchli and Cysouw 2012) of WHEN were generated from a massively parallel corpus of 1400+ linguistic varieties (Mayer and Cysouw 2014), to capture systematic variation in the ways languages tend to divide the semantic space of English WHEN by using different lexical items for its different meanings. One of the greatest limitations of a purely token-based typological approach to the study of temporal subordination in the world’s languages is that it does not

²Small caps WHEN is used to refer to the semantic concept of ‘generic temporal subordination’, rather than the English lemma *when* (written in italics).

³The following abbreviations are used in glosses throughout this paper: GER = gerund, 3 = third person, SG = singular, PL = plural, SBJ = subject, VIS = visible (speaker’s area), SS = same subject, DS = different subject, DISTR = distributive, NARR = narrative, NSBJ = non-subject, LOC = locative, AS2 = secondary assertion, *pro* = prominent, PFV = perfective.

¹<https://glottolog.org>

allow to account for variation *within* the semantic space covered by non-lexified WHEN-clauses cross-linguistically. That is, it will merely allow us to observe that particular *subsets* of *when*-occurrences are more likely to lack a parallel token in the target languages, without further identifying typologically widespread constructions (or *gram types*; Dahl and Wälchli 2016) within the semantic sub-space of non-lexified WHEN-clauses.

While languages using predominantly or exclusively morphological means to express generic temporal subordination are relatively uncommon among European languages, non-lexified WHEN-clauses are instead particularly frequent among Latin American languages, as evidenced by the plethora of areal studies on converbal, clause-bridging, and, especially, switch-reference morphology in the region (among others, van Gijn et al. 2011; van Gijn 2012, 2016; Overall 2014, 2016).

This paper zooms in on the languages of Latin American and the Caribbean, given the particular computational challenges posed by their common, extensive use of non-lexified WHEN-clauses (exclusively so or in addition to lexified means). As in previous experiments, Mayer and Cysouw’s (2014) massively parallel corpus of New Testament translations is used, and probabilistic semantic maps are adopted as a base method to induce typologically relevant dimensions within the semantic space of WHEN, since they allow capturing the gradience and overlap between different means in any given language, as well as the language-internal variation which is inherent to the very concept of competition. The goal of this paper is twofold:

- a. incorporate associations between character *n*-grams and English *when* for capturing differences among WHEN-clauses that are expressed morphologically *as well as* lexically, and generate probabilistic semantic maps based on the parallel dataset thus refined. As detailed in Section 2, this method builds on Asgari and Schütze’s (2017) ‘SuperPivot’ approach, but with substantial changes to their pipeline. Crucially, it gets rid of the assumption that there should be at most one ‘pivot’ (i.e. a marker in a parallel language) per linguistic feature (e.g. ‘past’ in Asgari and Schütze’s 2017 example), reflecting instead the existing typological knowledge about the nature of generic temporal subordination as a phenomenon with great language-internal

variation. The code to achieve this is released alongside this paper as a generalized tool, which starts from one or several lexical items in a source language and can be used to look for systematic cross-linguistic variation in a parallel dataset, both at the lexical and morphological level;

- b. generate probabilistic semantic maps that are built exclusively on the basis of the languages of the region, thus avoiding bias towards the many world’s languages that exclusively or predominantly use lexified connectors. The resulting maps and parallel data enriched with *n*-gram annotation are also released to facilitate future computational experiments.⁴

2 Methods

Dataset creation The Latin American and Caribbean parallel language data used in this experiment is a subset of Mayer and Cysouw’s (2014) massively parallel corpus. To identify Latin American and Caribbean varieties in the massively parallel corpus, a GeoJSON dataset was manually created using <https://geojson.io/> to define the geographical region of interest. The approximate coordinates for each language variety in the dataset were taken from Glottolog and assigned to each New Testament translation based on its associated ISO 639-3 code. All varieties whose approximate coordinates were outside of the polygon defined by the GeoJSON dataset were filtered out from the corpus. The resulting data consisted of 335 varieties, representing approximately one-third of all the languages (1,005) recorded for Latin America and the Caribbean by Glottolog.⁵ Figure 1 shows the areal distribution of the languages in our dataset among all the languages with an ISO 639-3 code from the region.

Word alignment & semantic mapping SyMGIZA++ (Junczys-Dowmunt and Szał 2012) was used to align the English version of the New Testament to each of the translations in our dataset at the token level, achieving a one-to-one token alignment for each language (i.e. each English token corresponds to at most one token in the target language, in contrast to possible one-to-many or many-to-one

⁴The code, datasets and all the maps, only a very small portion of which is presented in this paper, can be found in the [associated repository](#).

⁵This number excludes sign languages, as we focus on textual data.




Figure 1: Approximate areal distribution of the languages in the dataset (orange) among the languages listed by Glottolog for the region (blue).

alignments). The occurrences of English *when* and its parallels in all Latin American and Caribbean languages in the dataset were then extracted. The quality of the automatic alignment was evaluated based on a sample of 300 *when*-clauses manually aligned to the Huichol translation, against which automatic alignment achieved a precision of 0.66, recall of 1, and F1-score of 0.79.⁶

Each instance of *when* and its parallel in every target language was treated as one usage point for WHEN. Hamming distance was applied as a measure of dissimilarity between pairs of usage points, by counting the number of languages using two different words, as opposed to the same word, for

⁶To calculate precision and recall, the presence of an aligned word in the target language was considered a ‘positive’, whereas the lack of an alignment (‘NULL’-alignment) was considered a ‘negative’. For an alignment to be considered a ‘true negative’, English *when* needed to have a NULL-alignment in Huichol in cases where Huichol does not use a subjunction to render the WHEN-clause, but expresses temporal subordination morphologically. Conversely, ‘false negatives’ corresponded to any NULL-alignment which should have been aligned to a ‘when’ word in Huichol. ‘False positives’, then, were considered cases in which *when* was aligned to a token in Huichol, despite the language using a morphological subordination strategy (i.e. switch-reference) or an independent clause, rather than a *quepaucua*-(‘when’)-clause. Finally, ‘true positives’ corresponded to all ‘when’ instances correctly aligned to a ‘when’ word in Huichol.

the two usage points in each pair. Multidimensional scaling (MDS) was then used to reduce the resulting Hamming-distance matrix to two dimensions, which were then treated as coordinates to plot the semantic map of WHEN as shown in Figure 2. Each dot in the semantic map represents a context for WHEN (i.e., a New Testament verse), and the farther apart two dots are, the more different their semantics is assumed to be, and the more likely they are to be encoded by different linguistic means cross-linguistically.




Figure 2: Unlabelled semantic map of WHEN.

Benchmarking As a form of evaluation of the methods and results, this experiment leveraged detailed typological and grammatical descriptions of the morphological system of one particular Latin American language, Huichol (or Wixárika). Huichol is among the several Latin American languages that show a clear division of labor between lexified and non-lexified WHEN-clauses (Pedrazzini 2023). In particular, Huichol uses *switch-reference* marking, a morphological system for tracking referents in an ongoing discourse (Roberts 2017, 538). In a ‘canonical’ switch-reference system (cf. Haiman and Munro 1983, ix), a clause is marked to signal whether its subject is co-referential or not with the subject of another, usually adjacent, clause, even though switch-reference has now long been shown to serve a much broader purpose than merely signaling referential (non-)identity (cf. Stirling 1993; McKenzie 2012, 2015a,b; Keine 2013). With subject co-reference, a same-subject marker is used (SS), else a different-subject marker is employed (DS). Switch-reference is overwhelmingly present in languages that allow and use *clause chaining*, which is the possibility of asyndetically stacking up several *deranked* verb forms (Stassen

1985; Croft 2002; Cristofaro 1998, 2019), that is, lacking marking of one or more tense, aspect, or mood distinctions compared to independent clauses in the same language, to signal their status as ‘medial’ clauses or ‘converbs’. Switch-reference marking is well-known to serve that purpose particularly commonly among South American languages (cf. van Gijn 2016). In other words, by capturing switch-reference markers, we also capture the morphological means (i.e. the n -grams, or most common morphemes) that signal subordination, in our case, specifically, temporal clauses. (1) is a Huichol example of canonical switch-reference from our dataset, where switch-reference markers are used on the dependent verb to signal its subordinate status, where the English version has a *when*-clause in both cases.⁷

(1) Huichol/Wixárika (Uto-Aztecan)

- a. *Hesüana me-’u’-axüa-cu*
to.him 3.PL.SBJ-VIS-arrive.PL-DS
müpaii
thus
ti-ni-va-ru-ta-hüave
DISTR-NARR-3.PL.NSBJ-PL-SG-say
‘When they came to him he said to them’
(Acts 20:18)
- b. *Hesüana me-’u’-axüa-ca*
to.him 3.PL.SBJ-VIS-arrive.PL-SS
müme, müpaii
men thus
me-te-ni-ta-hüave
3.PL.SBJ-DISTR-NARR-SG-say
‘When the men had come to him they said’
(Luke 7:20)

Huichol additionally has a lexified ‘when’ subordinator (*quepaucua*), in which case switch-reference marking is absent, as in (2).

- (2) *Mericiüsü quepaucua yemuri-sie*
then when mountain-LOC
m-a-ca-ne, teüteri yumüiretii
AS2-PRO-down-arrive.PFV, people many
me-ca-n-i-veiya-caitiüni
3.PL.SBJ-NARR-NARR-3.SG.OBJ-follow-IPFV
‘When he came down from the mountain, great
crowds followed him’ (Matthew 8:1)

The concurrent presence of both a lexified connector and easily isolable morphemes for morphological subordination makes the language an ideal

⁷In the Huichol examples, the spelling of the Bible translation in Mayer and Cysouw’s (2014) corpus was kept. Note, however, that this is not the most common orthography found in most studies on Huichol today.

initial benchmark for experimenting with automatically detecting morphological *and* lexified markers of temporal subordination in the parallel corpus. As a form of evaluation for the character n -gram search system described below, the Huichol translation of the New Testament was enriched with annotation for different switch-reference markers. The markers were identified by using existing descriptions of Huichol switch-reference (i.e. Comrie 1983, 1982; Bierge 2017). The language has easily isolable switch-reference morphemes, namely *-ku* and *-ka* (spelled as *-cu* and *ca* in our dataset), for ‘different-subject’ and ‘same-subject’ marker, so the placeholders DS and SS were inserted before any word in the Huichol text ending with the respective forms, thus allowing the alignment model to capture the placeholders as dummy subordinators. Based on the annotated dataset, the location of SS and DS markers in the semantic map (Figure 3) can be compared with the location of morphological markers identified automatically via character n -gram search (Figure 4 in Section 3).




Figure 3: Probabilistic semantic map of WHEN, showing the location of lexified subordinators and switch-reference markers in Huichol after direct annotation (used as benchmark).

N-gram search Character n -grams were leveraged to identify potential morphological markers that are highly correlated with English *when*-clauses in our dataset, *in addition* to lexified means. As mentioned in the Introduction, the identification of potentially meaningful n -grams (i.e. those expressing a particular meaning of WHEN) is based on the approach by Asgari and Schütze (2017), albeit with additional steps and different n -gram ranges. Similarly to Asgari and Schütze (2017), χ^2 is used as a score of association between a ‘head pivot’

(in our case always *when*) and a character n -gram, and it is calculated based on how many times *when* is aligned to a word containing that n -gram, how many times it is aligned to other n -grams and the frequency of both *when* and the n -gram. The raw alignments by SymGIZA++ were used as a starting point to identify tokens on which the n -gram search should be carried out. The following steps were followed to subsequently refine the parallel dataset with potentially meaningful n -grams:

1. a bespoke list of stopwords in English was established, based on their being either extremely frequent (*Jesus, Herod, Peter, Paul*) or very likely to introduce noise in a study on temporal subordination because of their distributional overlap with subordinators in terms of absolute position in a sentence (*and, behold, then*). χ^2 was used to find highly associated forms and parallel forms with an associated p -value of 0 were removed from the target language;
2. associations were identified between *when* and all tokens aligned to *when* by SymGIZA++. Only tokens with the highest score and a p -value of 0 were kept as they were and did not undergo the next steps;
3. using spaCy's (Honnibal and Montani 2017) English model `en_core_web_sm`, the English source text was automatically annotated for syntactic dependency to identify the head of the token *when*. This allowed for the *verb* of the *when*-clause to be extracted and the parallel verb in the translation to be identified. This choice was informed by the observation that languages marking subordination on the verb itself (i.e. non-lexified WHEN-clauses) are much more likely to have an empty token `<NOMATCH>` aligned to English *when* rather than the verb itself, so that the latter must be included in the search for meaningful character n -grams associated with *when*;
4. associations were identified between *when* and n -grams of any size between 2 and 9 for all remaining tokens aligned to either *when* or its head verb;
5. the top-scoring 200 n -grams (by χ^2) were then sorted by the number of times *when* was found to cooccur with the n -gram. The top-

scoring 20 n -grams among the latter were then extracted as potentially meaningful n -grams;

6. the 20 extracted n -grams were clustered to attempt capturing groups of n -grams that are likely to be allomorphs of the same morpheme. Clustering was done using DBSCAN after converting the list of n -grams to a matrix of TF-IDF features. DBSCAN was selected after comparison with several other clustering algorithms (i.e. K-Means, K-Means++, Agglomerative Clustering, and Gaussian Mixture Modelling).
7. Each cluster of n -grams was assigned the placeholder label `ngram_1...ngram_N`, where N is the number of potentially meaningful n -gram clusters found for any given language.

Geostatistical interpolation Ordinary Kriging was then used to interpolate the linguistic items (i.e. the parallel token, if any, to *when*, or the n -gram placeholder label) used in each data point by each language in the dataset, to look for semantically relevant cross-linguistic dimensions. The Kriging model was implemented using the PyKrig library (Müller et al. 2023), with a Gaussian variogram model, a single averaging bin for the variogram (`nlag`), and `coordinates_type` set to `geographic`. The optimal range, sill, and nugget values for the Kriging models were set through a trial-and-error calibration process. Different combinations of these parameters were tested, and the ones used to produce the maps presented in Section 3 were chosen based on the interpretability of the resulting contour maps, with particular attention to the map for Huichol, thanks to the additional automatic annotation performed on the language using external knowledge bases. The contour levels generated through Kriging were normalized between 0 and 1 to facilitate the interpretation of the relative intensity of a linguistic means in the semantic space so that the closer the contour level to 1, the more intense the concentration of the respective means in the area. In the maps in Section 3, contours are plotted at all levels between 0.8 and 1.

The advantage of employing a geostatistical approach, such as Ordinary Kriging, for mapping language patterns is its ability to account for spatial autocorrelation (cf. Getis 2008), which facilitates the nuanced weighting of variables based on their prevalence and intensity across geograph-

ical space. While one linguistic means might be more widespread in terms of raw occurrence count in a given region of the semantic map, Kriging allows us to discern the spatial intensity of competing means. This, in turn, can clarify whether other means, despite being less prevalent overall, are more concentrated in that area and therefore more directly representative of the meaning associated with the respective space in the semantic map.

In the Kriging maps in Section 3, the placeholders for the n -grams are used instead of the actual list of n -grams.⁸

3 Results

Huichol Figure 4 shows the Kriging map generated from the Huichol data automatically refined with the n -gram search method. This can be compared with the labeled map in Figure 3, which, as explained in the previous section, is instead based on the Huichol data directly annotated with switch-reference markers as presented in typological descriptions of the language.




Figure 4: Kriging map of WHEN for Huichol.

Kriging detected relatively clearly separate areas (i.e. contexts or usage points) for lexified means (*quepaucua*), clustering at the bottom right of the map, and non-lexified means, corresponding to $ngram_1$ and $ngram_2$ in the map and clustering at the top of the map. *NOMATCH* indicates the absence of a parallel to English *when*, which suggests either a misalignment or the usage of a non-subordinate construction (e.g. an independent clause or a prepositional phrase, e.g. ‘during din-

⁸The reader can find which n -grams each group contains for any given language in the [associated repository](#).

ner’). It is clear that the two automatically identified groups of n -grams, $ngram_1$ and $ngram_2$, in the Huichol map correspond to DS and SS markers respectively. The $ngram_1$ group includes *u, su, usu, cusu, icusu, ricusu, ericusu*, whereas $ngram_2$ includes *ca, aca, eca, ieca, yaca, iyaca, xeyyaca, eiyaca, nieca*, which match the known switch-reference markers *-ku* and *-ka* (spelled as *-cu* and *ca* in our dataset) for DS and SS respectively (Comrie 1983).

Based on the Huichol results, automatic word-alignment combined with the n -gram search method achieves a precision of 0.90, recall of 0.99, and F1-score of 0.94, calculated upon comparison with another manually annotated random sample of 300 English-Huichol WHEN-clauses with added switch-reference distinctions (i.e. English *when* was manually aligned to either *quepaucua* ‘when’, DS, or SS).

Switch-reference languages A clear validation of our method comes from the Quechuan languages in our dataset. According to van Gijn (2016, 168-169), all Quechuan languages have switch-reference marking, albeit with some differences in the markers used and their semantic scope. A closer inspection of the maps reveals that all Quechuan languages in our dataset show, in fact, a clear division of labor between the bottom and top of the map. Most commonly, the former is a *NOMATCH* area, whereas the top areas are instead most clearly under the scope of switch-reference markers. This is clearly the case, for example, in Ambo-Pasco Quechua (Figure 5a), from Peru, where the $ngram_1$ group at the top of the map includes *r, ar, ur, cur, ycur, aycur, car*, all of which contain the distinctive *-r* SS marker of some Quechuan I subgroups (cf. van Gijn 2016, 168).

Another example is the map for Bolivar-North Chimborazo Highland Quichua (Figure 5b), from Ecuador. In this case, an $ngram_1$ Kriging area was detected alongside a potentially lexified subordinator *ña*. The n -gram group includes *aca, paca, hpaca, shpaca, ushpaca, ashpaca*, where the Quechuan II SS marker, */s/*, spelled *sh*, can be discerned (van Gijn 2016, 171).⁹

A similar split, where the top area of the map is dominated by a n -gram group, is also found outside of Quechuan. This is the case, for instance, of Cavineña (Figure 5c), a Pano-Tacanan language of

⁹The *-ca* ending is, in all likelihood, a personal ending that is particularly frequent in the source text.

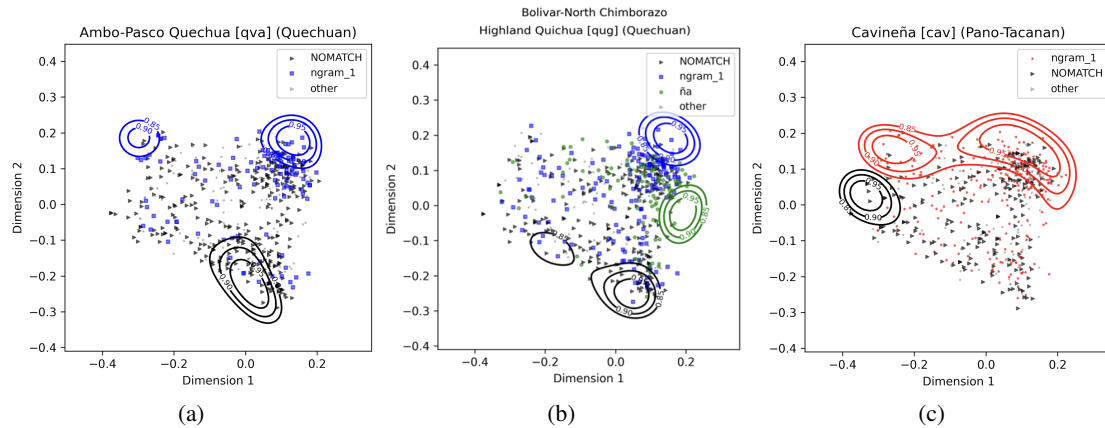


Figure 5: Kriging maps of WHEN for three Latin American languages.

the Amazonian plains of northern Bolivia, where *ngram_1* includes *u*, *su*, *tsu*, *atsu*, *aatsu*, *catsu*, *baatsu*, *acatsu*, *bacatsu*, *itsu*, where the SS marker *-tsu* (cf. Guillaume 2008, 2011) can be seen.

The semantic maps for several other varieties from different language families show a division of labor similar to the Huichol one, between lexified means at the bottom of the map and *n*-gram groups (i.e. likely morphologically encoded WHEN-clauses) at the top of the map, as in Chuy (Mayan, Guatemala; Figure 6a), Comaltepec Chinantec (Otomanguean, Mexico; Figure 6b), or Terena-Kinikinao-Chane (Arawakan, Bolivia; Figure 6c).

Beyond switch-reference The integration of character *n*-grams to the semantic map of when was primarily driven by the aim of capturing morphological means of marking generic temporal subordination, which these examples from Latin American languages indicate as promising, especially in light of the known switch-reference markers captured in the maps. However, as mentioned in Section 1, there is great linguistic variation in the Latin American and Caribbean region and the new semantic maps helped capture more than just *n*-gram groups overlapping with the switch-reference markers in Huichol or Quechuan languages. Several languages, for instance, show an inverted pattern to the Huichol one, with a lexicalized means at the top of the map and an *n*-gram area at the bottom, as in Ticuna (Ticuna-Yuri, Western Amazon; Figure 6d) or Lomeriano-Ignaciano Chiquitano (Chiquitano, Bolivia; Figure 6e).

Yet others only use one *n*-gram for both the bottom and top areas, as in Tabasco Chontal (Mayan, Mexico; Figure 6f), or only lexified means, as in San Mateo del Mar Huave

(Huavean/Isolate, Mexico; Figure 6g), Nivaclé (Matacoan, Argentina and Paraguay; Figure 6h), Kaqchikel (Mayan, Guatemala; Figure 6i), Guerrero Amuzgo (Otomanguean, Mexico; Figure 6j), Pichis Ashéninka (Arawakan, Peru; Figure 6k), and Chamacoco (Zamucoan, Paraguay; Figure 6l).

4 Conclusion & Future Work

Summary and findings This paper has presented probabilistic semantic maps of WHEN-clauses based on a parallel corpus of New Testament translations in Latin American and Caribbean languages. The rationale behind this study was the observation that WHEN-clauses in the Latin American region are often encoded *morphologically* (exclusively or predominantly so, i.e. in addition to lexified subordinators), which in previous token-based experiments (i.e. based only on full-token correspondences between languages) represented one of the main hurdles for the detection of systematic cross-linguistic variation in the expression of generic temporal subordination.

It built on previous approaches based on correspondences between a source word (English *when*) and character *n*-grams, using association measures to detect meaningful groups of *n*-grams that are likely to represent a particular morphological marker encoding temporal subordination in each target language. The approach has yielded results that are clearly helpful in identifying morphologically-encoded WHEN-clauses in languages where switch-reference markers (same-subject or different-subject marking) are employed to mark a predicate as subordinate to their matrix clause. The identification of groups of *n*-grams as switch-reference markers in some of the languages in the corpus was achieved by consulting descrip-

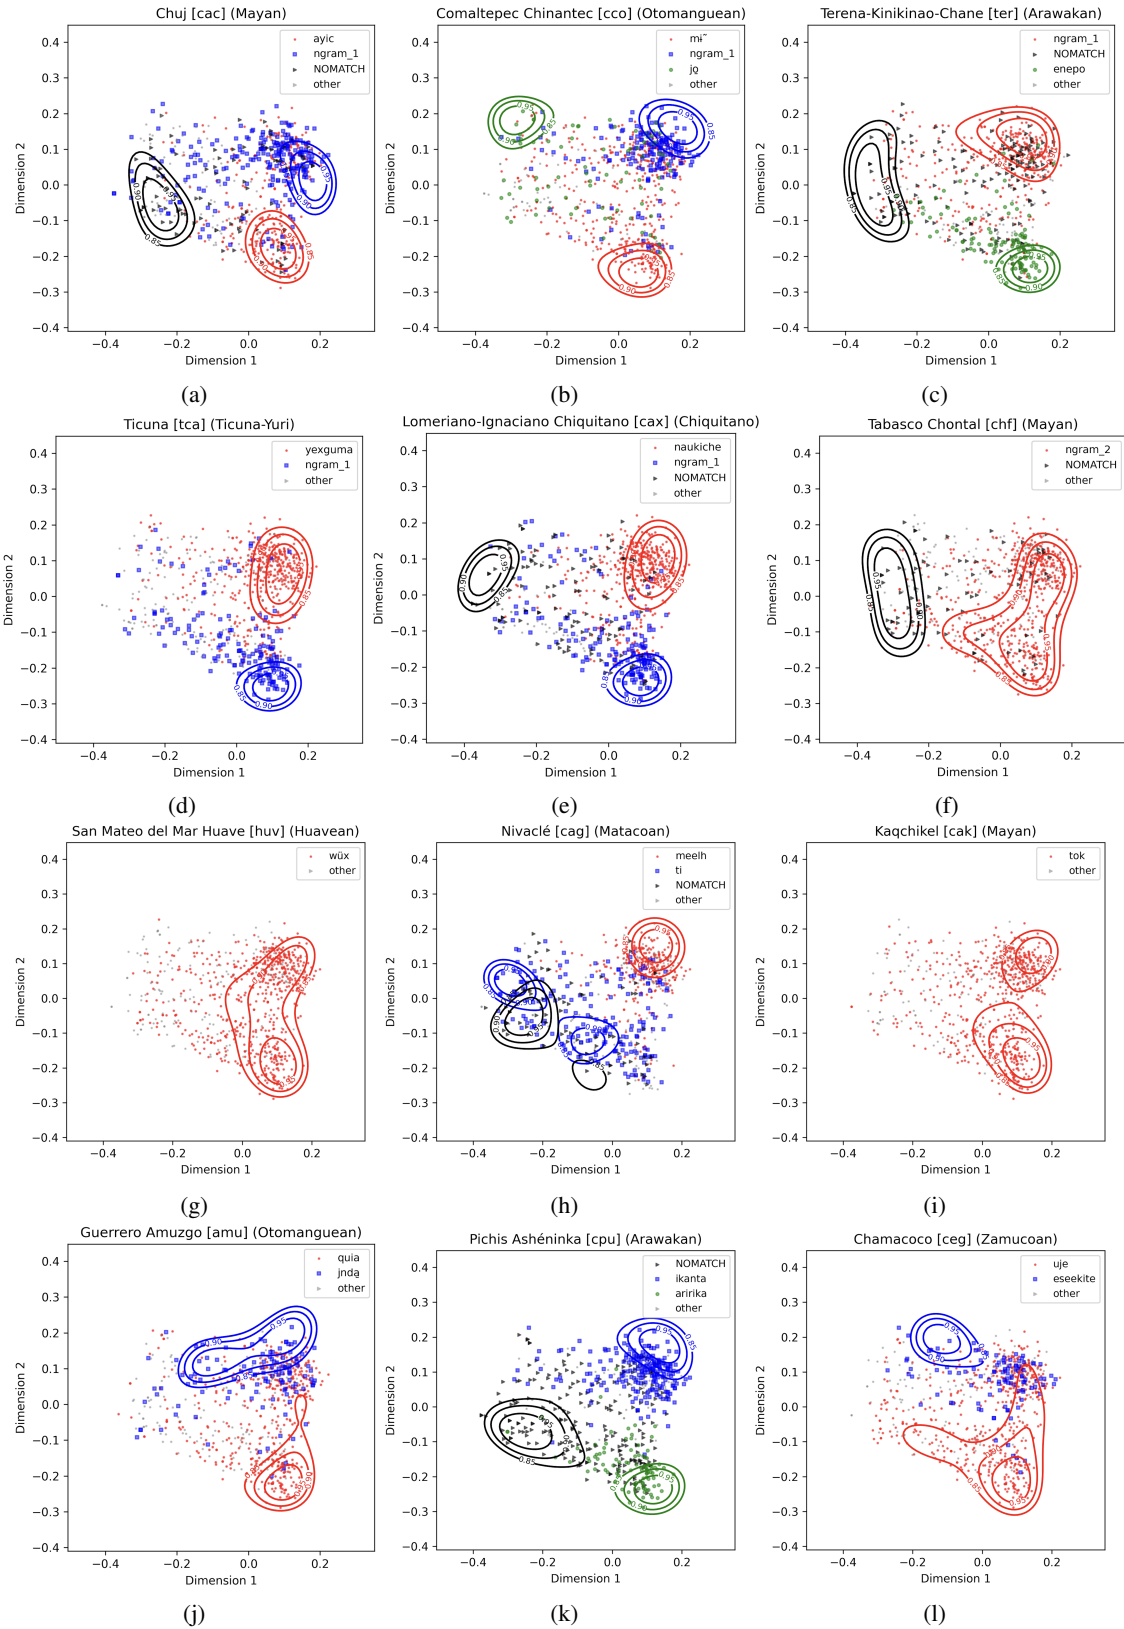


Figure 6: Kriging maps of WHEN showing some of the systematic variation in the dataset.

tive grammars and language-specific typological studies (e.g. on the Quechuan morphological system), but also because of the use of Huichol, a Mex-

ican language with switch-reference morphology, as a point of reference to build a small benchmark and optimize hyperparameters during the genera-

tion of the semantic maps.

Future research Future studies may want to experiment with different n -gram sizes and different association measures and Kriging parameters, as well as use languages other than Huichol as benchmarks for the calibration of the Kriging models. Languages showing an opposite pattern to that of Huichol (i.e. a lexified means where Huichol has a morphological means, and vice versa) would particularly benefit from a close-reading evaluation to ascertain whether the method did manage to capture morphologically-expressed *when*-clauses as accurately as their opposite pattern.

Finally, the semantic dimensions in the maps have not been fully analyzed, and future studies will take a systematic approach to identifying clusters of observations that are frequently co-expressed, whether morphologically or lexically, across the languages of the corpus, and will establish whether such clusters represent cross-linguistically relevant gram types.

Limitations

The main limitation of this experiment is that evaluation, including hyperparameter optimization for the Kriging models, was based on one particular language, Huichol, because of the well-studied subordination system and the presence of a lexified subordinator in addition to the widely employed morphological means (switch-reference). Moreover, not only is switch-reference only one of the several attested morphological means to convey generic temporal subordination cross-linguistically, but there are also major differences between switch-reference systems (both in terms of the set of markers available to a language, but also their range of functions). The hyperparameters tuning based on Huichol has likely introduced some bias towards languages that have a similar system (i.e. one lexified counterpart to English *when* alongside switch-reference morphology), potentially obscuring other relevant typological dimensions (e.g. systematic clause-bridging marking).

The n -gram approach identifies *groups* of character n -grams, but does not yet provide a straightforward way of selecting one particular set of characters as the representative morpheme from a series of potential allomorphs. A tentative solution could be extracting the shortest allomorph, or the allomorph representing the common denominator among all n -grams in a set. However, this has not been tested

and we have simply numbered each group of n -grams while keeping track of what forms each group contains for subsequent easier retrieval and inspection, if needed.

Acknowledgements

This work was supported by the Ecosystem Leadership Award under the EPSRC Grant EP/X03870X/1 & The Alan Turing Institute, particularly the Turing Research Fellowship scheme under that grant.

References

- Ehsaneddin Asgari and Hinrich Schütze. 2017. *Past, present, future: A computational investigation of the typology of tense in 1000 languages*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 113–124, Copenhagen, Denmark. Association for Computational Linguistics.
- Stefanie Ramos Bierge. 2017. *Clause Types and Transitivity in Wixárika (Huichol): a Uto-Aztecan Language*. Ph.D. thesis, University of Colorado.
- Bernard Comrie. 1982. *Grammatical relations in Huichol*. In *Studies in Transitivity*, pages 95–115. Brill, Leiden, The Netherlands.
- Bernard Comrie. 1983. *Switch-reference in Huichol: A typological study*. In John Haiman and Pamela Munro, editors, *Switch Reference and Universal Grammar*, pages 17–38. John Benjamins, Amsterdam/Philadelphia.
- Sonia Cristofaro. 1998. *Deranking and Balancing in Different Subordination Relations: a Typological Study*. *Sprachtypologie und Universalienforschung*, 51:3–42.
- Sonia Cristofaro. 2019. *Subordination*. Oxford University Press, Oxford.
- William Croft. 2002. *Typology and Universals*, 2nd edition. Cambridge University Press, Cambridge.
- William Croft and Keith T. Poole. 2008. *Inferring universals from grammatical variation: Multidimensional scaling for typological analysis*. *Theoretical Linguistics*, 34(1):1–37.
- Östen Dahl and Bernhard Wälchli. 2016. *Perfects and iamitives: two gram types in one grammatical space*. *Letras de Hoje*, pages 325–348.
- Barrie Evans. 2017. *A teaching grammar of Pular*. Ms.
- Arthur Getis. 2008. *A history of the concept of spatial autocorrelation: A geographer's perspective*. *Geographical Analysis*, 40(3):297–309.
- Antoine Guillaume. 2008. *A Grammar of Cavineña*. Mouton de Gruyter, Berlin.

- Antoine Guillaume. 2011. [Subordinate clauses, switch-reference, and tail-head linkage in Cavineña narratives](#). In Pieter Muysken Rik van Gijn, Katharina Haude, editor, *Subordination in Native South American languages*, pages 109–140. John Benjamins, Amsterdam/Philadelphia.
- John Haiman and Pamela Munro. 1983. [Introduction](#). In John Haiman and Pamela Munro, editors, *Switch Reference and Universal Grammar*, pages ix–xv. John Benjamins, Amsterdam.
- Harald Hammarström, Robert Forkel, Martin Haspelmath, and Sebastian Bank. 2023. [Glottolog 4.8](#). Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Dag Haug and Nilo Pedrazzini. 2023. [The semantic map of when and its typological parallels](#). *Frontiers in Communication*, 8.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.
- Marcin Junczys-Dowmunt and Arkadiusz Szał. 2012. [SyMGiza++: Symmetrized Word Alignment Models for Machine Translation](#). In *Security and Intelligent Information Systems (SIIS)*, volume 7053 of *Lecture Notes in Computer Science*, pages 379–390, Warsaw, Poland. Springer.
- Stephan Keine. 2013. [Deconstructing switch-reference](#). *Natural Language and Linguistic Theory*, 31:767–826.
- Natalia Levshina. 2019. [Token-based typology and word order entropy: A study based on Universal Dependencies](#). *Linguistic Typology*, 23(3):533–572.
- Natalia Levshina. 2022. [Corpus-based typology: applications, challenges and some solutions](#). *Linguistic Typology*, 26(1):129–160.
- Thomas Mayer and Michael Cysouw. 2014. [Creating a massively parallel Bible corpus](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3158–3163, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Andrew McKenzie. 2012. [The role of contextual restriction in reference-tracking](#). Ph.D. thesis, Amherst: University of Massachusetts.
- Andrew McKenzie. 2015a. [A survey of switch-reference in North America](#). *International Journal of American Linguistics*, 81:409–448.
- Andrew McKenzie. 2015b. [Austinian situations and switch-reference: The role of context in reference-tracking](#). Ms.
- Sebastian Müller, Roman Yurchak, Benjamin Murphy, nannau, Malte Ziebarth, Sudipta Basak, Marcelo Albuquerque, Mark Vrijlandt, Matthew Peveler, Daniel Mejía Raigosa, Harry Matchette-Downes, Jordan Porter, Rhilip, Scott Staniewicz, Will Chang, and kvanlombeek. 2023. [Geostat-framework/pykrige: v1.7.1](#).
- Sebastian Nordhoff and Harald Hammarström. 2011. [Glottolog/Langdoc: Defining dialects, languages, and language families as collections of resources](#). In *Proceedings of ISWC 2011*.
- Simon E. Overall. 2014. [Clause-chaining, switch-reference and nominalisations in Aguaruna \(Jivaroan\)](#). In Rik van Gijn, Jeremy Hammond, Dejan Matic, Saskia van Putten, and Ana Vilacy Galucio, editors, *Information Structure and Reference Tracking in Complex Sentences*, pages 309–340. John Benjamins, Amsterdam/Philadelphia.
- Simon E. Overall. 2016. [Switch-reference and case-marking in Aguaruna \(Jivaroan\) and beyond](#). In Rik van Gijn and Jeremy Hammond, editors, *Switch Reference 2.0*, page 453–472. John Benjamins, Amsterdam/Philadelphia.
- Nilo Pedrazzini. 2023. [A quantitative and typological study of Early Slavic participle clauses and their competition](#). Ph.D. thesis, University of Oxford.
- John R. Roberts. 2017. [A Typology of Switch Reference](#). In Alexandra Y. Aikhenvald and R. M. W. Editors Dixon, editors, *The Cambridge Handbook of Linguistic Typology*, Cambridge Handbooks in Language and Linguistics, page 538–573. Cambridge University Press.
- Leon Stassen. 1985. *Comparison and Universal Grammar*. Basil Blackwell, Oxford.
- Lesley Stirling. 1993. *Switch-Reference and Discourse Representation*. Cambridge University Press.
- Rik van Gijn. 2012. [Switch-attention \(aka switch-reference\) in South-American temporal clauses](#). *Linguistic Discovery*, 10(1):112–27.
- Rik van Gijn. 2016. [Switch reference in Western South America](#). In Rik van Gijn and Jeremy Hammond, editors, *Switch Reference 2.0*, page 153–206. John Benjamins, Amsterdam/Philadelphia.
- Rik van Gijn, Katharina Haude, and Pieter Muysken, editors. 2011. *Subordination in native South American languages*. John Benjamins, Amsterdam/Philadelphia.
- Bernhard Wälchli and Michael Cysouw. 2012. [Lexical typology through similarity semantics: Toward a semantic map of motion verbs](#). *Linguistics*, 50:671–710.