

CodemixedNLP: An Extensible and Open NLP Toolkit for Code-Mixing

Sai Muralidhar Jayanthi, Kavya Nerella, Khyathi Raghavi Chandu, Alan W Black

Language Technologies Institute

Carnegie Mellon University

{sjayanth, knerella, kchandu, awb}@cs.cmu.edu

Abstract

The NLP community has witnessed steep progress in a variety of tasks across the realms of monolingual and multilingual language processing recently. These successes, in conjunction with the proliferating mixed language interactions on social media have boosted interest in modeling code-mixed texts. In this work, we present CODEMIXEDNLP, an open-source library with the goals of bringing together the advances in code-mixed NLP and opening it up to a wider machine learning community. The library consists of tools to develop and benchmark versatile model architectures that are tailored for mixed texts, methods to expand training sets, techniques to quantify mixing styles, and fine-tuned state-of-the-art models for 7 tasks in Hinglish¹. We believe this work has a potential to foster a distributed yet collaborative and sustainable ecosystem in an otherwise dispersed space of code-mixing research. The toolkit is designed to be simple, easily extensible, and resourceful to both researchers as well as practitioners².

1 Introduction

Code-mixing refers to fluid alteration between two or more languages in a given utterance. This phenomenon is ubiquitous and more natural in multilingual communities, and is highly prevalent in social media platforms. Developing tools that can comprehend mixed texts can have a multitude of advantages, ranging from socially responsible NLP applications such as moderating abusive content in social media to improve naturalness of ubiquitous technologies such as conversational AI assistants and further to develop socio-cultural studies around human cognition, such as why and when people code-mix.

NLP tools for monolingual and multilingual language processing have rapidly progressed in

the past few years; thanks to the transformer-based models such as Multilingual BERT (Devlin et al., 2019) & XLM-RoBERTa (Conneau et al., 2020), and their *pretraining* techniques. On various mixed datasets, recent studies have shown that adopting multilingual pretrained models can perform better than their previous deep learning counterparts (Pires et al., 2019; Khanuja et al., 2020; Aguilar et al., 2020; Chakravarthy et al., 2020; Jayanthi and Gupta, 2021). While this looks promising for multilingual, the same is not translated to code-mixing. Hence, a critical investigation is required to understand generalizable modeling strategies to enhance performance on mixed texts (Winata et al., 2021; Aguilar and Solorio, 2020; Sitaram et al., 2020).

At the same time, practitioners who require an off-the-shelf tool into their downstream mixed text application (eg. sentiment or language identification), currently have to resort to monolingual toolkits such as NLTK, Flair, IndicNLP and iNLTK. On the other hand, while there have been several episodic works on mixed text processing, such as proposing novel datasets or shared-tasks or training strategies, there haven't been many initiatives to collate these resources into a common setting; doing so can benefit both researchers and practitioners, thereby accelerating NLP for mixed texts.

In this work, we address some of these shortcomings by creating an extensible and open-source toolkit for a variety of semantic and syntactic NLP applications in mixed languages. Our toolkit offers-

- simple **plug-and-play command line interfaces** with fine grained control over inputs, models and tasks for developing, quantifying, benchmarking, and re-using versatile model architectures tailored for mixed texts (§ 2.1, § 2.2, § 2.3)
- easy to use single stop interfacing for a variety

¹Demo is available at: <https://bit.ly/3rzOcWb>

²The library and pretrained models are available at github.com/murali1996/CodemixedNLP.

of **data augmentation** techniques including transliteration, spelling variations, expansion with monolingual corpora etc., by leveraging a collation of publicly available tools

- a toolkit library to **import fine-tuned and ready-to-use models** for 7 different tasks in Hinglish, along with an easy-to-setup **web interface** wrapper based on flask server (§ 4)

We believe the fine grained plug and play interfacing of the toolkit can serve a multitude of purposes in both academia and industry. Such fine control over the individual components of the model can enable accelerated experimentation in training different model architectures, such as multi-tasking, representation-fusion, and language-informed modeling. This in-turn helps our understanding of utilizing pretrained transformer-models for mixed datasets. In addition, our toolkit also offers computation of metrics to quantify code-mixing such as Code-Mixing Index, Language Entropy, etc., which can be utilized to find peculiarities of low-performing subsets.

Like a curse in disguise, though code-mixing is widely prevalent and available on social media, it is accompanied with non-standard spellings, mixed scripts and ill-formed sentences are common in code-mixing. To combat this, our toolkit offers techniques to augment the training sets with multiple views of each input corresponding to the above problems.

Among many potential applications, we first demonstrate our toolkit’s utility in *benchmarking* (§ 3). In addition, we publish state-of-the-art models for different NLP tasks in Hinglish and wrap them into a command line / deployable web interface (§ 4). Our toolkit is easily extensible—practitioners can incorporate new pretrained as well as fine-tuned models, include text processors such as tokenizers, transliterators and translators, and add wrappers on existing methods for downstream NLP applications.

2 Toolkit

Our toolkit is organized into components as depicted in Figure 1. In a nutshell, an end-to-end model architecture consists of one or more `encoder` components, a component for combining encodings, and one or more `adaptor` plus `task` components.

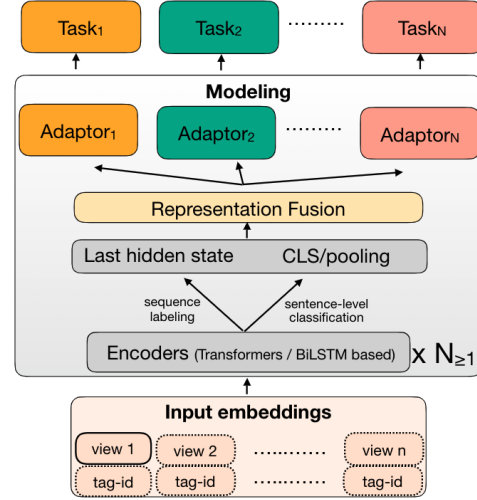


Figure 1: Customizable components in our toolkit. Marked in dashed box is an optional component.

2.1 Input Embeddings

Multi-view Integration: Tokens in mixed texts are often manifested in cross-script and mixed forms, that we refer to as *views*. This infusion motivates integration of text representations in varied forms, such as transliterated, translated, script-normalized, and tokens belonging to one of the participating languages. Especially in the context of pretrained multilingual models, this technique means extracting a holistic representation of a mixed text. To this end, the toolkit facilitates combining representations from different *views* of an input.

Text Tokenization: Motivated by some recent related works on using different word-level and sub-word-level embeddings (Winata et al., 2019; Aguilar and Solorio, 2020), our toolkit offers different tokenization methods for encoding text. Among the encoders available in our toolkit (§ 2.2), pretrained transformer-based encoders can either be tokenized using their default tokenization technique (i.e. subwords) or by using a character-CNN architecture (Boukkouri et al., 2020). LSTM-based models can take inputs in the form of tensor representations—eg. word-level FastText (Bojanowski et al., 2017) or semi-character (Sakaguchi et al., 2017) representations, or character-level representations—eg. char-BiLSTM³.

Tag-Informed Modeling: Studies in the past have shown the usefulness of language tag-aware modeling for mixed and cross-lingual

³Sequence Tagging with Tensorflow

Benchmarking transformer based models (F1 / Accuracy)										
Model	Text Classification Tasks							Sequence Tagging Tasks		
	Sentiment Classification			Aggression Identification	Hate Speech Identification	Offensiveness Identification	Youtube Comments Classification	Language Identification	NER	POS
	HIN-ENG ₁	HIN-ENG ₂	SPA-ENG [†]	HIN-ENG	HIN-ENG	TAM-ENG [†]	HIN-ENG	HIN-ENG [†]	HIN-ENG [†]	HIN-ENG
mBert	65.9 / 65.7	58.8 / 60.7	50.1 / 51.3	48.1 / 48.7	47.1 / 61.7	76.8 / 79.0	83.3 / 83.4	96.9 / 96.9	95.1 / 95.6	75.0 / 75.8
w/ Task adaptive	67.4 / 67.2	61.4 / 61.5	53.2 / 55.0	50.3 / 51.2	67.9 / 70.3	76.9 / 78.9	83.6 / 83.7	97.1 / 97.1	97.1 / 97.1	77.1 / 77.6
w/ Domain adaptive	71.4 / 71.3	62.5 / 63.0	—	50.7 / 51.4	65.5 / 70.3	—	85.2 / 85.3	97.3 / 97.3	97.2 / 97.3	75.8 / 76.3
XLM-RoBERTa	68.9 / 69.1	61.5 / 61.5	54.4 / 54.6	49.0 / 49.6	64.4 / 69.6	76.7 / 77.4	85.7 / 85.8	97.1 / 97.1	96.1 / 96.3	73.7 / 74.8
w/ Task adaptive	70.4 / 70.4	63.0 / 63.0	54.4 / 54.8	55.3 / 55.4	64.6 / 69.8	77.1 / 78.8	85.8 / 85.9	97.6 / 97.6	97.1 / 97.2	76.5 / 76.9
w/ Domain adaptive	72.1 / 72.2	65.6 / 65.7	—	56.7 / 57.1	65.2 / 70.4	—	87.3 / 87.4	97.5 / 97.5	96.9 / 97.0	74.9 / 75.7

Table 1: Results are reported for eight different tasks, namely, Sentiment Classification (HIN-ENG₁ (Patwa et al., 2020), HIN-ENG₂ (Patra et al., 2018), SPA-ENG(Aguilar et al., 2020)), Aggression Identification (Kumar et al., 2018), Hate Speech Identification (Bohra et al., 2018), Offensiveness Identification (Chakravarthi et al., 2020), Youtube Comments Classification (Kaur et al., 2019), Language Identification (Aguilar et al., 2020), Named Entity Recognition and Parts of Speech Tagging (Khanuja et al., 2020). † Implies results on *dev* split, otherwise on *test* splits.

Model	Sentiment Classification	
	HIN-ENG ₁	HIN-ENG ₂
XLM-RoBERTa	68.9 / 69.1	61.5 / 61.5
w/ multi-view integration	71.1 / 71.3	62.0 / 62.8
w/ language-tag informed	68.9 / 69.3	62.8 / 63.1
w/ fasttext-BiLSTM fusion	69.9 / 70.0	61.2 / 62.1
w/ char-BiLSTM fusion	69.3 / 69.1	62.0 / 62.3
w/ semi-char-BiLSTM fusion	69.4 / 68.9	60.0 / 60.8
w/ data noising	70.5 / 70.5	61.9 / 62.2
w/ monolingual corpora	68.9 / 69.3	68.2 / 68.3

Table 2: Results of various modelling techniques (F1 / Accuracy) when used with a pretrained transformers-based encoder. HIN-ENG₁ refers (Patwa et al., 2020) and HIN-ENG₂ refers to (Patra et al., 2018)

texts (Chandu et al., 2018; Lample and Conneau, 2019). However, their usefulness in the context of pretrained models and code-mixing is not thoroughly investigated. To this end, we offer a more generalized method in our toolkit to conduct *any* tag-aware fine-tuning, wherein representations for different kinds of tags can be added to the text representations. Examples of such tags include POS tags, Language IDs, etc.

2.2 Models

Encoders: An Encoder in our toolkit can consist of a transformer-based or BiLSTM-based architecture. Specifically, for the former, we utilize pretrained models from the HuggingFace library (Wolf et al., 2020) and the latter is implemented in Pytorch (Paszke et al., 2019).

Representation Fusion: Encodings from different encoders can be combined, and if required be augmented with (non-trainable) representations before passing through an adaptor. To combine encodings, one can either simply concat them or obtain a (trainable) weighted average, a more parameter-efficient choice than the former. Both choices are available in our toolkit.

Adaptors: An adaptor is a task-specific neural

layer and currently, BiLSTM and Multi-Layer Perceptron (MLP) choices are available as part of our toolkit. The inputs to adaptors are fused representations if multiple encoders are specified, else output from a single encoder. These adaptors serve as task-specific learnable parameters.

Multitasking: Multi-task learning can help models to pick relevant cues from one task to be applied to another. Such a setting was also previously investigated in the context of mixed texts, which showed promising improvements (Chandu et al., 2018). Furthermore, it is also shown in monolingual NLP that incorporating explicit semantics as an auxiliary task can enhance BERT’s performance (Zhang et al., 2020). Motivated by these, our toolkit offers support to conduct training of one or more tasks. Once a final representation is produced by adaptors of each task, we use a training criterion to compute loss and perform gradient backpropagation.

2.3 Tasks

Tasks: Our toolkit currently supports two kinds of tasks— sentence-level text classification and word-level sequence tagging, the flow for each is demonstrated in Figure 1. The *decoupled* design of our toolkit helps in seamlessly creating multi-task training setups. The kinds of tasks for which we offer support currently are listed in Table 1.

Adaptive Pretraining: Following the successes of task-adaptive and domain-adaptive pretraining in monolingual and multilingual NLP tasks (Gururangan et al., 2020), users of our toolkit can also perform such adaptive pretrainings using mixed texts on top of pretrained transformer-based models.

2.4 Codemixed Quantification

Our toolkit offers 5 standardized metrics for quantifying mixing in text, namely Code-Mixing Index (Gambäck and Das, 2014), Average switch-points (Khanuja et al., 2020), Multilingual Index, Probability of Switching and Language Entropy (Guzmán et al., 2017). We offer simple command line methods to compute these metrics and also offer metric-based data sampling.

2.5 Data Augmentation

Our toolkit also offers techniques to do data augmentation. While data augmentation is useful in cases where there is training data scarcity, for mixed datasets, it is also essential to produce a more generalized model. As part of this feature, this toolkit currently offers augmentation through transliteration, spelling variations and monolingual corpora. We currently support transliteration of Indic languages through an off-the-shelf tool- `indic-trans` (Bhat et al., 2015). Spelling variations include noising spelling, such as randomly removing/replacing vowel characters. Monolingual corpora augmentation is task specific. For a given task, such as sentiment classification, we augment publicly available monolingual corpora based on the task type from one or all of the mixing languages and use it while fine-tuning models.

2.6 Data Format

Due to diverse data formats of existing mixed datasets, benchmarking and comparing results across tasks is not readily feasible. To this end, we propose a standardized data format for syntactic, semantic level understanding and generation tasks, and our toolkit offers command line methods to adopt a user’s dataset to this standard format.

3 Experiments

Among many potential research applications of our toolkit, in this section, we demonstrate one – *benchmarking*. Table 1 presents performances of selected model architectures obtained using our toolkit on some popular mixed datasets. In Table 2, we also demonstrate the performances of different architectural choices implemented through our toolkit on two Hinglish datasets. For domain-adaptive pretraining of Hinglish datasets, we collate around 160K mixed sentences from several of the publicly available Hinglish datasets.

```
""" load toolkit """
from CodemixedNLP.cli import ModelCreator as HinglishToolkit
myToolkit = HinglishToolkit()

""" see what's available """
myToolkit.show_available_tasks()
# >>> ['sentiment', 'aggression', 'hatespeech', 'lid', 'pos', 'ner', 'mt']

""" download models """
myToolkit.download_finetuned("all")

""" load model for your application and use """
# Sentiment Classification
sentiment_classifier = myToolkit.from_pretrained("sentiment")
results = sentiment_classifier.get_predictions("wo team kaafi acha hei, not all days are yours")
# >>> results: positive

# Language identification
lid_classifier = myToolkit.from_pretrained("lid")
results = lid_classifier.get_predictions("Aur mere projects bhi bohot hain", prettify=True)
# >>> results: Aur/hi mere/hi projects/en bhi/hi bohot/hi hain/hi

# Machine Translation model
translator = myToolkit.from_pretrained("mt")
results = translator.get_predictions("anyways, wo team kaafi acha hei, glad to be a part of it")
# >>> results: anyways, that team is very good, glad to be a part of it
```

Figure 2: Command line interface for utilizing fine-tuned models. We provide several functionality compatible with the popular Huggingface and Fairseq libraries. Marked in boxes are customizable input arguments.

For task-adaptive pretraining, we just use the training and testing data available in the dataset of interest. For training, we use standard optimizers and model configurations.⁴

4 Demo

We fine-tune and publish transformer-based models for 7 tasks in Hinglish. We include 3 task types– (1) **Semantic** (Sentiment Classification, Hate Speech and Aggression Identification), (2) **Syntactic** (NER, POS and Language Identification), and (3) **Generation** (Hinglish→English Machine Translation). We present some examples of utilizing these models in Figure 2.

5 Conclusion

In this work, we presented a unified toolkit for modeling code-mixed texts. Additionally, the toolkit contains various functionalities such as data augmentation, code-mixing quantification, and ready-to-use fine-tuned models for 7 different NLP tasks in Hinglish. Our toolkit is simple enough for practitioners to integrate new features as well as develop wrappers around its existing functionalities. We believe this contribution facilitates a sustainable and extensible ecosystem of models by adding novel pretraining techniques tailored for mixed texts, text normalization techniques to counter spelling variations, error analysis tools to identify peculiarities in incorrect predictions and so on.

⁴Due to the space limitations, we direct the reader to check github.com/murali1996/CodemixedNLP for toolkit usage patterns and for the list of modeling choices.

References

- Gustavo Aguilar, Sudipta Kar, and Tamar Solorio. 2020. [LinCE: A centralized benchmark for linguistic code-switching evaluation](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1803–1813, Marseille, France. European Language Resources Association.
- Gustavo Aguilar and Tamar Solorio. 2020. [From English to code-switching: Transfer learning with strong morphological clues](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8033–8044, Online. Association for Computational Linguistics.
- Irshad Ahmad Bhat, Vandan Mujadia, Aniruddha Tammewar, Riyaz Ahmad Bhat, and Manish Shrivastava. 2015. [Iiit-h system submission for fire2014 shared task on transliterated search](#). In *Proceedings of the Forum for Information Retrieval Evaluation, FIRE '14*, pages 48–53, New York, NY, USA. ACM.
- Aditya Bohra, Deepanshu Vijay, Vinay Singh, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018. A dataset of hindi-english code-mixed social media text for hate speech detection. In *Proceedings of the second workshop on computational modeling of people’s opinions, personality, and emotions in social media*, pages 36–41.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#).
- Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne, Hiroshi Noji, Pierre Zweigenbaum, and Junichi Tsujii. 2020. [Characterbert: Reconciling elmo and bert for word-level open-vocabulary representations from characters](#).
- Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadharshini, and John Philip McCrae. 2020. [Corpus creation for sentiment analysis in code-mixed Tamil-English text](#). In *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL)*, pages 202–210, Marseille, France. European Language Resources association.
- Sharanya Chakravarthy, Anjana Umapathy, and Alan W Black. 2020. [Detecting entailment in code-mixed Hindi-English conversations](#). In *Proceedings of the Sixth Workshop on Noisy User-generated Text (W-NUT 2020)*, pages 165–170, Online. Association for Computational Linguistics.
- Khyathi Chandu, Thomas Manzini, Sumeet Singh, and Alan W. Black. 2018. [Language informed modeling of code-switched text](#). In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pages 92–97, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Flair. Flair toolkit. <https://github.com/flairNLP/flair>. Accessed: 2021-03-09.
- Björn Gambäck and Amitava Das. 2014. On measuring the complexity of code-mixing. In *Proceedings of the 11th International Conference on Natural Language Processing, Goa, India*, pages 1–7.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of ACL*.
- Gualberto A Guzmán, Joseph Ricard, Jacqueline Serigos, Barbara E Bullock, and Almeida Jacqueline Toribio. 2017. Metrics for modeling code-switching across corpora. In *INTERSPEECH*, pages 67–71.
- IndicNLP. Indic nlp library. https://anoopkunchukuttan.github.io/indic_nlp_library/. Accessed: 2021-03-09.
- iNLTK. Natural language toolkit for indic languages (inltk). <https://github.com/goru001/inltk>. Accessed: 2021-03-09.
- Sai Muralidhar Jayanthi and Akshat Gupta. 2021. [SJ_AJ@DravidianLangTech-EACL2021: Task-adaptive Pre-Training of Multilingual BERT models for Offensive Language Identification](#).
- Gagandeep Kaur, Abhishek Kaushik, and Shubham Sharma. 2019. Cooking is creating emotion: a study on hinglish sentiments of youtube cookery channels using semi-supervised approach. *Big Data and Cognitive Computing*, 3(3):37.
- Simran Khanuja, Sandipan Dandapat, Anirudh Srinivasan, Sunayana Sitaram, and Monojit Choudhury. 2020. [GLUECoS: An evaluation benchmark for code-switched NLP](#). In *Proceedings of the 58th Annual Meeting of the Association*

- for *Computational Linguistics*, pages 3575–3585, Online. Association for Computational Linguistics.
- Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. [Benchmarking aggression identification in social media](#). In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 1–11, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#).
- NLTK. Nltk sentiment analysis toolkit. <http://www.nltk.org/howto/sentiment.html>. Accessed: 2021-03-09.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*.
- Braja Gopal Patra, Dipankar Das, and Amitava Das. 2018. Sentiment analysis of code-mixed indian languages: An overview of sail_code-mixed shared task@ icon-2017. *arXiv preprint arXiv:1803.06745*.
- Parth Patwa, Gustavo Aguilar, Sudipta Kar, Suraj Pandey, Srinivas PYKL, Björn Gambäck, Tanmoy Chakraborty, Thamar Solorio, and Amitava Das. 2020. Semeval-2020 task 9: Overview of sentiment analysis of code-mixed tweets. In *Proceedings of the 14th International Workshop on Semantic Evaluation (SemEval-2020)*, Barcelona, Spain. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Keisuke Sakaguchi, Kevin Duh, Matt Post, and Benjamin Van Durme. 2017. [Robsut wrod reocginiton via semi-character recurrent neural network](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3281–3287. AAAI Press.
- Sunayana Sitaram, Khyathi Raghavi Chandu, Sai Krishna Rallabandi, and Alan W Black. 2020. [A survey of code-switched speech and language processing](#).
- Genta Indra Winata, Samuel Cahyawijaya, Zihan Liu, Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2021. [Are multilingual models effective in code-switching?](#)
- Genta Indra Winata, Zhaojiang Lin, Jamin Shin, Zihan Liu, and Pascale Fung. 2019. Hierarchical meta-embeddings for code-switching named entity recognition. *arXiv preprint arXiv:1909.08504*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020. [Semantics-aware bert for language understanding](#).