

Multilingual BERT has an accent: Evaluating English influences on fluency in multilingual models

Isabel Papadimitriou* and Kezia Lopez* and Dan Jurafsky

Computer Science Department

Stanford University

{isabelvp,keziak1,jurafsky}@stanford.edu

Abstract

While multilingual language models can improve NLP performance on low-resource languages by leveraging higher-resource languages, they also reduce average performance on all languages (the ‘curse of multilinguality’). Here we show another problem with multilingual models: grammatical structures in higher-resource languages bleed into lower-resource languages, a phenomenon we call *grammatical structure bias*. We show this bias via a novel method for comparing the fluency of multilingual models to the fluency of monolingual Spanish and Greek models: testing their preference for two carefully-chosen variable grammatical structures (optional pronoun-drop in Spanish and optional Subject-Verb ordering in Greek). We find that multilingual BERT is biased toward the English-like setting (explicit pronouns and Subject-Verb-Object ordering) as compared to our monolingual control language model. With our case studies, we hope to bring to light the fine-grained ways in which multilingual models can be biased, and encourage more linguistically-aware fluency evaluation.

1 Introduction

Multilingual language models share a single set of parameters between many languages, opening new pathways for multilingual and low-resource NLP. However, not all training languages have an equal amount, or a comparable quality of training data in these models. In this paper, we investigate if the hegemonic status of English influences other languages in multilingual language models. We propose a novel method for evaluation, whereby we ask if model predictions for lower-resource languages exhibit structural features of English. This is similar to asking if the model has learned some languages with an “English accent”, or an English *grammatical structure bias*.

We demonstrate this bias effect in Spanish and Greek, comparing the monolingual models BETO




Figure 1: Our method for evaluating English structural bias in multilingual models. We compare monolingual and multilingual model predictions on two sets of natural sentences in the target language: one which is structurally parallel to English, and one which is not.

(Cañete et al., 2020) and GreekBERT (Koutsikakis et al., 2020) to multilingual BERT (mBERT), where English is the most frequent language in the training data. We show that *mBERT prefers English-like sentence structure in Spanish and Greek* compared to the monolingual models. Our case studies focus on Spanish pronoun drop (pro-drop) and Greek subject-verb order, two structural grammatical features. We show that multilingual BERT is structurally biased towards explicit pronouns rather than pro-drop in Spanish, and subject-before-verb order in Greek: the structural forms parallel to English.

Though the effect we showcase here is likely not captured by the downstream classification tasks often used to evaluate multilingual models (Hu et al., 2020), it demonstrates the type of fluency that can be lost with multilingual training — something that current evaluation methods miss. In fact, though we choose two clear-cut syntactic features to investigate, there are many less-measurable features that make language production fluent: subtleties in lexical choice, grammatical choice, and discourse expression, among many others. With this paper, beyond showing a trend for two specific grammatical features, we wish to highlight fluency discrepancies in multilingual models, and also call for more evaluations focused on fluency.

S_{parallel} : English-like structure	$S_{\text{different}}$: Different structure
Spanish explicit pronoun (pron in orange, verb in blue)	Spanish prodrop (verb in blue)
Yo <u>volveré</u> para averiguarlo I will return to figure it out El 2004, <u>ella</u> <u>hizo</u> doblaje a el Inglés [...] In 2004, <u>she</u> <u>did</u> dubbing to English [...] <u>Ella</u> <u>decide</u> pasar sus vacaciones en la granja <u>She</u> <u>decides</u> to spend her vacation in the country	Jamás <u>dan</u> soluciones y siempre [...] [They] Never <u>give</u> solutions and always [...] <u>Jugó</u> de centrocampista en el Real Zaragoza [He/She/You] <u>Played</u> as a midfielder in Real Zaragoza <u>Habita</u> en Perú. [He/She/You] <u>Lives</u> in Peru
Greek Subject-Verb (subject in orange, verb in blue)	Greek Verb-Subject (subject in orange, verb in blue)
<u>Πηγές</u> της Αντιπολίτευσης <u>αναφέρουν</u> ότι [...] Sources of the Opposition <u>mention</u> that [...] Σε άλλες πλευρές ο <u>ποταμός</u> <u>κυλά</u> από ψηλούς βράχους On other sides, the <u>river</u> <u>flows</u> from tall boulders Η <u>εκπαίδευση</u> και η μόρφωση <u>απέκτησαν</u> επιτέλους προτεραιότητα <u>Training</u> and education have finally <u>acquired</u> priority	Το σκορ του αγώνα <u>άνοιξε</u> ο <u>Γουέν</u> Ρούνι The score of the game <u>opened</u> <u>Wayne</u> Rooney Εδώ πρέπει να <u>γίνουν</u> μεγαλύτερες <u>προσπάθειες</u> . Here must <u>happen</u> bigger <u>efforts</u> Απασχόληση στο εξωτερικό <u>ψάχνουν</u> οι μισοί <u>Έλληνες</u> σε παραγωγική ηλικία Employment in foreign countries <u>search</u> half of <u>Greeks</u>

Table 1: Examples from our dataset for S_{parallel} and $S_{\text{different}}$ in Spanish and Greek, along with roughly word-by-word gloss translations in English. In all cases, we’ve underlined $w(x)$, the word we use to represent the construction in our calculations. These examples are not randomly selected and have been chosen to be significantly shorter than the average sentence in our datasets in order to be presentable in a table.

Our proposed method can be expanded, without the need for manual data collection, to any language with a syntactic treebank and a monolingual model. Since our method focuses on fine-grained linguistic features, some expert knowledge of the target language is necessary for evaluation. Multilingual evaluation so far has been largely translated or automatically curated, and the methods for creating such datasets have allowed for the creation of resources in many languages for which there were none. Fluency evaluation requires some linguistic expertise to set up, and as such is more restricted in the languages the research community can reach. Nevertheless, such evaluation has been missing from the multilingual NLP literature, and our work bridges this gap by proposing fluency testing for multilingual models.

Our work builds off of a long literature on multilingual evaluation which has until now mostly focused on downstream classification tasks (Conneau et al., 2018; Ebrahimi et al., 2022; Clark et al., 2020; Liang et al., 2020; Hu et al., 2020; Raganato et al., 2020; Li et al., 2021). With the help of these evaluation methods, research has pointed out the problems for both high- and low-resource languages that come with adding many languages to a single model (Wang et al., 2020; Turc et al., 2021;

Lauscher et al., 2020, inter alia). Methods for creating more equitable models have been proposed, through identifying or reserving language-specific parameters for each language (Ansell et al., 2022; Pfeiffer et al., 2022), through training models without typologically distant languages that dominate the training data (Ogueji et al., 2021; Virtanen et al., 2019; Ògúnrè mí and Manning, 2023), as well as through adding model capacity (Conneau et al., 2020; Xue et al., 2021; Lepikhin et al., 2021; Liang et al., 2023). We hope that our work can add to these analyses and methodologies by pointing out issues beyond downstream classification performance that can arise with multilingual training, and aid towards building and evaluating more equitable multilingual models.

2 Method

Our method relies on finding a variable construction in the target language which can take two structural surface forms: one which is parallel to English (S_{parallel}) and one which is not ($S_{\text{different}}$). Surface forms parallel to English are those which mirror English structure. For example, English has strict Subject-Verb-Object word order, and so a *parallel* structure in another language is one where the verb

and its arguments appear in Subject-Verb-Object order, while a *different* structure is one where the verb appears before the subject (see Table 1 for examples).

Once we have identified such a construction in our target language, we can ask: are multilingual models biased towards S_{parallel} ? For a native speaker of the target language, structural, semantic, and discourse features determine whether they will use S_{parallel} or $S_{\text{different}}$ in a given context — with the alternative option usually being less fluent. We assume that a BERT-sized monolingual model in the target language will have a sufficiently accurate representation of this fluent variation between S_{parallel} and $S_{\text{different}}$ without being influenced by other languages. Therefore, to understand if multilingual models have an English structural bias, we now just have to answer: do multilingual models prefer S_{parallel} over $S_{\text{different}}$ *more* than the fluent distribution defined by a monolingual model?

2.1 Collecting model judgements

By design, both S_{parallel} and $S_{\text{different}}$ are constructions that occur naturally in the target language. Therefore, we should be able to use the syntactic treebank annotations to pick out sentences that exhibit the structures S_{parallel} or $S_{\text{different}}$. We can put these extracted sentences into two corpora, C_{parallel} and $C_{\text{different}}$. Note that the sentences in C_{parallel} and $C_{\text{different}}$ are unrelated and not paired, and that the two corpora can have different sizes. Crucially, we have to use natural sentences for both of our corpora: we cannot artificially alter sentences from S_{parallel} to $S_{\text{different}}$, or use templates to create sentences. This is because our evaluation is about the subtleties of fluency, while altered or templated stimuli are not naturally produced and are therefore often awkward, confounding any effect we might want to measure.

Each model gives us a ratio r_{model} : the average probability of a sentence in C_{parallel} divided by the average probability of a sentence in $C_{\text{different}}$ according to the model. That is:

$$r_{\text{model}} = \frac{\sum_{x \in C_p} P_{\text{model}}(x) / |C_p|}{\sum_{x \in C_d} P_{\text{model}}(x) / |C_d|} \quad (1)$$

We want to compare judgements on these corpora from two models: a monolingual model `mono` and a multilingual model `multi`. Our experimental question then boils down to asking if r_{multi} is significantly larger than r_{mono} .

2.2 From model outputs to construction probability

How can we calculate $P_{\text{model}}(x)$ for a given sentence x , focusing on the probability of a specific construction in x ? Looking at model judgements over long natural sentences introduces a lot of noise that is unrelated to the structural construction in question, reducing the statistical power of our experiment. Furthermore, since we are looking at encoder-only bidirectional models, there is no canonical or controlled way of extracting the probability of a whole sentence. To get a better model judgement for each sentence, we can extract the probability of *one word* in each sentence that best represents the construction. For example, if we are looking at pronoun drop, it makes sense to use main verb of the sentence as the target word, as this is the syntactic head of the pronoun that is present or dropped. Using a carefully chosen word as a proxy for the probability of a construction is a methodological choice also made in reading time psycholinguistics experiments (Levy, 2011; Levy and Keller, 2013).

Going back to our problem of calculating $P_{\text{model}}(x)$, we define w to be a function that returns the structurally-relevant word from each sentence. Using this, we approximate $P_{\text{model}}(x)$ in Eq. (1) with $P_{\text{model}}(w(x)|x)$. The probability $P(w(x)|x)$ is simple to calculate for BERT-style masked language models: it is simply the logit of the word $w(x)$ when we encode the sentence x using `model`.

2.3 Extending to more languages

Extending our fluency evaluation to a new language requires three language-specific steps: (1) decide on an appropriate construction with two structural forms S_{parallel} and $S_{\text{different}}$, (2) decide on an appropriate $w(x)$: which word in each structural form can represent the form, and (3) use treebank annotations to pull out sentences which exhibit S_{parallel} or $S_{\text{different}}$, and identify the relevant word. Below, we detail these steps for our two case studies.

2.4 Case Study: Spanish Pro-drop

In Spanish, the subject pronoun is often dropped: person and number are mostly reflected in verb conjugation, so the pronoun is realized or dropped depending on semantic and discourse factors. English, on the other hand, does not allow null subjects except in rare cases, and expletive syntactic subjects like “there” are even added when there




Figure 2: Results from our experiment on the Spanish GSD treebank, along with two examples from the treebank to illustrate S_{parallel} (with pronoun) and $S_{\text{different}}$ (pro-drop). We compare model logits for the main verb of the sentence, which is bold and highlighted in the examples. Error bars represent 95% bootstrap confidence intervals. We find that r_{mono} is significantly smaller than r_{multi} (bootstrap sampling, $p < 0.05$).

is no clear subject. For our Spanish experiment, we define S_{parallel} to be sentences which have the subject pronoun of the main verb, as is necessary in English, and $S_{\text{different}}$ to be pro-drop sentences which have a main verb with no realized subject. We define w to be the main verb of the sentence, which is always present in our extracted examples.

To extract our corpora C_{parallel} and $C_{\text{different}}$, we use the Spanish GSD treebank from the Universal Dependencies dataset (De Marneffe et al., 2021). We ignore all sentences not verb-rooted (i.e. noun phrases), those rooted with “haber” (which in its copula-like existential form cannot take an explicit subject, “There is” in English), and those using the impersonal-“se” passive construction (e.g. “se nos fue permitido”, “it was permitted of us”). We then take all sentences with a pronoun subject (i.e. a pronoun dependent of the root verb) and add them to C_{parallel} and all sentences where there is no nsubj relation to root verb and add them to $C_{\text{different}}$. We always pick the main root verb of the sentence as our w . We collect 283 sentences in C_{parallel} and 2,656 sentences in $C_{\text{different}}$.

2.5 Case Study: Greek Subject-Verb order

English is a fixed word order language: with few exceptions, the order of a verb and its arguments is Subject-Verb-Object. Greek, on the other hand, has mostly free word order (Mackridge, 1985), meaning that the verb and arguments can appear in any order that is most appropriate given discourse context. For our experiment, we define S_{parallel} to be cases in Greek when the subject precedes the verb, as is the rule in English. $S_{\text{different}}$ is then the cases when the verb precedes the subject, which almost never happens in English.

We define w to be the first element of the subject and verb: the subject when the subject comes first or the verb when the verb comes first. This first element is closer to the surrounding context, and so gives us a word-order-sensitive measurement of how the subject-verb construction is processed as a whole within the context. Though this choice means that our w is a noun in S_{parallel} and a verb in $S_{\text{different}}$, this does not constitute a confounder between models: we are comparing the same noun-verb probability ratio between different models.

To extract our corpora C_{parallel} and $C_{\text{different}}$, we use the Greek Dependency Treebank, the Universal Dependencies treebank for Greek (Prokopidis and Papageorgiou, 2017). We take all sentences where the main verb has a lexical subject, and we add to C_{parallel} if the subject appears before the verb and to $C_{\text{different}}$ if it appears after. We collect 1,446 sentences in C_{parallel} and 425 sentences in $C_{\text{different}}$.

3 Results

Results are shown in Figures 2 and 3, showing for both of our case studies that multilingual BERT has a greater propensity for preferring English-like sentences which exhibit S_{parallel} . Multilingual BERT significantly prefers pronoun sentences over pro-drop compared with monolingual BETO (bootstrap sampling, $p < 0.05$), and significantly prefers subject-verb sentences over verb-subject sentences over GreekBERT (bootstrap sampling, $p < 0.05$).

4 Discussion

In this paper, we proposed fluency evaluation as a further way of understanding the curse of multilinguality: what can be lost when we train many languages together. The discrepancies that we point out in these experiments are not going to seriously affect multilingual LM performance, especially in the more coarse-grained classification tasks that




Figure 3: Results from our experiment on the Greek Dependency Treebank, along with two examples from the treebank to illustrate S_{parallel} (Subject-Verb) and $S_{\text{different}}$ (Verb-Subject). We measure and compare model logits for the bold words: the subject in subject-verb sentences and the verb in verb-subject sentences. Error bars represent 95% bootstrap confidence intervals. r_{mono} is significantly smaller than r_{multi} (bootstrap sampling, $p < 0.05$).

are most commonly used for evaluation. But, as we demonstrate here, not all levels of language learning can be evaluated from such datasets.

Our experiments do not pinpoint the reasons behind the effects that we measure: there are different possible explanations for the English-like trends that we showcase. On the one hand, the effects we measure might stem from training with a language that’s more dominant in the training data, like English is for many multilingual models. Such training could lead to an English-biased representation space which the representations of other languages conform to. On the other hand, the effects we show might be down to the data: the non-English datasets used to train a multilingual model may be more limited in domain, may contain a high proportion of data that’s actually been translated from English (Multilingual Wikipedia is often translated, Adar et al., 2009), or might be more polluted with irrelevant or non-linguistic elements. Domain limitations and translationese stemming from the data are separate, but related issues to fluency: fluency can be grammatical, but

also involves proficiency in a range of registers or possibilities. It is also possible that the effects we show are due to a combination of both multilingual representation learning artifacts, and training data quality. Further controlled fluency experimentation on the limits and abilities of multilingual models is needed to disentangle these effects. We hope the case studies in this paper can inspire more fine-grained evaluation of multilingual models, so that we understand the “accent”-like effects of hegemonic languages more fully.

5 Limitations

This study is meant to highlight the kinds of modeling flaws that have so far gone undetected and that can arise for lower-resource languages in multilingual models. However, our study does not focus on languages that are truly low-resource. In fact, as designed it could not do so: our methodology relies on having an available monolingual model, which of course requires a large amount of training data. This is because our method requires a control: we can only judge multilingual models against what we can believe to be a non-biased language model in the language. There are ways to test for fluency in low-resource languages that would not require a monolingual model as a control, but would require dataset collection in the target language for features that reflect fluency and linguistic acceptability (similar to what Warstadt et al. (2019) achieve with the CoLA dataset for English). We hope our study can create motivation for such work in linguistically-aware, fine-grained multilingual evaluation for languages of all resource levels.

Our experiments focus on BERT-style models, since this is mostly the size of model available for monolingual, non-English models (in our case BETO and GreekBERT). However, it is not necessary from these experiments that our findings extrapolate to larger models that are commonplace at the time of writing.

Lastly, both pro-drop and subject-verb order are largely discourse-dependent constructions. For example, pro-drop is more likely when the subject of the sentence is very clear from the discourse, while subject-verb order in Greek is changed to achieve different discourse focus, similar to how intonation changes the focus of a sentence in English (e.g., stressing the verb in “Mary helped John” puts the focus on the verb, which in Greek can be done by putting the verb first). Despite this, all of our ex-

periments are done on isolated sentences from the UD treebanks and do not contain discourse content. Though this means that the models do not have the full relevant context for each input, we do not expect that having more context should favor one model more than another for our evaluation. Since this work compares models on the same inputs, we did not consider this a significant confounder.

6 Acknowledgements

We thank Anjalie Field and Mirac Suzgun for comments on drafts. This research was funded in part by NSF award number IIS-2128145.

References

- Eytan Adar, Michael Skinner, and Daniel S. Weld. 2009. [Information arbitrage across multi-lingual wikipedia](#). In *Proceedings of the Second ACM International Conference on Web Search and Data Mining, WSDM '09*, page 94–103, New York, NY, USA. Association for Computing Machinery.
- Alan Ansell, Edoardo Ponti, Anna Korhonen, and Ivan Vulić. 2022. Composable sparse fine-tuning for cross-lingual transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1778–1796.
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Joun-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish Pre-trained BERT Model and Evaluation Data. In *PMLADC at ICLR 2020*.
- Jonathan H Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard 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.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating Cross-lingual Sentence Representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485.
- Marie-Catherine De Marneffe, Christopher D Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal Dependencies. *Computational linguistics*, 47(2):255–308.
- Abteem Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios Gonzales, Ivan Meza-Ruiz, et al. 2022. AmericasNLI: Evaluating Zero-shot Natural Language Understanding of Pre-trained Multilingual Models in Truly Low-resource Languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6279–6299.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- John Koutsikakis, Ilias Chalkidis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2020. Greek-BERT: The Greeks visiting sesame street. In *11th Hellenic Conference on Artificial Intelligence*, pages 110–117.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. From Zero to Hero: On the limitations of zero-shot language transfer with multilingual transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. [{GS}hard: Scaling giant models with conditional computation and automatic sharding](#). In *International Conference on Learning Representations*.
- Roger Levy. 2011. Integrating surprisal and uncertainty in online sentence comprehension: Formal techniques and empirical results. In *Proceedings of the 49th annual meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1055–1065.
- Roger P. Levy and Frank Keller. 2013. [Expectation and locality effects in German verb-final structures](#). *Journal of Memory and Language*, 68(2):199–222.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. 2021. Mtop: A comprehensive multilingual task-oriented semantic parsing benchmark. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2950–2962.
- Davis Liang, Hila Gonen, Yuning Mao, Rui Hou, Naman Goyal, Marjan Ghazvininejad, Luke Zettlemoyer, and Madian Khabsa. 2023. Xlm-v: Overcoming the vocabulary bottleneck in multilingual masked language models. *arXiv preprint arXiv:2301.10472*.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroou

- Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- P. Mackridge. 1985. *The Modern Greek Language: A Descriptive Analysis of Standard Modern Greek*. Oxford University Press.
- Kelechi Ogueji, Yuxin Zhu, and Jimmy Lin. 2021. [Small Data? No Problem! Exploring the viability of pretrained multilingual language models for low-resourced languages](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 116–126, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tolúloṗé Ògúnṛẹ̀mí and Christopher D. Manning. 2023. [Mini but Mighty: Efficient multilingual pretraining with linguistically-informed data selection](#). In *Findings of EACL 2023*.
- Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022. [Lifting the curse of multilinguality by pre-training modular transformers](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3479–3495, Seattle, United States. Association for Computational Linguistics.
- Prokopis Prokopidis and Haris Papageorgiou. 2017. [Universal Dependencies for Greek](#). In *Proceedings of the NoDaLiDa 2017 Workshop on Universal Dependencies (UDW 2017)*, pages 102–106, Gothenburg, Sweden. Association for Computational Linguistics.
- Alessandro Raganato, Tommaso Pasini, Jose Camacho-Collados, and Mohammad Taher Pilehvar. 2020. [XLWiC: A multilingual benchmark for evaluating semantic contextualization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. The Association for Computational Linguistics.
- Iulia Turc, Kenton Lee, Jacob Eisenstein, Ming-Wei Chang, and Kristina Toutanova. 2021. [Revisiting the Primacy of English in Zero-shot Cross-lingual Transfer](#). *CoRR*, abs/2106.16171.
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. [Multilingual is not enough: BERT for Finnish](#). *arXiv preprint arXiv:1912.07076*.
- Zirui Wang, Zachary C Lipton, and Yulia Tsvetkov. 2020. [On negative interference in multilingual models: Findings and a meta-learning treatment](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4438–4450.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *NAACL-HLT*.