

Subword-Level Language Identification for Intra-Word Code-Switching

Manuel Mager¹, Özlem Çetinoğlu¹ and Katharina Kann²

¹Institute for Natural Language Processing,
University of Stuttgart, Germany

²Center for Data Science, New York University, USA
{manuel.mager, ozlem}@ims.uni-stuttgart.de, kann@nyu.edu

Abstract

Language identification for code-switching (CS), the phenomenon of alternating between two or more languages in conversations, has traditionally been approached under the assumption of a single language per token. However, if at least one language is morphologically rich, a large number of words can be composed of *morphemes* from more than one language (intra-word CS). In this paper, we extend the language identification task to the subword level, such that it includes splitting mixed words while tagging each part with a language ID. We further propose a model for this task, which is based on a segmental recurrent neural network. In experiments on a new Spanish–Wixarika dataset and on an adapted German–Turkish dataset, our proposed model performs slightly better than or roughly on par with our best baseline, respectively. Considering only mixed words, however, it strongly outperforms all baselines.

1 Introduction

In settings where multilingual speakers share more than one language, mixing two or more languages within a single piece of text, for example a tweet, is getting increasingly common (Grosjean, 2010). This constitutes a challenge for natural language processing (NLP) systems, since they are commonly designed to handle one language at a time.

Code-switching (CS) can be found in multiple non-exclusive variants. For instance, sentences in different languages can be mixed within one text, or words from different languages can be combined into sentences. CS can also occur on the subword level, when speakers combine morphemes from different languages (*intra-word CS*). This last phenomenon can mostly be found if at least one of the languages is morphologically rich. An example for intra-word CS between the Ro-

(a) *ne'iwa* *pecansadoxi*
WIX MIXED
my.brother you-are.tired.PPFV

(b) *ne'iwa* *pe* *cansado* *xi*
WIX WIX ES WIX
my.brother you-are tired PPFV

‘My brother, you are tired.’

Figure 1: Intra-word CS between Spanish and Wixarika, (a) standard LID for CS, (b) our task. PPFV stands for past perfective.

mance language Spanish and the Yuto-Aztecan language Wixarika¹ is shown in Figure 1.

CS language identification (LID), i.e., predicting the language of each token in a text, has attracted a lot of attention in recent years (cf. Solorio et al. (2014); Molina et al. (2016)). However, intra-word mixing is mostly not handled explicitly: words with morphemes from more than one language are simply tagged with a `MIXED` label.

While this works reasonably well for previously studied language pairs, overlooking intra-word CS leads to a major loss of information for highly polysynthetic languages. A mixed word is unknown for NLP systems, yet a single word contains much more information, cf. Figure 1 (b). Furthermore, we find intra-word CS to be much more frequent for Spanish–Wixarika than for previously studied language pairs, such that it is crucial to handle it.

Motivated by these considerations, we extend the LID task to the subword level (from (a) to (b) in Figure 1). We introduce a new CS dataset for Spanish–Wixarika (ES–WIX) and modify an existing German–Turkish (DE–TR) CS corpus (Çetinoğlu, 2016) for our purposes. We then intro-

¹Wixarika, also known as Huichol, is a polysynthetic Mexican indigenous language.

duce a segmental recurrent neural network (SegRNN) model for the task, which we compare against several strong baselines. Our experiments show clear advantages of SegRNNs over all baselines for intra-word CS.

2 Related Work

The task of LID for CS has been frequently studied in the last years (Al-Badrashiny and Diab, 2016; Rijhwani et al., 2017; Zhang et al., 2018), including two shared tasks on the topic (Solorio et al., 2014; Molina et al., 2016). The best systems (Samih et al., 2016; Shirvani et al., 2016) achieved over 90% accuracy for all language pairs. However, intra-word CS was not handled explicitly, and often systems even failed to correctly assign the mixed label. For Nepali–English, Barman et al. (2014) correctly identified some of the mixed words with a combination of linear kernel support vector machines and a k -nearest neighbour approach. The most similar work to ours is Nguyen and Cornips (2016), which focused on detecting intra-word CS for Dutch–Limburgish (Nguyen et al., 2015). The authors utilized Morfessor (Creutz and Lagus, 2002) to segment all words into morphemes and Wikipedia to assign LID probabilities to each morpheme. However, their task definition and evaluation are on the word level. Furthermore, as this method relies on large monolingual resources, it is not applicable to low-resource languages like Wixarika, which does not even have its own Wikipedia edition.

Subword-level LID consists of both segmentation and tagging of words. An earlier approach to handle a similar scenario was the connectionist temporal classification (CTC) model developed by Graves et al. (2006). The disadvantage of this model was the lack of prediction of the segmentation boundaries that are necessary for our task. Kong et al. (2016) later proposed the SegRNN model that segments and labels jointly, with successful applications on automatic glossing of polysynthetic languages (Micher, 2017, 2018). Segmentation of words into morphemes alone has a long history in NLP (Harris, 1951), including semi- or unsupervised methods (Goldsmith, 2001; Creutz and Lagus, 2002; Hammarström and Borin, 2011; Grönroos et al., 2014), as well as supervised ones (Zhang and Clark, 2008; Ruokolainen et al., 2013; Cotterell et al., 2015; Kann et al., 2018).

3 Task and Data Description

3.1 Task Description

Formally, the task of subword-level LID consists of producing two sequences, given an input sequence of tokens $X = \langle x_1, \dots, x_i, \dots, x_{|X|} \rangle$. The first sequence contains all words and splits $X^s = \langle x_1^s, \dots, x_i^s, \dots, x_{|X|}^s \rangle$, where each x_i^s is an m -tuple of variable length $0 < m \leq |x_i|$, where $|x_i|$ is the number of characters in x_i . The second sequence is such that $T^s = \langle t_1^s, \dots, t_i^s, \dots, t_{|X|}^s \rangle$, where $|T^s| = |X^s| = |X|$ and each $t_i^s \in T^s$ is an n -tuple of tags from a given set of LID tags. An input–output example for a DE–TR mixed phrase is shown in Figure 2.

Input		⟨ ‘Yerim’, ‘seni’, ‘,’,	‘danke’, ‘Schatzzy’ ⟩
Output		⟨ (Yerim), (seni), (,),	(danke), (Schatzzy, m) ⟩
		⟨ (TR), (TR), (OTHER), (DE),	(DE, TR) ⟩

Figure 2: Subword-level LID in German–Turkish.

3.2 Datasets

German–Turkish The German–Turkish Twitter Corpus (Çetinoğlu and Çöltekin, 2016) consists of 1029 tweets with 17K tokens. They are manually normalized, tokenized, and annotated with language IDs. The language ID tag set consists of TR (Turkish), DE (German), LANG3 (other language), MIXED (intra-word CS), AMBIG (ambiguous language ID in context), and OTHER (punctuation, numbers, emoticons, symbols, etc.). Named entities are tagged with a combination of NE and their language ID: NE . TR, NE . DE, NE . LANG3. In the original corpus, some Turkish and mixed words undergo a morphosyntactic split,² with splitting points not usually corresponding to language boundaries. For the purpose of subword-level LID, these morphosyntactic splits are merged back into single words. We then manually segment MIXED words at language boundaries, and replace their labels with more fine-grained language ID tags. The total percentage of mixed words is 2.75%. However, the percentage of sentences with mixed words is 15.66%. The complete dataset statistics can be found in Table 1.

Spanish–Wixarika Our second dataset consists of 985 sentences and 8K tokens in Spanish and Wixarika. Wixarika is spoken by approximately

²E.g., separating copular suffixes from roots they are attached to, cf. Çetinoğlu and Çöltekin (2016) for details.

Tokens	All	%	Unique	Unique %
DE	3992	20.37	1360	20.43
TR	9913	50.59	4071	61.16
LANG3	112	0.57	83	1.25
AMBIG	32	0.16	23	0.18
OTHER	4345	22.17	294	4.42
NE.TR	417	2.13	275	4.13
NE.DE	389	1.99	244	3.67
NE.AMBIG	16	0.08	12	1.25
NE.LANG3	112	0.57	95	1.43
MIXED	231	1.18	183	2.75
<i>DE TR</i>	231	100.0	183	100.0

Table 1: The frequency breakdown of tokens by language IDs in the German-Turkish dataset. *All*: the total number of tokens per tag, *%*: the percentage of them with respect to the total number of tokens; *Unique*: the number of unique word types, and *Unique %*: the percentage of them with respect to the total number of unique word types.

50,000 people in the Mexican states of Durango, Jalisco, Nayarit and Zacatecas (Leza and López, 2006) and is polysynthetic, with most morphemes occurring in verbs. The data is collected from public postings and comments from Facebook accounts. To ensure the public characteristic of these posts, we manually collect data that is accessible publicly without being logged in to Facebook, to comply with the terms of use and privacy of the users. These posts and comments are taken from 34 users: 14 women, 10 men, and the rest does not publically reveal their gender. None of them have publically mentioned their age. To get a dataset that focuses on the LID task, we only consider threads where the CS phenomenon appears. We replace usernames with @username in order to preserve privacy. Afterwards, we tokenize the text, segment mixed words, and add language IDs to words and segments.

The tag set is parallel to that of German-Turkish: ES (Spanish), WIX (Wixarika), EN (English), AMBIG (ambiguous) OTHER (punctuation, numbers, emoticons, etc), NE.ES, NE.WIX and NE.EN (named entities). Mixed words are segmented and each segment is labeled with its corresponding language (ES, WIX, EN). Table 2 shows a detailed description of the dataset. The percentage of mixed words is higher than in the DE-TR dataset: 3.13% of the tokens and 4.26% of the types. The most common combination is Spanish roots with Wixarika affixes. Furthermore, 16.55% of the sentences contain mixed words.

We split the DE-TR corpus and the ES-WIX corpus into training and test sets of sizes 800:229

Tokens	All	%	Unique	Unique %
ES	4218	50.73	1527	45.76
WIX	2019	24.28	1191	35.69
EN	24	0.29	21	0.63
AMBIG	28	0.34	25	0.75
OTHER	1664	20.01	288	8.63
NE.ES	96	1.15	85	2.55
NE.WIX	77	0.93	49	1.47
NE.EN	11	0.13	9	0.27
MIXED	177	2.13	142	4.26
<i>ES WIX</i>	35	19.77	31	21.83
<i>WIX ES</i>	122	68.93	93	65.49
<i>WIX ES WIX</i>	17	9.60	31	10.56
<i>WIX EN</i>	1	0.07	1	0.07
<i>EN ES</i>	1	0.07	1	0.07

Table 2: Number of tokens classified by language tags seen in the Spanish-Wixarika dataset. We show the total number of *Tokens* per tag, their proportion (%) with the total tokens, the *Unique* word types, and the proportion (*Unique %*) of them with the total number of unique word types.

and 770:216, respectively. Error analysis and hyperparameter tuning are done on the training set via 5-fold cross-validation. We present results on the test sets. Both datasets are available at <https://www.ims.uni-stuttgart.de/institut/mitarbeiter/ozlem/NAACL2019.html>

4 Experiments

Our main system is a neural architecture that jointly solves the segmentation and language identification tasks. We compare it to multiple pipeline systems and another joint system.

4.1 SegRNN

We suggest a SegRNN (Kong et al., 2016) would be the best fit for our task because it models a joint probability distribution over possible segmentations of the input and labels for each segment.

The model is trained to optimize the following objective, which corresponds to the joint log-likelihood of the segment lengths e and the language tags t :

$$\mathcal{L}(\theta) = \sum_{(x,t,e) \in \mathcal{D}} -\log p(t, e|x) \quad (1)$$

\mathcal{D} denotes the training data, θ is the set of model parameters, x the input, t the tag sequence and e the sequence of segment lengths.

Our inputs are single words.³ As hyperparame-

³We also experimented with entire phrases as inputs, and the achieved scores were slightly worse than for word-based inputs.

	DE-TR							ES-WIX						
	Segmentation			Tagging			Char Acc.	Segmentation			Tagging			Char Acc.
	P	R	F1	P	R	F1		P	R	F1	P	R	F1	
SegRNN	60.4	46.8	53.0	78.8	60.2	74.0	72.9	75.6	62.7	68.5	85.3	70.5	77.2	84.6
BiLSTM+Seq2Seq	46.1	33.4	38.7	84.3	66.8	74.5	67.7	66.6	52.2	58.5	82.4	66.2	73.4	78.4
BiLSTM+CRF	49.4	34.5	40.6	84.3	66.8	74.5	68.0	61.1	48.4	54.0	82.4	66.3	73.4	76.7
CRFTag+Seq2Seq	12.7	6.8	8.9	27.8	15.2	19.7	37.2	47.2	31.9	38.1	63.5	43.0	51.3	69.6
CRFTag+CRF	11.0	5.9	7.7	27.8	15.2	19.7	36.8	47.6	32.4	38.6	63.5	43.0	51.3	69.4
CharBiLSTM	19.1	26.6	22.2	32.9	45.7	38.2	61.1	49.7	52.2	50.8	63.1	68.1	66.5	75.7

Table 3: Segmentation and LID test results for mixed words only.

ters we use: 1 RNN layer, a 64-dimensional input layer, 32 dimensions for tags, 16 for segments, and 4 for lengths. For training, we use Adam (Kingma and Ba, 2014).

4.2 Baselines

BiLSTM+Seq2Seq/BiLSTM+CRF Our first baselines are pipelines. First, the input text is tagged with language IDs. Language IDs of a mixed word are directly predicted as a combination of all language ID tags of the word (i.e., WIX_ES). Second, a subword-level model segments words with composed language ID tags. For word-level tagging, we use a hierarchical bidirectional LSTM (BiLSTM) that incorporates both token- and character-level information (Plank et al., 2016), similar to the winning system (Samih et al., 2016) of the Second Code-Switching Shared Task (Molina et al., 2016).⁴ For the subword level, we use two supervised segmentation methods: a CRF segmenter proposed by Ruokolainen et al. (2013), that models segmentation as a labeling problem and a sequence-to-sequence (Seq2Seq) model trained with an auxiliary task as proposed by Kann et al. (2018).

CRFTag+Seq2Seq/CRFTag+CRF Since our datasets might be small for training neural networks, we substitute the BiLSTM with a CRF tagger (Müller et al., 2013, CRFTag) in the first step. For segmentation, we use the same two approaches as for the previous baselines.

CharBiLSTM We further employ a BiLSTM to tag each character with a language ID. For training, each character inherits the language ID of the word or segment it belongs to. At prediction time, if the characters of a word have different language IDs, the word is split.

⁴For all BiLSTM models input dimension is 100 with a hidden layer size of 100. For training we use a stochastic gradient descent (Bottou, 2010), 30 epochs, with a learning rate of 0.1. A 0.25 dropout factor is applied.

4.3 Metrics

We use two metrics for evaluation. First, we follow Kong et al. (2016) and calculate precision (P), recall (R), and F1, using segments as units (an unsegmented word corresponds to one segment). We also report a tagging accuracy (Char Acc.) by assigning a language ID to each character and calculating the ratio of correct language tags over all characters.

	DE-TR			ES-WIX		
	Seg. F1	Tag. F1	Char Acc.	Seg. F1	Tag. F1	Char Acc.
SegRNN	98.7	94.0	93.6	97.8	92.5	92.4
BiLSTM+Seq2Seq	98.6	95.1	94.3	98.1	90.9	90.7
BiLSTM+CRF	98.7	94.9	94.4	97.9	87.8	90.6
CRFTag+Seq2Seq	98.4	93.7	93.1	97.7	90.4	90.1
CRFTag+CRF	98.4	93.7	93.1	97.6	90.3	90.1
CharBiLSTM	87.7	88.0	92.5	87.7	87.9	91.3

Table 4: Test set results for entire datasets.

4.4 Results and Discussion

Table 4 shows all test results for the entire datasets. We find the following: (i) For ES-WIX, SegRNN performs slightly better for tagging than the best baseline, both in terms of F1 and character accuracy. For DE-TR, SegRNN and BiLSTM+CRF are the best segmentation models, but the BiLSTM models slightly outperform SegRNN for tagging. (ii) The CRF pipelines perform slightly worse than the best word-level BiLSTM models for both datasets and all evaluations.

Table 3 shows the results of tagging and segmentation only for the mixed words in our datasets. Here, we can see that: (i) Our SegRNN model achieves the best performance for segmentation. Differences to the other approaches are $\geq 10\%$, showing clearly why these models are good for the task when the number of words belonging to two languages is high. (ii) The pipeline BiLSTM models work best for tagging of the DE-TR data with a slight margin, but underperform




Figure 3: Confusion matrices of the two best models on both datasets. The x axis represents tags seen in the gold standard, and the y axis shows the corresponding predicted tags. Values are rounded up, therefore not all columns add up to 1.

on the ES–WIX dataset as compared to the SegRNN models. (iii) Both CRFTag models achieve very low results for both segmentation and tagging. (iv) CharBiLSTM performs better than the CRFTag models on both tasks, but is worse than all other approaches in our experiments.

More generally, we further observe that recall on mixed words for the DE–TR pair is low for all systems, as compared to ES–WIX. This effect is especially strong for the CRFTag and CharBiLSTM models, which seem to be unable to correctly identify mixed words. While this tendency can also be seen for the ES–WIX pair, it is less extreme. We suggest that the better segmentation and tagging of mixed words for ES–WIX might mostly be due to the higher percentage of available examples of mixed words in the training set for ES–WIX.

Overall, we conclude that SegRNN models seem to work better on language pairs that have more intra-word CS, while pipeline approaches might be as good for language pairs where the number of mixed words is lower.

Error analysis. Figure 3 shows confusion matrices for SegRNN and BiLSTM+Seg2Seg. Both models achieve good results assigning monolingual tags (ES, WIX, DE, TR) and punctuation symbols (OTHERS). The hardest labels to classify are named entities (NE, NE.TR, NE.WIX, NE.ES), as well as third language and ambiguous tags (LANG3, EN, AMBIG). Performance on multilingual tags (DE TR, WIX ES, ES WIX, WIX ES WIX) is mixed. For DE TR, BiLSTM+Seq2Seq gets slightly better classifications, but for the ES–WIX tags SegRNN achieves better results.

Regarding oversegmentation problems, BiLSTM+Seq2Seq (0.8% for DE–TR and 2.0%

for ES–WIX) slightly underperforms SegRNN (0.7% for DE–TR and 1.13% for ES–WIX). The BiLSTM+Seq2Seq (2.4%) makes fewer undersegmentation errors for DE–TR than SegRNN (2.7%). However, for ES–WIX, SegRNN performs better with 3.81% undersegmentation errors compared to 4.2% of BiLSTM+Seq2Seq.

5 Conclusion

In this paper, we extended the LID task to the subword level, which is particularly important for code-switched text in morphologically rich languages. We further proposed a SegRNN model for the task and compared it to several strong baselines. Investigating the behaviour of all systems, we found that pipelines including a BiLSTM tagger work well for tagging DE–TR, where the number of mixed tokens is not that high, but that our proposed SegRNN approach performs better than all other systems for ES–WIX. Also, SegRNNs have clear advantages over all baselines if we consider mixed words only. Our subword-level LID datasets for ES–WIX and DE–TR are publicly available.

Acknowledgments

We would like to thank Mohamed Balabel, Sam Bowman, Agnieszka Falenska, Ilya Kulikov and Phu Mon Htut for their valuable feedback. We also want to thank Jeffrey Micher for his help with the setup of the SegRNN code. This project has benefited from financial support to Manuel Mager and Özlem Çetinoğlu by DFG via project CE 326/1-1 “Computational Structural Analysis of German-Turkish Code-Switching”, to Manuel Mager by DAAD Doctoral Research Grant, and to Katharina Kann by Samsung Research.

References

- Mohamed Al-Badrashiny and Mona Diab. 2016. Lili: A simple language independent approach for language identification. In *COLING*.
- Utsab Barman, Amitava Das, Joachim Wagner, and Jennifer Foster. 2014. Code mixing: A challenge for language identification in the language of social media. *EMNLP*.
- Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*. Springer.
- Özlem Çetinoğlu. 2016. A Turkish-German code-switching corpus. In *LREC*.
- Özlem Çetinoğlu and Çağrı Çöltekin. 2016. Part of speech annotation of a Turkish-German code-switching corpus. In *LAW-X*, Berlin, Germany.
- Ryan Cotterell, Thomas Müller, Alexander Fraser, and Hinrich Schütze. 2015. Labeled morphological segmentation with semi-markov models. In *CoNLL*.
- Mathias Creutz and Krista Lagus. 2002. Unsupervised discovery of morphemes. In *SIGMORPHON*.
- John Goldsmith. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics*, 27(2):153–198.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML*. ACM.
- Stig-Arne Grönroos, Sami Virpioja, Peter Smit, and Mikko Kurimo. 2014. Morfessor flatcat: An HMM-based method for unsupervised and semi-supervised learning of morphology. In *COLING*.
- F. Grosjean. 2010. *Bilingual: Life and Reality*. Harvard University Press.
- Harald Hammarström and Lars Borin. 2011. Unsupervised learning of morphology. *Computational Linguistics*, 37(2).
- Zellig S Harris. 1951. Methods in structural linguistics.
- Katharina Kann, Manuel Mager, Ivan Meza, and Hinrich Schütze. 2018. Fortification of neural morphological segmentation models for polysynthetic minimal-resource languages. In *NAACL-HLT*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *ICLR*.
- Lingpeng Kong, Chris Dyer, and Noah A Smith. 2016. Segmental recurrent neural networks. *Interspeech*.
- José Luis Iturrioz Leza and Paula Gómez López. 2006. *Gramática wixarika*, volume 1. Lincom Europa.
- Jeffrey Micher. 2017. Improving coverage of an inuktitut morphological analyzer using a segmental recurrent neural network. In *ComputEL*.
- Jeffrey Micher. 2018. Using the Nunavut Hansard data for experiments in morphological analysis and machine translation. In *PYLO*.
- Giovanni Molina, Fahad AlGhamdi, Mahmoud Ghoneim, Abdelati Hawwari, Nicolas Rey-Villamizar, Mona Diab, and Tamar Solorio. 2016. Overview for the second shared task on language identification in code-switched data. In *CS Workshop*.
- Thomas Müller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient higher-order CRFs for morphological tagging. In *EMNLP*.
- Dong Nguyen and Leonie Cornips. 2016. Automatic detection of intra-word code-switching. In *SIGMORPHON*.
- Dong Nguyen, Dolf Trieschnigg, and Leonie Cornips. 2015. Audience and the use of minority languages on twitter. In *ICWSM*.
- Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss. In *ACL*.
- Shruti Rijhwani, Royal Sequiera, Monojit Choudhury, Kalika Bali, and Chandra Shekhar Maddila. 2017. Estimating code-switching on twitter with a novel generalized word-level language detection technique. In *ACL*.
- Teemu Ruokolainen, Oskar Kohonen, Sami Virpioja, and Mikko Kurimo. 2013. Supervised morphological segmentation in a low-resource learning setting using conditional random fields. In *CoNLL*.
- Younes Samih, Suraj Maharjan, Mohammed Attia, Laura Kallmeyer, and Tamar Solorio. 2016. Multilingual code-switching identification via LSTM recurrent neural networks. In *CS Workshop*.
- Rouzbeh Shirvani, Mario Piergallini, Gauri Shankar Gautam, and Mohamed Chouikha. 2016. The howard university system submission for the shared task in language identification in Spanish-English codeswitching. In *CS Workshop*.
- Tamar Solorio, Elizabeth Blair, Suraj Maharjan, Steven Bethard, Mona Diab, Mahmoud Gohneim, Abdelati Hawwari, Fahad AlGhamdi, Julia Hirschberg, Alison Chang, et al. 2014. Overview for the first shared task on language identification in code-switched data. *CS Workshop*.
- Yuan Zhang, Jason Riesa, Daniel Gillick, Anton Bakalov, Jason Baldrige, and David Weiss. 2018. A fast, compact, accurate model for language identification of codemixed text. In *EMNLP*.

Yue Zhang and Stephen Clark. 2008. Joint word segmentation and POS tagging using a single perceptron. *ACL-HLT*.