

Challenge Dataset of Cognates and False Friend Pairs from Indian Languages

Diptesh Kanojia^{†,♣,*}, Pushpak Bhattacharyya[†], Malhar Kulkarni[†], and Gholemreza Haffari^{*}

[†]Indian Institute of Technology Bombay, India

[♣]IITB-Monash Research Academy, India

^{*}Monash University, Australia

[†]{diptesh, pb, malhar}@iitb.ac.in, ^{*}gholamreza.haffari@monash.edu

Abstract

Cognates are present in multiple variants of the same text across different languages (*e.g.*, “hund” in German and “hound” in English language mean “dog”). They pose a challenge to various Natural Language Processing (NLP) applications such as Machine Translation, Cross-lingual Sense Disambiguation, Computational Phylogenetics, and Information Retrieval. A possible solution to address this challenge is to identify cognates across language pairs. In this paper, we describe the creation of two cognate datasets for twelve Indian languages, namely Sanskrit, Hindi, Assamese, Oriya, Kannada, Gujarati, Tamil, Telugu, Punjabi, Bengali, Marathi, and Malayalam. We digitize the cognate data from an Indian language cognate dictionary and utilize linked Indian language Wordnets to generate cognate sets. Additionally, we use the Wordnet data to create a False Friends’ dataset for eleven language pairs. We also evaluate the efficacy of our dataset using previously available baseline cognate detection approaches. We also perform a manual evaluation with the help of lexicographers and release the curated gold-standard dataset with this paper.

Keywords: cognate sets, Indian languages, cognate dataset, true cognates, false friends, gold data

1. Introduction and Motivation

Cognates are words that have a common etymological origin. For *e.g.*, the French and English word pair, *Liberté - Liberty*, reveals itself to be a cognate through orthographic similarity. Automatic Cognate Detection (ACD) is the task of detecting such etymologically related words or word sets among different languages¹. They share a formal and/or semantic affinity. Cognate words can facilitate the Second Language Acquisition (SLA) process, particularly between related languages. They have similar meanings and, therefore, can support the acquisition and/or learning of a non-native language. However, although they can accelerate vocabulary acquisition, learners also have to be aware of false friends and partial cognates. False friends are similar words that have distinct, unrelated meanings. For example, “*gift*” in German means “*poison*” unlike its English meaning. We illustrate the occurrence of one such example each for cognates and false friends’ for a pair of Indian languages in Table 1. In other cases, there are partial cognates *i.e.*, similar words that have a common meaning only in some contexts. For example, the word “*police*” in French can translate to “*police*”, “*policy*” or “*font*”, depending on the context. Dictionaries often include information about cognates and false friends, and there are dictionaries (Hammer and Monod, 1976; Prado, 1993) exclusively devoted to them.

In this paper, we describe the creation of three cognate datasets. First, we describe the digitization of one such Cognate dictionary named, “Tatsama Shabda Kosha” and its annotation with linked Wordnet IDs. With the help of a lexicographer, we perform the digitization of this dictionary. Further, we annotate the cognate sets from the dic-

	Hindi (Hi)	Marathi (Mr)	Hindi Meaning	Marathi Meaning
Cognate	ank	ank	Number	Number
False Friend	shikshA	shikshA	Education	Punishment

Table 1: An example each of a cognate pair and a false friend pair from the closely related Indian languages Hindi (Hi) and Marathi (Mr)

tionary with Wordnet synset IDs based on manual validation, where the lexicographer checks each Wordnet in the existing linked sense. Based on Kanojia et al. (2019b)’s approach, we use linked Indian Wordnets to generate true cognate data and create another cognate dataset. Additionally, we use the same Wordnet data to produce a list of False Friends and release² all the three datasets publicly. Our cognate sets can be utilized for lookup in phrase tables produced during Machine Translation to assess the quality of the translation system in question. They can be utilized as candidate translations for words, and our false friends’ list can be utilized by language learners to avoid pitfalls during the acquisition of a second language. False Friend and Cognate detection techniques can use these lists to train automatic cognate detection approaches for Indian languages. Automatic Cognate Detection (ACD) techniques help phylogenetic inference by helping isolate diachronic sound changes and thus detecting the words of a common origin (Rama, 2014). The Indo-Aryan and Dravidian language families prevalent in South Asia are examples of language families with a few ancestors (Sanskrit/Persian and Proto-Dravidian, respectively). Indian language pairs borrow a large number of cognates and false friends due to this shared ancestry. Knowing and utilising these cognates/false friends can help improve the performance of computational phylogenetics (Rama et al., 2018) as well as cross-lingual

¹Cognates can also exist in the same language. Such word pairs/sets are commonly referred to as *doublets*.

²Github Link

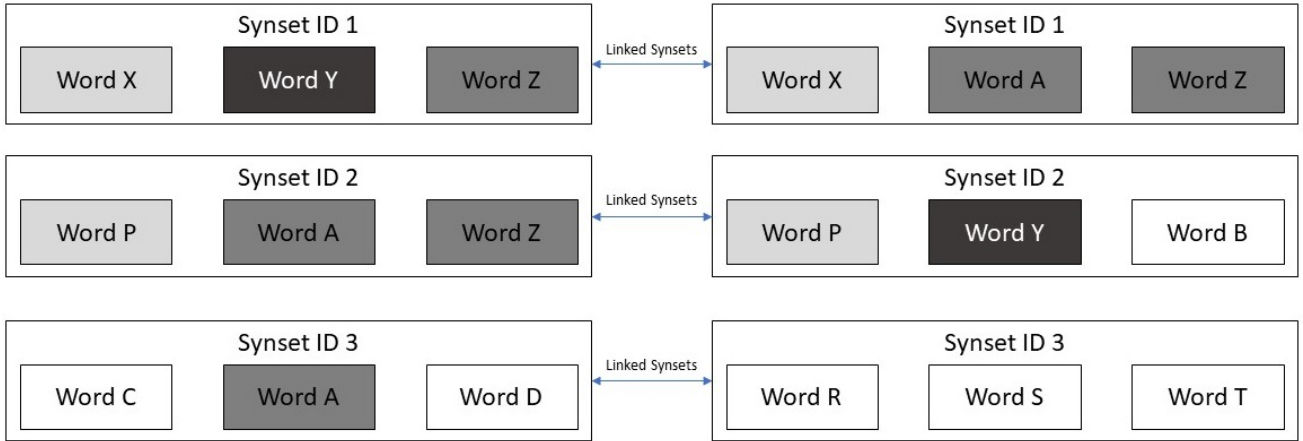


Figure 1: The difference between True Cognates (Word X and Word P), False Friends (Word Y) and Partial Cognates (Word A and Word Z) explained for creating our Datasets (D2 and D3).

information retrieval (Meng et al., 2001) in the Indian setting, thus encouraging us to investigate this problem for this linguistic area³. Some other applications of cognate detection in NLP have been sentence alignment (Simard et al., 1993; Melamed, 1999), inducing translation lexicons (Mann and Yarowsky, 2001; Tufis, 2002), improving statistical machine translation models (Al-Onaizan et al., 1999), and identification of confusable drug names (Kondrak and Dorr, 2004). All these applications depend on an effective method of identifying cognates by computing a numerical score that reflects the likelihood that the two words are cognates. Our work provides cognate sets for Indian languages, which can help the automated cognate detection methodologies and can also be used as possible translation candidates for applications such as MT.

2. Related Work

Wu and Yarowsky (2018) release cognate sets for Romance language family and provide a methodology to complete the cognate chain for related languages. Our work releases similar data for Indian languages. Such a cognate set data has not been released previously for Indian languages, to the best of our knowledge. Additionally, we release lists of false friends’ for language pairs. These cognates can be used to challenge the previously established cognate detection approaches further. Kanojia et al. (2019a) perform cognate detection for some Indian languages, but a prominent part of their work includes *manual verification and segregation* of their output into cognates and non-cognates. Identification of cognates for improving IR has already been explored for Indian languages (Makin et al., 2007). Orthographic/String similarity-based methods are often used as baseline methods for cognate detection, and the most commonly used method amongst them is the Edit distance-based similarity measure (Melamed, 1999). Research in automatic cognate detection using various aspects involves computation of similarity by decomposing

phonetically transcribed words (Kondrak, 2000), acoustic models (Mielke et al., 2012), clustering based on semantic equivalence (Hauer and Kondrak, 2011), and aligned segments of transcribed phonemes (List, 2012). Rama (2016) employs a Siamese convolutional neural network to learn the phonetic features jointly with language relatedness for cognate identification, which was achieved through phoneme encodings. Jäger et al. (2017) use SVM for phonetic alignment and perform cognate detection for various language families. Various works on orthographic cognate detection usually take alignment of substrings within classifiers like SVM (Ciobanu and Dinu, 2014; Ciobanu and Dinu, 2015) or HMM (Bhargava and Kondrak, 2009). Ciobanu and Dinu (2014) employ dynamic programming based methods for sequence alignment. Among cognate sets, common overlap set measures like set intersection, Jaccard (Järvelin et al., 2007) or XDice (Brew et al., 1996) could be used to measure similarities and validate the members of the set.

3. Dataset Creation

We create three different datasets to help the NLP tasks of cognate and false friends’ detection. In this section, we describe the creation of these three datasets for twelve Indian languages, namely Sanskrit, Hindi, Assamese, Oriya, Kannada, Gujarati, Tamil, Telugu, Punjabi, Bengali, Marathi, and Malayalam.

3.1. D1 - True Cognate Sets

The first dataset is created with the help of manual annotation. We digitize the book “Tatsama Shabda Kosh” with the help of a lexicographer. The dictionary is a collection of cognates from 15 Indian languages, but focus our work on 12 languages due to the unavailability of Wordnets for the rest of the languages. The lexicographer then also annotates each cognate set with a Wordnet sense ID after manual validation of each cognate in the twelve linked Wordnets. This helps us capture an appropriate sense for the cognate word provided via the dictionary. The annotation was performed manually with the data collected in a CSV format in a text editor. By definition, cognates are supposed to spell and mean the same. Our manual annotation using the Wordnet

³The term linguistic area or Sprachbund (Emeneau, 1956) refers to a group of languages that have become similar in some way as a result of proximity and language contact, even if they belong to different families. The best-known example is the Indian (or South Asian) linguistic area.

Language Pair	Hi-Bn	Hi-Gu	Hi-Mr	Hi-Pa	Hi-Sa	Hi-Ml	Hi-Ta	Hi-Te	Hi-As	Hi-Kn	Hi-Or
Potential Candidates	50959	81834	47718	25044	33921	18084	5203	16230	14240	12480	54014
Cognates (D2)	15312	17021	15726	14097	21710	9235	3363	936	3478	4103	11894
Percent Agreement	0.9877	0.9849	0.9838	0.9754	0.9617	0.9223	0.9033	0.9553	0.9167	0.9122	0.8833
Cohen’s kappa	0.7851	0.7972	0.8628	0.7622	0.7351	0.7046	0.6436	0.7952	0.7591	0.7953	0.8333

Table 2: Number of Potential Cognates, Number of cognates retained on both annotators’ agreement [Cognates (D2)], Percent agreement among the annotators and Cohen’s kappa score for each language pair in our dataset

Language Pair	Hi-Bn	Hi-Gu	Hi-Mr	Hi-Pa	Hi-Sa	Hi-Ml	Hi-Ta	Hi-Te	Hi-As	Hi-Kn	Hi-Or
Potential Candidates	11128	10378	14430	9062	9285	5192	1018	7149	9374	3384	5011
False Friends (D3)	4380	6204	5826	4489	2193	1076	783	699	3872	926	2602
Percent Agreement	0.8912	0.9122	0.9233	0.9500	0.9018	0.8125	0.9288	0.8492	0.8825	0.9367	0.9133
Cohen’s kappa	0.8827	0.8245	0.7815	0.9255	0.9452	0.9064	0.7244	0.8901	0.8432	0.8167	0.9548

Table 3: Number of Potential False Friends, Number of False Friend pairs retained on both annotators’ agreement [False Friends (D3)], Percent agreement among the annotators and Cohen’s kappa score for each language pair in our dataset

	Nouns	Verbs	Adjectives	Adverbs
D1	78.20	0.06	19.00	0.60
D2	76.35	2.41	20.11	1.10

Table 4: The percentage share of parts-of-speech categories in cognate datasets D1 and D2

IDs helps provide an appropriate sense to each cognate set in the dataset and thus can help automatic cognate detection techniques utilize the synset information.

This dataset consists of 1021 cognate sets with a total of 12252 words. The book consisted of a total of 1556 cognate sets, but during manual validation, 535 were found to be partial cognates and have been ignored from this dataset. The percentage share of parts-of-speech categories for Wordnet annotated cognates released is shown in Table 4. Partial cognates, as previously explained, mean differently in different contexts. NLP tasks such as Machine Translation will benefit the most from gold-standard translation candidates. Keeping the application of our dataset in mind, we ignore the inclusion of partial cognates from this dataset.

3.2. D2 - True Cognate Pairs via IndoWordnet

In their paper, Kanojia et al. (2019b) identify IndoWordnet (Bhattacharyya, 2017) as a potential resource for the task of cognate detection. They utilize deep neural network based approaches to validate their approach for cognate detection. We build this dataset using a simple orthographic similarity based approach from the IndoWordnet dataset. Our approach combines Normalized Edit Distance (NED) (Nerbonne and Heeringa, 1997) and Cosine Similarity (CoS) (Salton and Buckley, 1988) between words. We compare synset words from every language pair using NED and populate a list of cognate sets where NED score is 0.7 and above. Similarly, we populate another list of cognate sets from every language pair using a shingle (n-gram) based Cosine Similarity with the same threshold. Due to the different methods using which NED and CoS similarity techniques compute scores, both NED and CoS output a different number of word pairs. We choose a common in-

tersection of cognate pairs from among both the lists, and populate a final ‘potential cognate set’ for eleven Indian language pairs. We take the help of two lexicographers and manually validate this output. We are aided by two lexicographers for each of the language pairs of Hindi (source) - (target) ‘other Indian languages’⁴ Each lexicographer was requested to annotate whether the given word pair is cognate or not, given the Wordnet synset information, which contained the definition of the concept and an example sentence. We retain in the final dataset, only cognate pairs, which were marked to be true cognates by both annotators. We provide the language pair wise cognate data statistics, percent agreement, and Cohen’s Kappa (IAA) values for the lexicographers’ annotation in Table 2. The percentage share of parts-of-speech categories for Wordnet annotated cognates released is shown in Table 4.

3.3. D3 - False Friends’ Pairs

The creation of such a False Friends’ dataset is another one of our novel contributions in this paper. We search for false friend candidate pairs by searching for commonly spelled words through the non-parallel synsets. These candidate pairs initially included partial cognates as well, since words which are commonly spelled and belong to different senses, could occur in both the contexts. We further prune this list by ensuring that these commonly spelled words do not occur in parallel synsets and also do not occur in the corresponding linked synset on either the source or the target side. Figure 1 explains our heuristic where Word Y is a False Friend among Synsets 1 and 2, Word X and Word P are True Cognates chosen for D2, and Word A / Word Z are ignored because they are partial cognates. Once we find out such unique False Friend pairs, which are exact matches in spelling but do not occur in parallel synsets, on either side, we populate our list of false friend pairs. We populate this list for eleven language pairs where Hindi is always the source language. Please note that false friends do not follow transitivity, i.e., if A and B are false friend pairs in languages X and Y, and A and C are false friends

⁴We intended to isolate the lexicographers of clues from other language cognate pairs. Hence, we create cognate data in language pairs.

in X and Z, it is not necessary that B and C would be false friends. Hence, we populate eleven different false friend lists and take lexicographers help for each language pair to manually validate this output. Post-manual validation we choose retain the false friend pair which were annotated as false friends’ by both the annotators. We report the statistics for D3 in Table 3, which include the number of potential false friend candidates, False friend pairs after manual validation of these potential candidates, Percent agreement among both the annotators and Cohen’s Kappa (IAA) score for the annotation performed.

4. Experiment Setup for Evaluation

We evaluate the challenge posed by our datasets using the tasks of automatic cognate detection and false friend detection, based on previously available approaches. In this section, we describe the task setup and approaches which show the challenges posed by these tasks. We also discuss how our dataset is a challenging dataset for these tasks, and better approaches are needed to tackle the problems posed by a morphologically richer dataset of cognates.

We combine D1 and D2 based on Wordnet Sense IDs and remove duplicates to form a single dataset of true cognates, which we evaluate through the task of cognate detection. We use various approaches to perform the cognate detection task *viz.* baseline cognate detection approaches like orthographic similarity based, phonetic similarity based, phonetic vectors with Siamese-CNN based proposed by Rama (2016), and deep neural network based approaches proposed by Kanojia et al. (2019b). We use the same hyperparameters and architectures, as discussed in these papers. For the Orthographic similarity based approach, we use the orthographic similarity between words as a feature. For the Phonetic similarity based approach, we compute the phonetic similarity between two words using phonetic vectors available via the IndicNLP Library⁵. To classify cognate pairs, we use a simple feed forward neural network with the respective feature scores passed to a fully connected layer with ReLU activations, followed by a softmax layer (in the first two approaches). We replicate the best reported systems from Rama (2016) *i.e.*, Siamese Convolutional Neural Network with phonetic vectors as features. To replicate Kanojia et al. (2019b)’s approach, we use the Recurrent Neural Network architecture with a combination of Normalized Edit Distance, Cosine Similarity, and Jaro-Winkler Distance as reported in their paper. We have already discussed the manual validation of our datasets in the previous section, which allows us to create a more curated dataset. We use the computational approaches on this curated dataset post manual validation. For the training dataset, we use the data provided by Kanojia et al. (2019b) and create training and validation sets with an 80-20 split. We then test the aforementioned approaches on our dataset.

For the False Friends detection task, since no such dataset is available for Indian languages, we annotate the data cre-

ated by us with positive labels and divide it into train and test sets. We then add true cognates to the training dataset with negative labels since intuitively, they are the best candidates for misclassification due to common spellings just like false friends do, but in case of true cognates, they also mean the same. We test our false friends’ dataset using similar approaches with baseline features like orthographic similarity, and phonetic similarity. We use a simple Feed Forward neural network as the classifier with the respective feature scores passed to a fully connected layer with ReLU activations, followed by a softmax layer. To ensure the learning algorithm has ‘context’ available to decipher the meaning among false friends, we use the notion of distributional semantics and employ a word vectors based approach proposed by Castro et al. (2018). We use their approach to test the efficacy of our dataset and show better approaches need to be devised for morphologically richer languages.

Since the approach proposed by Castro et al. (2018) requires monolingual word embeddings to be used, we train the monolingual word embeddings using the standard Wikimedia dumps⁶. We extract text from the Wiki dumps and tokenize the data. We, then, train twelve monolingual word embedding models for each Indian language we are dealing with. In the next section, we discuss the results of our dataset evaluations.

5. Results of Our Evaluation

In table 5, we show the results for the cognate detection task. We observe that on our combined cognate dataset, the current approaches do not perform well. These approaches have reported better performance for their own datasets. In most of the cases (Language pairs), Kanojia et al. (2019b)’s approach performs better, but for the Hindi-Telugu language pair, Rama (2016)’s approach performs better. Although both the approaches perform the same for Hi-As, Hi-Mr, and Hi-Ta language pairs, the scores are still lower than what has been previously reported. We believe that these approaches perform well when on a limited dataset, moreover, when the dataset consists of words which are stripped on morphological inflections. NLP tasks such as Machine Translation and Cross-lingual Information Retrieval do not use synthetic data, which is stripped of morphological information. If the cognate detection task has to be a part of a pipeline aiding the NLP tasks, then approaches that perform the task should be robust enough to tackle a dataset such as ours. Hence, we claim our dataset to be a more challenging dataset, which should help develop better approaches.

In table 6, we report the results for the task of False Friends’ detection. We observe that the approach proposed by Castro et al. (2018) does not perform as well as it does for Spanish and Portuguese, as reported previously. We believe that this approach inherently lacks the linguistic intuition which is needed for the false friends’ detection task. Please recall that False friends are word pairs that spell the same but do not mean the same. But for the approach to

⁵https://anoopkunchukuttan.github.io/indic_nlp_library/

⁶as on 15th October, 2019

Approaches	Hi-Bn	Hi-As	Hi-Or	Hi-Gu	Hi-Mr	Hi-Pa	Hi-Sa	Hi-Ml	Hi-Ta	Hi-Te	Hi-Kn
Orthographic Similarity	0.36	0.34	0.38	0.25	0.29	0.21	0.24	0.28	0.20	0.16	0.19
Phonetic Similarity	0.42	0.38	0.39	0.29	0.32	0.24	0.25	0.31	0.24	0.22	0.25
Rama et. al. (2016)	0.65	0.71	0.61	0.67	0.72	0.47	0.53	0.62	0.53	0.65	0.57
Kanojia et. al. (2019)	0.68	0.71	0.62	0.75	0.72	0.73	0.72	0.66	0.53	0.63	0.58

Table 5: Results of the Cognate Detection Task (in terms of F-Scores) for D1+D2. We use the same architecture, features and hyperparameters as discussed in the papers for Rama et. al. (2016) and Kanojia et. al. (2019) and observe that these systems do not perform as well on our dataset, as claimed by the authors.

Language Pairs	Hi-Bn	Hi-As	Hi-Or	Hi-Gu	Hi-Mr	Hi-Pa	Hi-Sa	Hi-Ml	Hi-Ta	Hi-Te	Hi-Kn
Orthographic Similarity	0.36	0.45	0.49	0.51	0.53	0.44	0.52	0.24	0.29	0.30	0.50
Phonetic Similarity	0.60	0.66	0.67	0.62	0.59	0.69	0.61	0.54	0.48	0.50	0.57
Castro et. al. (2018)	0.66	0.64	0.59	0.65	0.69	0.73	0.72	0.65	0.52	0.69	0.64

Table 6: Results of the False Friends’ Detection Task (in terms of F-Scores) for D3. We use the same architecture, features and hyperparameters as discussed in the paper by Castro et. al. (2018) and observe that these systems do not perform as well on our False Friends’ dataset.

perform well, monolingual embeddings may not be an appropriate feature. Cross-lingual word embeddings project monolingual word embeddings into a common space and thus should be able to decipher the ‘meaning’ or the ‘sense’ of two different words better, when they belong to different languages. Given the recent advancements in word representation models, cross-lingual word embedding based models should be employed for such a task. Please also note that we do not propose a new approach for the task of False friends’ detection and hence do not perform any experimentation with cross-lingual word embeddings. However, Merlo and Rodriguez (2019) show that cross-lingual word embeddings obtained using the VecMap (Artetxe et al., 2016) approach have shown promise and can be used to obtain a semantic comparison between two words from different languages.

6. Conclusion and Future Work

In this paper, we describe the creation of a challenging dataset of true cognates which encompasses of cognates from two different sources. First, we digitize a cognate dictionary and annotate it with Wordnet Sense IDs for twelve Indian languages to create Dataset 1 (D1). We also use linked Indian Wordnets to create a true cognate dataset, as described in the paper. For both the datasets, we ensure a quality check with the help of manual validation. We report the percent agreement and Inter-annotator agreement for D1 and D2 in this paper, and for D2, we retain the cognate pairs, which were marked to be cognates by both the annotators; we were aided by two annotators for each language pair. Additionally, we release a curated list of False friends for eleven language pairs where the Hindi language is always the source, and other Indian languages are the target languages. We evaluate the efficacy of all these datasets using previously available approaches for the tasks of Cognate and False Friends’ detection. We show that these approaches do not perform as well on our dataset, given the same hyperparameters and settings as described in their papers. We discuss these results in the previous section. We also believe that this work provides a challenging gold-standard dataset for the tasks for Cognate and False Friends’ detection, which can also be used to

aid the NLP tasks of Machine Translation, Cross-lingual Information Retrieval, and Computational Phylogenetics. We hope better approaches are developed for these tasks which can perform well on our challenge dataset.

In the near future, we shall include partial cognates in our dataset creation approach and release another dataset on the same repository. Partial cognates mean different given different contexts and can confuse an NLP task. Hence, we believe it is also important to have a challenging dataset for partial cognates as well which can be evaluated via the bilingual bootstrapping approach described by Frunza (2006). We would also like to evaluate our dataset on other NLP tasks and report its efficacy in aiding the tasks of MT, CLIR, Cross-lingual Question Answering *etc.*

7. Bibliographical References

- Al-Onaizan, Y., Curin, J., Jahr, M., Knight, K., Lafferty, J., Melamed, D., Och, F.-J., Purdy, D., Smith, N. A., and Yarowsky, D. (1999). Statistical machine translation. In *Final Report, JHU Summer Workshop*, volume 30.
- Artetxe, M., Labaka, G., and Agirre, E. (2016). Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294, Austin, Texas, November. Association for Computational Linguistics.
- Bhargava, A. and Kondrak, G. (2009). Multiple word alignment with profile hidden markov models. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Student Research Workshop and Doctoral Consortium*, pages 43–48. Association for Computational Linguistics.
- Bhattacharyya, P. (2017). Indowordnet. In *The WordNet in Indian Languages*, pages 1–18. Springer.
- Brew, C., McKelvie, D., et al. (1996). Word-pair extraction for lexicography. In *Proceedings of the 2nd International Conference on New Methods in Language Processing*, pages 45–55.

- Castro, S., Bonanata, J., and Rosá, A. (2018). A high coverage method for automatic false friends detection for spanish and portuguese. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2018)*, pages 29–36.
- Ciobanu, A. M. and Dinu, L. P. (2014). Automatic detection of cognates using orthographic alignment. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 99–105.
- Ciobanu, A. M. and Dinu, L. P. (2015). Automatic discrimination between cognates and borrowings. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 431–437.
- Emeneau, M. B. (1956). India as a linguistic area. *Language*, 32(1):3–16.
- Frunza, O. (2006). Automatic identification of cognates, false friends, and partial cognates [tesis de maestría]. *Ottawa, University of Ottawa, Master on Computer Science*.
- Hammer, P. and Monod, M. (1976). English-french cognate dictionary. In *English-French Cognate Dictionary*. ERIC.
- Hauer, B. and Kondrak, G. (2011). Clustering semantically equivalent words into cognate sets in multilingual lists. In *Proceedings of 5th international joint conference on natural language processing*, pages 865–873.
- Jäger, G., List, J.-M., and Sofroniev, P. (2017). Using support vector machines and state-of-the-art algorithms for phonetic alignment to identify cognates in multi-lingual wordlists. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 1205–1216.
- Järvelin, A., Järvelin, A., and Järvelin, K. (2007). s-grams: Defining generalized n-grams for information retrieval. *Information Processing & Management*, 43(4):1005–1019.
- Kanojia, D., Kulkarni, M., Bhattacharyya, P., and Haffari, G. (2019a). Cognate identification to improve phylogenetic trees for indian languages. In *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, pages 297–300. ACM.
- Kanojia, D., Patel, K., Bhattacharyya, P., Kulkarni, M., and Haffari, R. (2019b). Utilizing wordnets for cognate detection among indian languages. In *Global Wordnet Conference (2019)*.
- Kondrak, G. and Dorr, B. (2004). Identification of confusable drug names: A new approach and evaluation methodology. In *Proceedings of the 20th international conference on Computational Linguistics*, page 952. Association for Computational Linguistics.
- Kondrak, G. (2000). A new algorithm for the alignment of phonetic sequences. In *Proceedings of the 1st North American chapter of the Association for Computational Linguistics conference*, pages 288–295. Association for Computational Linguistics.
- List, J.-M. (2012). Lexstat: Automatic detection of cognates in multilingual wordlists. In *Proceedings of the EACL 2012 Joint Workshop of LINGVIS & UNCLH*, pages 117–125. Association for Computational Linguistics.
- Makin, R., Pandey, N., Pingali, P., and Varma, V. (2007). Approximate string matching techniques for effective clir among indian languages. In *International Workshop on Fuzzy Logic and Applications*, pages 430–437. Springer.
- Mann, G. S. and Yarowsky, D. (2001). Multipath translation lexicon induction via bridge languages. In *Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies*, pages 1–8. Association for Computational Linguistics.
- Melamed, I. D. (1999). Bitext maps and alignment via pattern recognition. *Computational Linguistics*, 25(1):107–130.
- Meng, H. M., Lo, W.-K., Chen, B., and Tang, K. (2001). Generating phonetic cognates to handle named entities in english-chinese cross-language spoken document retrieval. In *IEEE Workshop on Automatic Speech Recognition and Understanding, 2001. ASRU'01.*, pages 311–314. IEEE.
- Merlo, P. and Rodriguez, M. A. (2019). Cross-lingual word embeddings and the structure of the human bilingual lexicon. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 110–120.
- Mielke, M. M., Roberts, R. O., Savica, R., Cha, R., Drubach, D. I., Christianson, T., Pankratz, V. S., Geda, Y. E., Machulda, M. M., Ivnik, R. J., et al. (2012). Assessing the temporal relationship between cognition and gait: slow gait predicts cognitive decline in the mayo clinic study of aging. *Journals of Gerontology Series A: Biomedical Sciences and Medical Sciences*, 68(8):929–937.
- Nerbonne, J. and Heeringa, W. (1997). Measuring dialect distance phonetically. In *Computational Phonology: Third Meeting of the ACL Special Interest Group in Computational Phonology*.
- Prado, M. (1993). Ntc’s dictionary of spanish false cognates. In *NTC’s dictionary of Spanish false cognates*. NTC.
- Rama, T., List, J.-M., Wahle, J., and Jäger, G. (2018). Are automatic methods for cognate detection good enough for phylogenetic reconstruction in historical linguistics? *arXiv preprint arXiv:1804.05416*.
- Rama, T. (2014). Gap-weighted subsequences for automatic cognate identification and phylogenetic inference. *arXiv preprint arXiv:1408.2359*.
- Rama, T. (2016). Siamese convolutional networks for cognate identification. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1018–1027.
- Salton, G. and Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523.

- Simard, M., Foster, G. F., and Isabelle, P. (1993). Using cognates to align sentences in bilingual corpora. In *Proceedings of the 1993 conference of the Centre for Advanced Studies on Collaborative research: distributed computing-Volume 2*, pages 1071–1082. IBM Press.
- Tufis, D. (2002). A cheap and fast way to build useful translation lexicons. In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Wu, W. and Yarowsky, D. (2018). Creating large-scale multilingual cognate tables. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.