

Approaching Reflex Predictions as a Classification Problem Using Extended Phonological Alignments

Tiago Tresoldi

Department of Linguistics and Philology

University of Uppsala

Uppsala

tiago.tresoldi@lingfil.uu.se

Abstract

This work describes an implementation of the “extended alignment” model for cognate reflex prediction submitted to the “SIGTYP 2022 Shared Task on the Prediction of Cognate Reflexes”. Similarly to List et al. (2022a), the technique involves an automatic extension of sequence alignments with multilayered vectors that encode informational tiers on both site-specific traits, such as sound classes and distinctive features, as well as contextual and suprasegmental ones, conveyed by cross-site referrals and replication. The method allows to generalize the problem of cognate reflex prediction as a classification problem, with models trained using a parallel corpus of cognate sets. A model using random forests is trained and evaluated on the shared task for reflex prediction, and the experimental results are presented and discussed along with some differences to other implementations.

1 Introduction

The Special Interest Group of Linguistic Typology (SIGTYP) organized in 2022 the “SIGTYP 2022 Shared Task on the Prediction of Cognate Reflexes” (List et al., 2022b), providing the community with cognate-coded wordlists from which varying amounts of cognate sets were withheld. Participants were asked to submit models capable of predicting the missing words and morphemes from the non-withheld members. This work describes the submission named `ceot-extalign-rf`.

Reflex prediction is particularly interesting to computational historical linguistics, since most approaches to reconstruction are rooted in comparative methods operating on cognate sets (Jäger, 2019; Greenhill et al., 2020; List et al., 2018), i.e., sets of words that are assumed to derive from a common proto-word. We can therefore define a cognate as a member of a set of words (a “cognate set”) that share an etymological origin via vertical descent. Such a definition highlights how

the defining property of a cognate set is the regular sound correspondences between the words involved, even when they are not judged as “similarly sounding” by human evaluators. Homologies don’t imply homogenies, as processes like word borrowing and chance resemblance can lead to similarities, while true cognates can be very dissimilar, as in the often cited example of English “two” and Armenian “erku”, both from a reconstructed form $*dwóh_1$ (Kroonen, 2013).

The task can be seen as a form of zero-shot learning (Xian et al., 2018), where a model must learn to predict the “reflexes” of a potentially unknown ancestral word form, with no examples of the relevant cognate set provided during the training phase. When considering the landscape of machine learning methods available and the approaches so far proposed (Dinu and Ciobanu, 2014; Bodt and List, 2022; Meloni et al., 2021; Beinborn et al., 2013; Dekker and Zuidema, 2021; Fourier et al., 2021; List et al., 2022a), including other submissions to this challenge (Jäger, 2022; Celano, 2022; Kirov et al., 2022), it is possible to identify two main strategies for the task. The first one treats the problem as one of classification, potentially refining sequence results with probabilities from a character model, while the second employs sequence transformation methods, especially those akin to *seq2seq* approaches (Sutskever et al., 2014), making the task one analogous to that of “translation”. This submission is of the first kind, with results provided by random forests (Kam et al., 1995; Hastie et al., 2009), but attention should focus on the proposed method of data transformation.

The method is based on “multitiered” proposals, originally based on an idea by J.-M. List, which have been described and implemented in other works (Tresoldi et al., 2018; List, 2019b; Bodt and List, 2022; Chacon and List, 2016), including for the baseline for the shared task at hand (List et al., 2022b). It is an approach for modeling historical

linguistic relationships developed to solve problems that are not addressed by pure-correspondence approaches, such as those involving cross-site generalization and the need to capture contextual and suprasegmental information. By considering how discrete representations of phonological sequences can be an unsuitable provision, as the latter are abstract representations of both discrete and continuous multidimensional phonological domains, the technique extends base alignments with multi-layered vectors, encoding additional and derived features. Here, such an informational extension is knowingly and intentionally similar to the analytical frameworks of Firthian (Mitchell, 1975) and autosegmental phonology (Goldsmith, 1990), where sequence representation is given by more than a single string of segments.

This work first summarizes the extended alignment technique, with a focus on how it can be used to predict cognate reflexes. Then, it describes the experimental setting used for the submission, and how the extended alignment method can apply to the task at hand. It concludes with a discussion on the results and future work.

2 Materials and Methods

2.1 Materials

Data for the experiment comprised 20 standardized datasets derived from the Lexibank project (List et al., forthcoming), each encompassing a single linguistic family and providing good geographic and typological variation. The datasets were split into five partitions each, with different ratios of words kept for training and evaluation, and were provided by the task’s organizers. They provided a more comprehensive description of the data and the way they prepared it in List et al. (2022b).

The submission also extensively uses the phonological information provided by the *mipa* and *tresoldi* models of the *maniphono* library (Tresoldi, 2021), which were incorporated into the code in order to simplify installation requirements. The information provided by these models is adapted from the data of CLTS (Anderson et al., 2018; List et al., 2021), with additional mappings and structures partly described in Tresoldi (2020).

2.2 Methods

2.2.1 Multitiered Extension

It is important to distinguish between the method for representing phonological data and the actual

model for reflex prediction. The first, building upon the theoretical discussions also shown in List et al. (2022a), is the most promising element because of its innovative treatment of alignment sites: instead of just being linear components of a sequence under an alignment set, they are treated as independent records in a database. Such a database is turned into a two-dimensional matrix, from which observed and predicted features are extracted and used by common classification models. The method thus facilitates the usage of established algorithms for machine learning, allowing to reframe tasks from historical linguistics as more common tasks of classification, regression, or transformation by using conventional methods and well-researched implementations.

Let’s consider the example given in the challenge’s call, reproduced in Table 1, with three cognate sets, identified by comparable concepts, involving German, English, and Dutch.

Cognate Set	German	English	Dutch
ASH	a ʃ ε	æʃ	ɑ s
BITE	b ai s ə n	b ai t	b ei t ə
BELLY	b au x	?	b œi k

Table 1: Exemplary cognate reflexes in German, English, and Dutch. Adapted from List et al. (2022b).

The first step in producing an alignment enriched with new tiers is to perform multiple sequence alignment (List, 2012), which in this case yields three independent alignment sets, as shown in Tables 2, 3, and 4. This is a step common to most classification methods for reflex prediction.

Language	#1	#2	#3
German	a	ʃ	ε
English	æ	ʃ	-
Dutch	ɑ	s	-

Table 2: Alignment for cognate set ASH for German, English, and Dutch reflexes.

Language	#1	#2	#3	#4	#5
German	b	ai	s	ə	n
English	b	ai	t	-	-
Dutch	b	ei	t	ə	-

Table 3: Alignment for cognate set BITE for German, English, and Dutch reflexes.

Language	#1	#2	#3
German	b	au	x
English	?	?	?
Dutch	b	œi	k

Table 4: Alignment for cognate set BELLY for German, English, and Dutch reflexes.

Even without extensions, by considering each alignment site an independent observation it is possible to combine multiple sets into a single data frame. The operation “transposes” the alignments, joining them into a single frame as shown in Table 5, following the common steps of this framework (List, 2019a; Tresoldi et al., 2018; Bodt and List, 2022; List et al., 2022a).

ID	Source	German segment	English segment	Dutch segment
1	ASH.1	a	æ	ɑ
2	ASH.2	ʃ	ʃ	s
3	ASH.3	ɛ	-	-
4	BITE.1	b	b	b
5	BITE.2	ai	ai	ɛi
6	BITE.3	s	t	t
7	BITE.4	ə	-	ə
8	BITE.5	n	-	-
9	BELLY.1	b	?	b
10	BELLY.2	au	?	œi
11	BELLY.3	x	?	k

Table 5: Extended data frame, without any contextual extension, from the aligned German, English, and Dutch reflexes for the cognate sets ASH, BITE, and BELLY. Note that “Source” is only provided here for ease of exposition, and is not part of actual implementation.

Such an organization of the data is already appropriate for training statistical methods, as each observation is independent, making it possible to identify correspondences and fill gaps by imputing values. For instance, the partial match between the sites of index 4 and 9 in Table 5 strongly suggests that we should impute the missing information for the English segment as /b/, or that site 11 should be a dorsal consonant. Despite disappointing performance in most cases given by the lack of information for a proper zero-shot classification, reflex prediction would already be possible due to the

comparatively high number of informative features extracted from a small set of alignments, especially if the prediction refers to phonological traits instead of atomic segments. For example, if given a hypothetical partial cognate set with a German reflex /faus/ and an English /fout/, most statistical methods could already classify the first site in the alignment as analogous to site 2 (given the /f/ to /ʃ/ correspondence between German and English), the second one to site 10 (even though English /ou/ is not attested in this example), and the third one to site 6 (given the /s/ to /t/ correspondence), yielding a hypothetical Dutch form /sœit/.

The data frame can be extended in two ways. First, it can be enriched with information specific to each alignment site, allowing machine learning methods to generalize from observed instances (for example, learning that a correspondence applies not just to one sound, but to one or more sets of sounds) and to restrict the effect of correspondences to certain word or syllable positions¹. In Table 6, such features are added by extending sites with tiers for the “sound class” (Dolgopolsky, 1986) under the SCA model (List, 2012) and the alignment position for each segment. In an actual implementation, more site-specific information would likely be added, such as tiers derived from distinctive features (both from commonly used models, like those derived from Chomsky and Halle 1968, and from binary models designed for machine learning, like those provided by `maniphono`) and indexes related to the position in the word and in the syllable when counting either left-to-right (“index”) and right-to-left (“rindex”).

The second type of extensions addresses the fact that making each alignment site independent loses contextual information. In Table 7, two contextual tiers are added for each segment tier, one specifying the previous segment (the segment one position to the left, thus L1) and one carrying information on the following sound class (the SCA one position to the right, thus R1). There is no limit on the amount of contextual information with which each alignment site can be enriched, and, in fact, when using a complete system of phonological features it is possible to encode complex phonological information such as “the preceding syllable has a nasal consonant” or “the word ends with a front vowel”.

Depending on the size of the context window, the

¹Note that position of the alignment site in the word is explicitly tested by List et al. (2022a), who report no improvements in performance.

ID	Index	German segment	German SC	English segment	English SC	Dutch segment	Dutch SC
1	1	a	A	æ	E	ɑ	A
2	2	ʃ	S	ʃ	S	s	S
3	3	ɛ	E	-	-	-	-
4	1	b	P	b	P	b	P
5	2	ai	A	ai	A	ɛi	E
6	3	s	S	t	T	t	T
7	4	ə	E	-	-	ə	E
8	5	n	N	-	-	-	-
9	1	b	P	?	?	b	P
10	2	au	A	?	?	œi	U
11	3	x	G	?	?	k	K

Table 6: Multitiered representation extended with information on the alignment index and the corresponding SCA sound class, from the aligned German, English, and Dutch reflexes for sets ASH, BITE, and BELLY.

number, and the type of tiers used for extending the alignment, the shape of the data frame can increase to hold hundreds or thousand of features. While the human inspection of such data will rapidly become impractical, this property should not be considered an issue because, at this stage, the representation is intended for machine consumption. Methods for extracting information for human consumption should rely on these enlarged data frames, without reducing their size beforehand. Nonetheless, before carrying any kind of statistical analysis, it is recommended to perform common tasks of data preparation, such as dimensionality reduction and scaling of the features. Data standardization and normalization can ensure that features extracted from different tiers are placed on a common scale, and transformation processes such as Principal Component Analysis (PCA) (Tipping and Bishop, 1999) can improve training times and performance.

The implementation presented in this work differs from the previous ones due to the greater attention to the principles of autosegmental phonology. A general advantage of the approach (outlined already in List 2019b), is that it allows to use the propagation strategy of contextual information also for suprasegmental features, such as stress and tone, to all alignment sites where it applies. With tones, for example, instead of marking the tone as a segment-like token at the end of the syllable or

as a property of the nucleus alone, it is possible to expand the alignment with one or more tiers regarding the relevant tonality, which will apply to all the segments that make up the relevant syllable. Such information is not restricted to the tone itself, but can be decomposed into properties like “tone contour” or “starting pitch”. With complex correspondences involving suprasegmental features, machine learning methods will not need to “look ahead” for a tone token, as one or more columns will carry the relevant information in the data frame record itself. In addition, the representation structure allows more easily composing results from different information tiers: instead of establishing models that only predict segments, as atomic units, the implementation allows to predict, contemporaneously or individually, two or more tiers. whose information can be combined for the final results. For example, especially with a binary model, it is possible to predict the manner and place of articulation of a segment independently, aggregating the results into a phoneme or sound class.

2.2.2 Cognate Prediction

Cognate prediction is performed by training classifiers on the data frames prepared by the code for extending alignments. When paying attention to issues such as scaling of features, missing data, and the encoding of multistate categorical features (usually with one- or multi-hot binary encoders), any machine learning method can be used.

The task will involve two subtasks: the first for training all the classifiers that are needed, and the second for generating output in the expected format once the classifiers have been prepared.

The steps for the first subtask are:

1. **Align raw data.** This step can be performed manually or with tools for linguistic alignments, such as LingPy (List and Forkel, 2021).
2. **Prepare extended data frames.** From the alignment sets, a single data frame is generated with all the requested additional tiers, both for in-site and contextual information.
3. **Prepare the training data for each language.** The training data comprises a data frame for input variables X , including all features save for those related to the language being predicted, and a vector for the output variable y , from the appropriate “segment” column. All

ID	Index	German segment	German SC	German segment L1	German SC R1	English segment	English SC	English segment L1	English SC R1	Dutch segment	Dutch SC	Dutch segment L1	Dutch SC R1
1	1	a	A	Ø	S	æ	E	Ø	S	ɑ	A	Ø	S
2	2	ʃ	S	a	E	ʃ	S	æ	-	s	S	ɑ	-
3	3	ɛ	E	ʃ	Ø	-	-	ʃ	Ø	-	-	s	Ø
4	1	b	P	Ø	A	b	P	Ø	A	b	P	Ø	E
5	2	ai	A	b	S	ai	A	b	T	ɛi	E	b	T
6	3	s	S	ai	E	t	T	ai	-	t	T	ɛi	E
7	4	ə	E	s	N	-	-	t	-	ə	E	t	-
8	5	n	N	ə	Ø	-	-	-	Ø	-	-	ə	Ø
9	1	b	P	Ø	A	?	?	Ø	?	b	P	Ø	U
10	2	au	A	b	G	?	?	?	?	œi	U	b	K
11	3	x	G	au	Ø	?	?	?	Ø	k	K	œi	Ø

Table 7: Data frame of alignment sites extended with information on the alignment index, the corresponding SCA sound class, the preceding segment, and the following SCA sound class, from the aligned German, English, and Dutch reflexes for the cognate sets ASH, BITE, and BELLY. Missing information, such as the segment to the left of the first site in an alignment, is marked with Ø.

other features related to the language under study are discarded.

4. **Train and save classifiers.** This will result in a collection with one classifier for each language in the dataset.

Once the classifiers are ready, it is possible to perform reflex prediction with the following steps:

1. **Align raw data.** As above.
2. **Prepare extended data frames.** As above; it is highly recommended, and depending on the machine learning method it is necessary, that the set of tiers added to the alignment is equal to or a subset of the one used in training.
3. **Prepare the X data frame and generate a y prediction.** For most classification methods, y will yield a probability for different segments.
4. **Prepare the output.** Build the sequences of predicted reflexes and organize them in the expected data structure for evaluation.

3 Implementation and Results

For the shared task, data and classifiers were prepared according to the workflows described in subsection 2.2.2.

Due to design decisions aiming at testing the method more than achieving the best performing model, only few additional tiers were used for extending the base alignments. These were the segments and SCA sound classes pertinent to each language, with a left and right order of 1 and 2 (i.e., L1, L2, R1, R2), thus increasing four times the number of phonological tiers. Indexing tiers related to the position in the alignment, counting both left-to-right and right-to-left, were also added. No pruning or preemptive dimensionality reduction was performed. Random forests (Breiman, 2001; Kam et al., 1995; Hastie et al., 2009) were trained using the default implementation in *scikit-learn* version 1.1.0 (Pedregosa et al., 2011). Other classification methods were explored using samples of the datasets, yielding good performance improvements in the cases of XGBoost (Chen and Guestrin, 2016), LightGBM (Ke et al., 2017), and multi-layer perceptrons (Hinton, 1990), particularly when performing hyperparameter optimization (Akiba et al., 2019). The decision to only submit the results using the random forest method was based on the time and computational constraints to perform a full training, as well as in the goal of establishing a baseline for this implementation of the method.

Comparing with both the baseline and other submitted methods, the final results suggest much

room for improvement. Performance was in general between the Baseline and Baseline-SVM method, mostly due to gains in prediction offered by Support Vector Machines (Koutroumbas and Theodoridis, 2008), for the partitions with most information, degrading rapidly when the amount of withheld data was increased, as illustrated by the global results reported in List et al. (2022b).

The implementation performed particularly poorly with the datasets `bantudvd` (Greenhill and Gray, 2015) (as illustrated by the 0.7053 B-Cubes for the 10% partition, versus 0.7835 of the Baseline submission), `felekesemitic` (Feleke, 2021) (0.6661 B-Cubes for the 10% partition, versus 0.6925 of the baseline). It performed better with datasets `beidazihui` (Běijīng Dàxué, 1962) (as illustrated by the 0.8356 B-Cubes for the 10% partition, versus 0.7279 of the baseline), `bodtkhobwa` (Bodt and List, 2022) (0.7993 B-Cubes for the 10% partition, versus 0.7566 of the baseline), `bremerberta` (Bremer, 2016) (0.7915 B-Cubes for the 10% partition, versus 0.7187 of the baseline), `deepadungpalaung` (Deepadung et al., 2015) (0.8143 B-Cubes for the 10% partition, versus 0.7597 of the baseline), `wangbai` (Wang and Wang, 2004) (0.8326 B-Cubes for the 10% partition, versus 0.8048 of the baseline), `hattorijaponic` (Hattori, 1973) (0.8127 B-Cubes for the 10% partition, versus 0.7889 of the baseline), `listsamplesize` (List, 2014) (0.5325 B-Cubes for the 10% partition, versus 0.4048 of the baseline). Full results are available along with the submission, with performance for other datasets comparable to the baseline.

4 Discussion

Similarly to the implementation in List et al. 2022a, the method of extending alignments implemented here can be conceived as a learned data augmentation strategy, since it is not based on statistical properties of each dataset (that is, each collection of alignment sets), but on the contribution from specialized linguistic knowledge. Such a strategy does not exclude the use of purely statistical methods, allowing them to more easily find the correspondences between sets, especially when they convey phonological traits that are sparse and non sequential. Once mature, we believe that the strategy will be beneficial for most machine learning methods, as it should allow for increased generalization due to increased complexity of the data sets, incorpo-

rating informational tiers beyond phonological or sequence properties.

Extended alignments have proven their usefulness for tasks like supervised phonological reconstruction and cognate reflex prediction. There are, however, many other problems to which they could also be applied, either by replacing or by complementing existing methods. The probability of correspondence between certain phonemes, for example, can be used for fine-grained decisions in cases of doubts in the domain of sequence alignment, using information which is “local” to the languages under study along with default correspondence score matrices that are generally computed for global usage. Likewise, the method can credibly be used in the detection of loanwords (Miller et al., 2020), especially in cases of large exchange of words between two languages and of “learned loanwords”. Finally, like the other models proposed for this challenge, the method can be used to manage linguistic data, identifying forms that are not predicted by a model trained on the data itself, which may be due to either particularly rich individual histories or different amounts data noise.

Acknowledgements

The author is supported by the Cultural Evolution of Texts project, with funding from the Riksbankens Jubileumsfond (grant agreement ID: MXM19-1087:1). Theoretical work was partly developed when the author received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. #715618, “Computer-Assisted Language Comparison”).

Supplementary Material

Code and data for reproducing the analyses is deposited at <https://doi.org/10.5281/zenodo.6586772>.

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. *Optuna: A next-generation hyperparameter optimization framework*. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD ’19, page 2623–2631, New York, NY, USA. Association for Computing Machinery.
- Cormac Anderson, Tiago Tresoldi, Thiago Costa Chacon, Anne-Maria Fehn, Mary Walworth, Robert

- Forkel, and Johann-Mattis List. 2018. [A Cross-Linguistic Database of Phonetic Transcription Systems](#). *Yearbook of the Poznań Linguistic Meeting*, 4(1):21–53.
- Lisa Beinborn, Torsten Zesch, and Iryna Gurevych. 2013. Cognate production using character-based machine translation. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 883–891.
- Timotheus Adrianus Bodt and Johann-Mattis List. 2022. [Reflex prediction. A case study of Western Kho-Bwa](#). *Diachronica*, 39(1):1–38.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- Nate D. Bremer. 2016. *A sociolinguistic survey of six Berta speech varieties in Ethiopia*. SIL International, Addis Ababa.
- Beijing University Běijīng Dàxué. 1962. *Hànyǔ fāngyīn zìhuì [Chinese dialect character pronunciation list]*. Wénzì Gǎigé, Běijīng.
- Giuseppe G. A. Celano. 2022. A Transformer architecture for the prediction of cognate reflexes. In *The Fourth Workshop on Computational Typology and Multilingual NLP*, Online. Association for Computational Linguistics.
- Thiago Costa Chacon and Johann-Mattis List. 2016. [Improved computational models of sound change shed light on the history of the tukanoan languages](#). *Journal of Language Relationship*, 13(3-4):177–204.
- Tianqi Chen and Carlos Guestrin. 2016. [XGBoost: A scalable tree boosting system](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, pages 785–794, New York, NY, USA. ACM.
- Noam Chomsky and Morris Halle. 1968. *The sound pattern of English*. ERIC.
- Sujaritlak Deepadung, Supakit Buakaw, and Ampika Rattanapitak. 2015. A lexical comparison of the Palaung dialects spoken in China, Myanmar, and Thailand. *Mon-Khmer Studies*, 44:19–38.
- Peter Dekker and Willem Zuidema. 2021. [Word prediction in computational historical linguistics](#). *Journal of Language Modelling*, 8(2):295–336.
- Liviu Dinu and Alina Maria Ciobanu. 2014. [Building a dataset of multilingual cognates for the Romanian lexicon](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1038–1043, Reykjavik, Iceland. European Language Resources Association (ELRA).
- A. B. Dolgopolsky. 1986. A probabilistic hypothesis concerning the oldest relationships among the language families of northern Eurasia. In Vitalij V. Shevoroshkin, editor, *Typology, Relationship and Time*, pages 27–50. Karoma Publisher, Ann Arbor.
- Tekabe Legesse Feleke. 2021. [Ethiosemitic languages: Classifications and classification determinants](#). *Ampersand*, page 100074.
- Clémentine Fourrier, Rachel Bawden, and Benoît Sagot. 2021. [Can cognate prediction be modelled as a low-resource machine translation task?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 847–861, Online. Association for Computational Linguistics.
- John A Goldsmith. 1990. *Autosegmental and metrical phonology*, volume 1. Basil Blackwell.
- Simon J Greenhill and Russell D Gray. 2015. Bantu Basic Vocabulary Database.
- Simon J Greenhill, Paul Heggarty, and Russell D Gray. 2020. Bayesian phylolinguistics. *The handbook of historical linguistics*, 2:226–253.
- Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.
- Shirō Hattori. 1973. Japanese dialects. In Henry M. Hoenigswald and Robert H. Langacre, editors, *Diachronic, areal and typological linguistics*, number 11 in Current Trends in Linguistics, pages 368–400. Mouton, The Hague and Paris.
- Geoffrey E Hinton. 1990. Connectionist learning procedures. In *Machine learning*, pages 555–610. Elsevier.
- Gerhard Jäger. 2019. [Computational historical linguistics](#). *Theoretical Linguistics*, 45(3-4):151–182.
- Gerhard Jäger. 2022. Bayesian phylogenetic cognate prediction. In *The Fourth Workshop on Computational Typology and Multilingual NLP*, Online. Association for Computational Linguistics.
- Ho Tin Kam et al. 1995. Random decision forest. In *Proceedings of the 3rd international conference on document analysis and recognition*, volume 1416, page 278282. Montreal, Canada, August.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. [Lightgbm: A highly efficient gradient boosting decision tree](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Christo Kirov, Richard Sproat, and Alexander Gutkin. 2022. Mockingbird at the SIGTYP 2022 Shared Task: Two types of models for the prediction of cognate reflexes. In *The Fourth Workshop on Computational Typology and Multilingual NLP*, Online. Association for Computational Linguistics.
- Konstantinos Koutroumbas and Sergios Theodoridis. 2008. *Pattern recognition*. Academic Press.

- Guus Kroonen. 2013. Etymological dictionary of proto-germanic. In *Etymological Dictionary of Proto-Germanic*. Brill.
- Johann-Mattis List. 2012. SCA: Phonetic alignment based on sound classes. In Marija Slavkovic and Dan Lassiter, editors, *New directions in logic, language, and computation*, pages 32–51. Springer, Berlin and Heidelberg.
- Johann-Mattis List. 2014. Investigating the impact of sample size on cognate detection. *Journal of Language Relationship*, 11:91–101.
- Johann-Mattis List. 2019a. Automatic inference of sound correspondence patterns across multiple languages. *Computational Linguistics*, 45(1):137–161.
- Johann-Mattis List. 2019b. Automatic sound law induction (Open problems in computational diversity linguistics 3). *The Genealogical World of Phylogenetic Networks*, 6(4).
- Johann-Mattis List, Cormac Anderson, Tiago Tresoldi, and Robert Forkel. 2021. *Cross-Linguistic Transcription Systems [Dataset, Version 2.1.0]*. Max Planck Institute for the Science of Human History, Jena.
- Johann-Mattis List and Robert Forkel. 2021. *LingPy. A Python library for quantitative tasks in historical linguistics [Software Library, Version 2.6.9]*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Johann-Mattis List, Robert Forkel, Simon J. Greenhill, Christoph Rzymiski, Johannes Englisch, and Russell D. Gray. forthcoming. Lexibank, A public repository of standardized wordlists with computed phonological and lexical features. *Scientific Data*, pages 1–31.
- Johann-Mattis List, Nathan W. Hill, and Robert Forkel. 2022a. A new framework for fast automated phonological reconstruction using trimmed alignments and sound correspondence patterns. In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 89–96, Dublin. Association for Computational Linguistics.
- Johann-Mattis List, Ekaterina Vylomova, Robert Forkel, Nathan W. Hill, and Ryan D. Cotterell. 2022b. The SIGTYP 2022 shared task on the prediction of cognate reflexes. In *The Fourth Workshop on Computational Typology and Multilingual NLP*, Online. Association for Computational Linguistics.
- Johann-Mattis List, Mary Walworth, Simon J Greenhill, Tiago Tresoldi, and Robert Forkel. 2018. Sequence comparison in computational historical linguistics. *Journal of Language Evolution*, 3(2):130–144.
- Carlo Meloni, Shauli Ravfogel, and Yoav Goldberg. 2021. Ab antiquo: Neural proto-language reconstruction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4460–4473, Online. Association for Computational Linguistics.
- John E. Miller, Tiago Tresoldi, Roberto Zariquiey, César A. Beltrán Castañón, Natalia Morozova, and Johann-Mattis List. 2020. Using lexical language models to detect borrowings in monolingual wordlists. *PLOS ONE*, 15(12):1–23.
- Terence Frederick Mitchell. 1975. *Principles of Firthian Linguistics*. Longman, London.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Michael E Tipping and Christopher M Bishop. 1999. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622.
- Tiago Tresoldi. 2020. A model of distinctive features for computer-assisted language comparison. *Computer-Assisted Language Comparison in Practice*, 3(6).
- Tiago Tresoldi. 2021. Maniphono, a library for the symbolic manipulation of phonological entities. version 0.3.3. <https://github.com/tresoldi/maniphono>.
- Tiago Tresoldi, Cormac Anderson, and Johann-Mattis List. 2018. Modelling sound change with the help of multi-tiered sequence representations. In *Proceedings of the 48th Poznań Linguistic Meeting*, Poznań.
- Feng Wang and William S.-Y. Wang. 2004. Basic words and language evolution. *Language and Linguistics*, 5(3):643–662.
- Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. 2018. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2251–2265.