

Please cite the Published Version

Štajner, Sanja, Ferrés, Daniel, Shardlow, Matthew , North, Kai, Zampieri, Marcos and Saggion, Horacio (2022) Lexical simplification benchmarks for English, Portuguese, and Spanish. Frontiers in Artificial Intelligence, 5. p. 991242. ISSN 2624-8212

DOI: https://doi.org/10.3389/frai.2022.991242

Publisher: Frontiers Media S.A.

Version: Published Version

Downloaded from: https://e-space.mmu.ac.uk/631705/

Usage rights: (cc) BY

Creative Commons: Attribution 4.0

Additional Information: This is an Open Access article which appeared in Frontiers in Artificial Intelligence, published by Frontiers Media

Data Access Statement: The original contributions presented in the study are publicly available. This data can be found here: https://github.com/LaSTUS-TALN-UPF/TSAR-2022-Shared-Task.

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines)

Check for updates

OPEN ACCESS

EDITED BY Grigori Sidorov, Instituto Politécnico Nacional (IPN), Mexico

REVIEWED BY Noman Ashraf, Dana-Farber Cancer Institute, United States Iqra Ameer, Instituto Politécnico Nacional (IPN), Mexico Maaz Amjad, Instituto Politécnico Nacional (IPN), Mexico

*CORRESPONDENCE Matthew Shardlow M.Shardlow@mmu.ac.uk

SPECIALTY SECTION This article was submitted to

Natural Language Processing, a section of the journal Frontiers in Artificial Intelligence

RECEIVED 11 July 2022 ACCEPTED 19 August 2022 PUBLISHED 22 September 2022

CITATION

Štajner S, Ferrés D, Shardlow M, North K, Zampieri M and Saggion H (2022) Lexical simplification benchmarks for English, Portuguese, and Spanish. *Front. Artif. Intell.* 5:991242. doi: 10.3389/frai.2022.991242

COPYRIGHT

© 2022 Štajner, Ferrés, Shardlow, North, Zampieri and Saggion. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Lexical simplification benchmarks for English, Portuguese, and Spanish

Sanja Štajner¹, Daniel Ferrés², Matthew Shardlow^{3*}, Kai North⁴, Marcos Zampieri⁴ and Horacio Saggion²

¹Other, Karlsruhe, Germany, ²LaSTUS Lab, TALN Group, Universitat Pompeu Fabra, Barcelona, Spain, ³Department of Computing and Mathematics, Manchester Metropolitan University, Manchester, United Kingdom, ⁴School of Computing, George Mason University, Fairfax, VA, United States

Even in highly-developed countries, as many as 15-30% of the population can only understand texts written using a basic vocabulary. Their understanding of everyday texts is limited, which prevents them from taking an active role in society and making informed decisions regarding healthcare, legal representation, or democratic choice. Lexical simplification is a natural language processing task that aims to make text understandable to everyone by replacing complex vocabulary and expressions with simpler ones, while preserving the original meaning. It has attracted considerable attention in the last 20 years, and fully automatic lexical simplification systems have been proposed for various languages. The main obstacle for the progress of the field is the absence of high-quality datasets for building and evaluating lexical simplification systems. In this study, we present a new benchmark dataset for lexical simplification in English, Spanish, and (Brazilian) Portuguese, and provide details about data selection and annotation procedures, to enable compilation of comparable datasets in other languages and domains. As the first multilingual lexical simplification dataset, where instances in all three languages were selected and annotated using comparable procedures, this is the first dataset that offers a direct comparison of lexical simplification systems for three languages. To showcase the usability of the dataset, we adapt two state-of-the-art lexical simplification systems with differing architectures (neural vs. non-neural) to all three languages (English, Spanish, and Brazilian Portuguese) and evaluate their performances on our new dataset. For a fairer comparison, we use several evaluation measures which capture varied aspects of the systems' efficacy, and discuss their strengths and weaknesses. We find that a state-of-the-art neural lexical simplification system outperforms a state-of-the-art non-neural lexical simplification system in all three languages, according to all evaluation measures. More importantly, we find that the state-of-the-art neural lexical simplification systems perform significantly better for English than for Spanish and Portuguese, thus posing a question if such an architecture can be used for successful lexical simplification in other languages, especially the low-resourced ones.

KEYWORDS

natural language processing, lexical simplification, benchmark datasets, evaluation methodologies, low-resource tasks, artificial intelligence for social good

1. Introduction

According to the adult literacy report conducted in 24 highly-developed countries (OECD, 2013), 16.7% of a population, on average, cannot understand texts that go beyond a basic vocabulary. This percentage is even higher for some countries, e.g., 21.7% for the U.S., and 28.3% for Spain (OECD, 2013). People who do not correctly understand written information cannot make informed decisions regarding critical processes such as healthcare choices, legal representation, education, or democratic rights. This prevents them from taking an active role in society.

The disparity between the typical level of vocabulary in written communications and the audiences they were intended for was already evident almost a century ago (Ogden, 1937). Since then, various campaigns advocated for producing easy-to-read texts that would be understood by more people. Similarly, in scenarios where critical information needs to be unambiguously communicated, such as technical manual writing or in disaster relief efforts, the use of controlled languages has been suggested (Temnikova et al., 2015).

Although the efforts to promote the use of plain English have been effective in communicating the need for easily understandable text, there has been little consensus on what standards should be adopted and a low uptake of existing standards. Basic English (Ogden, 1937) was suggested for international communication as a means of producing more accessible texts by limiting the vocabulary and the variety of syntactic structures used. Fifty years later, the Plain English Campaign (Crystal, 1987) offered a Crystal Mark scheme which entails manually checking documents for simplified English. The campaign was followed by releases of several guidelines for producing easy-to-read English texts for people with intellectual disabilities (Freyhoff et al., 1998; Mencap, 2002), making public information more accessible by using plain English (PlainLanguage, 2011), and making web content more accessible (W3C, 2008; Cooper et al., 2010). Most guidelines were initially proposed for English, and later adapted to other languages. For example, Rational French (Barthe et al., 1999) was inspired by AECMA Simplified English, a controlled language used in the aerospace industry (Štajner, 2021).

Presently, easy-to-read news articles are offered on specialized websites in many countries, e.g., Noticias fácil in Spain¹, DR in Denmark², News Web Easy in Japan³. Several websites, e.g., Newsela⁴, and News in Levels⁵ offer news articles at several complexity levels and are designed for English

3 https://www3.nhk.or.jp/news/easy/

language learners. Those websites, however, offer only a few current articles at any one time as manual adaptation cannot keep up with new information being released on the web.

1.1. Lexical simplification

Lexical simplification is a natural language processing (NLP) task of automatically simplifying words and phrases in a given text or sentence, so that the information is more understandable to the reader. At the same time, lexical simplification needs to preserve the original meaning of the text or sentence. Depending on the level of simplification applied, some non-essential nuances of the original meaning are allowed to be lost. Nevertheless, the core information needs to be preserved. Due to a high potential of being used for social good and improving social inclusion of many people, lexical simplification has been attracting growing attention from the NLP community (Saggion, 2017; Štajner, 2021).

The main reason for focusing on automated tools for lexical simplification is that manual transformation of complex vocabulary into basic vocabulary in a given text is expensive, time-consuming, and requires professional editors. Furthermore, it has been shown that different target groups, e.g., native and non-native speakers, need different words to be simplified (Yimam et al., 2017). In the case of people with cognitive or reading impairments, needs for vocabulary simplification are even more heterogeneous (Orăsan et al., 2018). In the case of language learners, which words need to be simplified depends not only on their language proficiency level, but also on their native language (Aprosio et al., 2018). Lexical simplification should thus ideally be personalized. Wellperforming lexical simplification systems would significantly lower the editing costs and enable publishing more up-to-date articles understandable to wider populations, and offer a possibility for developing personalized readers.

Although it has been attracting the attention of the NLP community since the late 1990s, lexical simplification has only recently gained popularity as the underlying technologies have advanced (Štajner, 2021). The majority of proposed systems focus on the English language (Paetzold and Specia, 2017b; Alva-Manchego et al., 2020), as it is the case with many tasks in the NLP world. Several works have proposed lexical simplification systems for Spanish (Bott et al., 2012; Baeza-Yates et al., 2015; Ferrés et al., 2017b), Portuguese (Hartmann et al., 2020), French (Hmida et al., 2018), Chinese (Qiang et al., 2021), Japanese (Kajiwara and Yamamoto, 2015; Hading et al., 2016), and Swedish (Abrahamsson et al., 2014).

Lexical simplification consists of four subtasks that can be modeled either separately (in a modular approach) or jointly (in an end-to-end approach): (1) complex word/phrase identification (CWI or CPI); (2) generation of possible substitutes (SG); (3) selection of substitutes that fit the context

¹ https://www.discapnet.es/lectura-facil/noticias-facil

² https://www.dr.dk/ligetil/

⁴ www.newsela.com

⁵ https://www.newsinlevels.com/

and preserve the original meaning (SS); and (4) ranking of substitutes (SR).

The goal of the first subtask is to ensure that only those words and phrases that are difficult for the target reader(s) are simplified. Some works opt for not having this subtask and instead treat all content words as potentially difficult words. In such lexical simplification systems, the other subtasks treat the original/target word as one of the candidate substitutions. Having a Complex Word Identification (CWI) module at the beginning of the lexical simplification pipeline has been shown to improve the performance of lexical simplification systems by avoiding unnecessary errors stemming from trying to simplify words that do not need to be simplified (Paetzold and Specia, 2015). In the NLP community, complex word identification is treated as a separate NLP task and has attracted a lot of attention through two shared tasks: SemEval 2016 CWI for English (Paetzold and Specia, 2016b), and the BEA 2018 CWI shared task for English, German and Spanish, as well as multilingual CWI including English, German, Spanish, and French (Yimam et al., 2018). The SemEval 2021 shared task on lexical complexity prediction (Shardlow et al., 2021) also provided a new dataset for complex word and multi-word expressions identification for English. In the second sub-task (SG), candidate substitutions are usually retrieved either from specialized dictionaries and thesaurus, or by leveraging statistical properties of large corpora. The approach chosen for this sub-task influences coverage of lexical simplification systems. The third sub-task (SS) is crucial for ensuring that the original meaning has been preserved during lexical simplification, as it checks whether or not the substitutes fit the context well and convey the same semantics as the original word. This sub-task also ensures grammaticality of the output. The method used for ranking the simplicity of the substitution candidate, in the fourth sub-task (SR), should be chosen depending on specific lexical simplification needs of the target population or target user. In an end-to-end approach to lexical simplification, this sub-task should ideally be modeled separately from the rest, to allow for adaptation of the system to particular needs of different users or target populations.

As there are already many high-quality datasets for the evaluation of complex word/phrase identification modules, in this study, we focus on datasets necessary for evaluating the three other aspects of the lexical simplification pipeline (generation of substitution candidates, selection of substitutes that fit the context and preserve the original meaning, and ranking of substitutes according to their simplicity).

1.2. Evaluation datasets for lexical simplification

The main bottleneck for building reliable and efficient lexical simplification systems is the absence of datasets for training

and evaluation. The absence of datasets for training has been mitigated by using unsupervised methods. Since 2015, lexical simplification systems that use word embedding vectors (Glavaš and Štajner, 2015; Paetzold and Specia, 2016c) and neural language models (Shardlow and Nawaz, 2019; Qiang et al., 2020a) dominate the field. The absence of reliable datasets for automatic evaluation is, nevertheless, a great issue, as evaluation with target users requires significant time and specifically trained human assessors, and is thus not optimal for prototyping lexical simplification systems (Štajner, 2021).

Evaluation datasets for lexical simplification systems only exist for a handful of languages: English (Specia et al., 2012; Horn et al., 2014; Paetzold and Specia, 2016a), French (Rolin et al., 2021), Portuguese (Hartmann and Aluísio, 2020), Spanish (Alarcón, 2021), Japanese (Kajiwara and Yamamoto, 2015), and Chinese (Qiang et al., 2021). The SemEval-2012 dataset (Specia et al., 2012) only evaluates one aspect of English lexical simplification systems, the ranking of the substitution candidates. The other datasets were used to evaluate all aspects of the lexical simplification pipeline: generation of the substitution candidates, their ranking, and fitting to the given context. Of the above-mentioned evaluation datasets, only the English ones have the status of the benchmark datasets. Evaluation datasets for other languages contain a small number of annotations by only one-to-five people. They were thus used only as a proxy for evaluating proposed lexical simplification systems in corresponding languages (see Section 2 for more details).

Another issue with the existing evaluation datasets is that they are not comparable across languages due to different procedures used to select and annotate the instances. Therefore, they cannot be used for the evaluation of multilingual lexical simplification systems and better understanding of strengths and weaknesses of different approaches to lexical simplification should they be adapted to other languages. To fill this gap, in this study, we present a multilingual benchmark dataset for evaluation of lexical simplification systems that consists of instances in English, Spanish, and Brazilian Portuguese, all selected and annotated in a comparable manner using best practices.

1.3. Contributions

The work presented in this study makes four contributions to the field of lexical simplification:

- Compilation of three comparable benchmark datasets for evaluation of lexical simplification systems for English, Spanish, and Portuguese;
- Comparison of several evaluation methods for assessing efficacy of automatic lexical simplification systems and discussion of their strengths and weaknesses;

- 3. Comparison of the performances of state-of-the-art lexical simplification systems for English, Spanish, and Portuguese on the new benchmark datasets;
- 4. Detailed description of the process of compiling the benchmark datasets to offer a possibility for compiling comparable datasets in other languages.

2. Related work

In this section, we give an overview of existing lexical simplification systems for English, Spanish, and Portuguese (Section 2.1), evaluation datasets for lexical simplification and their limitations (Section 2.2), and commonly used evaluation metrics for measuring performance of lexical simplification systems (Section 2.3). These sections lay out the reasons for the choices made in this work, i.e., the choice of lexical simplification systems used in experiments, design of the annotation experiments, and the choice of evaluation metrics used to showcase the usefulness of the newly compiled dataset.

2.1. State-of-the-art lexical simplification systems for English, Spanish, and Portuguese

Since 2015 and up to 2020, neural LS systems which leverage word embeddings for the retrieval of the substitution candidates and their ranking (Glavaš and Štajner, 2015; Paetzold and Specia, 2016c) were considered the state-of-the-art for LS in English. Due to their unsupervised nature and the use of word embeddings trained on vast amounts of data, they have significantly better coverage and adaptability than the previously proposed non-neural systems (Biran et al., 2011; Horn et al., 2014). The systems proposed by Glavaš and Štajner (2015) and Paetzold and Specia (2016c) show similar performances on several benchmarks (Paetzold and Specia, 2017b), the former being computationally lighter as it uses 200-dimensional pretrained word embeddings, while the latter requires training of 1300-dimensional word embeddings. Two other neural systems (Paetzold and Specia, 2017a; Gooding and Kochmar, 2019) also rank highly on the common LS benchmarks (Qiang et al., 2020b), but they are supervised. This lowers their potential to be adapted to other languages which do not offer much training data.

Currently, the best performing LS system for English is the LSBert system (Qiang et al., 2020a), which uses pre-trained transformer language model BERT (Devlin et al., 2019) and a masking technique for finding suitable simplifications for complex words. This approach was further extended by Przybyła and Shardlow (2020) to build a multi-word LS system for English. The LSBert system (Qiang et al., 2020a) and our

adaptation of it to Spanish and Portuguese will be described in more details in Section 4.2.

Lexical simplification in languages other than English attracted less attention. For Spanish, several LS systems have been proposed so far:

- LexSiS (Bott et al., 2012) An unsupervised lexical simplification system for Spanish that uses an online dictionary and Web as a corpus to compute three features (word vector model, word frequency, and word length) for finding the best substitution candidates. Morphological generation of the right inflection for the best substitute is done by a combination of hand-crafted rules and dictionary look-up.
- CASSA (Baeza-Yates et al., 2015) An unsupervised lexical simplification approach for Spanish that uses Google Books Ngram Corpus, the Spanish OpenThesaurus, and web frequencies for finding the best substitution candidates. This approach only finds the best lemma and does not perform morphological generation of the right inflection.
- TUNER (Ferrés et al., 2017a) An unsupervised lexical simplification approach for Spanish, Portuguese, Catalan, and Galician. This system achieves the state-of-the-art results in lexical simplification for Spanish. It will be described in more details in Section 4.1, as one of the systems we adapt to all three languages and use in our experiments.
- EASIER (Alarcón et al., 2021) Neural lexical simplification systems for Spanish, which leverage pretrained word embedding vectors and BERT models. The systems were evaluated only for three sub-tasks: CWI, SG, and SS. The CWI sub-task was evaluated using the CWI 2018 shared task dataset for Spanish (Yimam et al., 2018). The other two sub-tasks (SG and SS) were evaluated using the EASIER-500 corpus (Alarcón et al., 2021). The fourth sub-task, ranking of substitutes (SR), was not evaluated as no Spanish lexical simplification datasets existed that could be used for that purpose (Alarcón et al., 2021).

For Portuguese (regardless of the language variant), only three systems that perform text simplification were proposed so far. All three were built and evaluated for Brazilian Portuguese. The system proposed by Specia (2010) is a machine translation-based sentence simplification system. It is a fully supervised system that relies on parallel original-simple sentences for training. It performs several transformations at the same time: lexical simplification, word and clause reordering, syntactic simplification, etc. The other two systems (Hartmann and Aluísio, 2020; Hartmann et al., 2020) are lexical simplification systems which are particularly designed to simplify texts for children.

Lang.	Dataset	#instances	#targ.	#syn.	#annotators				Annotator type
					CWI	SG	SS	SR	
EN	LSeval (De Belder and Moens, 2012)	430	43	5.04	0	5	(SG+SS)	5	AMT, US-based, >95% acc. rate
EN	LexMTurk* (Horn et al., 2014)	500	459	12.58	0	50	(SG+SS)	0	AMT, US-based, >95% acc. rate
EN	CEFR-LS (Uchida et al., 2018)	406	406	2.35	0	0 1 0		0	native English speaker, expert
PT(BR)	SIMPLEX-PB 3.0 (Hartmann and	1,719	757	7.31	0	2	(SG+SS)	NA	Expert linguists
	Aluísio, 2020)								
FR	FrenLyS (Rolin et al., 2021)	196	196	4.03	?	0	3	20	SG: automatic, SS: expert linguists,
									SR: native speakers
ZH	HanLS (Qiang et al., 2021)	524	524	8.51	2	5	(SG+SS)	mix	Native speakers
JP	SNOW E4 (Kajiwara and Yamamoto,	2,330		4.50	0	1	5	5	Crowdsourced
	2015)								
JP	BCCWJ (Kodaira et al., 2016)	2,010	210	4.30	0	5	mix	mix	native, >95% acc.rate
ES	EASIER-500 (Alarcón et al., 2021)	500	500	3	1	1	(SG+SS)	NA	Expert linguist

TABLE 1 Datasets used in evaluation of lexical simplification systems.

Commonly used benchmark datasets are denoted with an "*" next to their name. The column "#targ," denotes the total number of unique target words (in many datasets, the same target word is given in 10 different contexts). The column "#syn." denotes the average number of unique simpler synonyms proposed by all annotators per instance. The column "#annotators" denotes the number of annotators that provided their annotations for the specific task: CWI (complex word identification), SG (generation of possible substitutes), SS (substitute selection), and SR (substitute ranking). "?" denotes that the information is not available in the paper that describes the corresponding dataset.

To showcase the usability of the Portuguese portion of our dataset, we adapt TUNER and LSBert to (Brazilian) Portuguese.

2.2. Existing evaluation datasets for lexical simplification

The main characteristics of the existing datasets for evaluation of lexical simplification systems are given in Table 1. The following characteristics were taken into account:

- *#instances*: The total number of instances/contexts;
- *#targ*.: Total number of target/complex words;
- *#syn*.: The average number of simpler synonyms per target word;
- *CWI*: The number of annotators (per instance) who pointed out complex words that need to be simplified;
- *SG*⁶: the number of annotators (per instance) who suggested potential substitutes;
- SS: The number of annotators (per instance) who selected substitutes (from the list of potential substitutes) that correctly preserve the original meaning and fit in the context (preserve both semantics and grammaticality);
- *SR*: The number of annotators (per instance) who ranked the selected substitutes based on their simplicity;

• *annotator type*: Description of the annotators as given in the respective papers that describe the datasets.

The existing datasets all follow different procedures for data selection and annotation. For example, the context sentences have been selected from various genres and topics:

- Wikipedia: LexMTurk (EN)
- Introduction parts from introductory textbooks on various topics (economics, psychology, sociology, etc.): CEFR-LS (EN)
- Internet texts (balanced): LSeval (EN), BCCWJ (JP)
- Mixture of textbooks and dialogues for children: SIMPLEX-PB 3.0 (PT-BR), FrenLyS (FR)
- Mixture of original and translated texts: HanLS (ZH)
- Newspapers: SNOW E4 (JP)

The procedures used to choose target words (i.e., complex words to be simplified) also vary across the datasets. Most of them use a fully automatic method. Nevertheless, they still differ in which fully automatic method they use. LSeval (EN), CEFR-LS (EN), SIMPLEX-PB 3.0 (PT-BR), BCCWJ (JP) select the target words by leveraging dictionaries of easy/complex words, LexMTurk (EN) by leveraging automatic alignment of original and Simple English Wikipedia, while SNOW E4 (JP) bases the choice of target words on dictionaries of easy words and word frequency counts in newspapers.

For FrenLyS (Rolin et al., 2021), instances were selected from two corpora, ALECTOR (Gala et al., 2020) and texts from various textbooks. For the instances originating from ALECTOR corpus, complex words were identified based on the

⁶ In some cases, the same annotators were asked to suggest substitutes that fit in the context (preserve the original meaning and grammaticality). In those cases, the number of annotators is given jointly for columns SG and SS and denoted as (SG+SS).

information gained from a reading experiment with dyslexic children. For the instances originating from various textbooks, complex words were identified based on a reading experiment with various readers and the recorded reading time. Given that dyslexic readers have different lexical simplification needs than neurotypical readers (Rello et al., 2013; Rello, 2014), it is important to identify the target audience when working with such corpora.

For obtaining a list of replacement candidates, some datasets relied on human annotation and some on automatic generation of the candidates. For LSeval (EN), LexMTurk (EN), HanLS (ZH), SNOW E4 (JP), candidates were generated and selected jointly, by asking annotators to suggest simpler synonyms that fit the context (by preserving original meaning and grammaticality). For creation of BCCWJ (JP) dataset, one set of annotators suggested replacement candidates, and the other set of annotators assessed whether or not those candidates are a good fit (preserving meaning and grammaticality). For CEFR-LS (EN) dataset, replacement candidates were obtained automatically based on thesaurus and dictionaries. For FrenLyS (FR), replacement candidates were also obtained automatically, but in this case, based on several sources: thesaurus, word embeddings, and neural language models.

The instructions given to annotators for how to select the right candidates, i.e., to judge candidate fitness in context also varied across the datasets. During the creation of FrenLyS (FR), the annotators were instructed to judge the candidate substitution correct if it does not change original meaning, and to accept hypernyms and hyponyms, and small changes in nuances as correct (Rolin et al., 2021). During the creation of CEFR-LS (EN), in contrast, the annotators were instructed to select the given substitution candidate, only if it successfully conveys the nuance of the target word in the specific context and does not affect the meaning of a sentence (Uchida et al., 2018). For SNOW E4 (JP) datasets, the annotators were instructed to select the substitution candidate if it sounds natural in the context and does not change the original meaning (Kajiwara and Yamamoto, 2015). In contrast to all other datasets, substitute selection in BCCWJ (JP) was done using majority vote on top of human suggestions.

The procedures for candidate ranking also varied across the evaluation datasets. LSeval (EN), FrenLys (FR), and SNOW E4 (JP) used human ranking with different type and number of annotators. Candidate ranking in LexMTurk (EN) was automatic, based on the frequency of the candidate being proposed by 50 crowdworkers. In CEFR-LS (EN) dataset, candidates were also ranked automatically, but in this case, based on leveraging special sources (word lists and CEFR language learning framework). In HanLS (ZH) and BCCWJ (JP) datasets, candidates were first ranked by humans, and then the final ranks were computed by using the mean value of the ranks given by each annotator during suggestion of adequate replacement candidates (HanLS), or maximum likelihood estimation on top of five human rankings (BCCWJ).

For all above-mentioned differences in procedures used to select and annotate instances in existing evaluation datasets, it is not possible to compare performances of lexical simplification systems in different languages. Furthermore, some of the datasets offer too few simpler synonyms per target word to make evaluation metrics applied to them reliable. Our dataset, in contrast, uses comparable procedures for selection and annotation of instances across three languages, which makes the results of lexical simplification systems comparable across languages. Due to a high number of annotators per instance (25), our dataset also offers a higher number of simpler synonyms per target word, which results in higher reliability of standard evaluation metrics applied on it.

2.3. Evaluation metrics for lexical simplification

The common metrics for the evaluation of lexical simplification systems are defined by Paetzold and Specia (2016a) as:

- **Potential** The percentage of instances for which at least one of the substitutions generated is present in the gold standard;
- **Precision** The percentage of generated candidates that are in the gold standard;
- **Recall** The percentage of gold-standard substitutions that are included in the generated substitutions;
- F1 The harmonic mean of Precision and Recall.

All metrics can be calculated taking into account all outputs of the system for each instance, or only taking the first k(ranked) outputs of the system for each instance in which case they are usually denoted as Potential@K, Precision@k, Recall@k, and F1@k. Using the metrics "@k" instead of the original ones is recommended when comparing systems of different architectures, especially those that notably differ in the number of simpler substitutes generated per instance. Systems that output a higher number of simpler substitutes would have notably higher Potential and Recall, and lower Precision than the systems that output only a few simpler substitutes.

Each metric aims to evaluate different aspects of lexical simplification systems. Potential is used to evaluate the substitution candidate generation (SG) phase of the systems, i.e., the potential of a system to generate at least one simpler substitute. Precision evaluates system's performance at generating and selecting substitution candidates (SG and SS phases). Recall evaluates versatility of generated simpler synonyms. When used only on the k best ranked simpler

TABLE 2 Statistics on the TSAR-ST 2022 Lexical Simplification Dataset.

Language	Instances Tokens		Contexts	Suggestions		
				Min	Max	Avg
English (EN)	386	330	369	2	22	10.55
Spanish (ES)	381	356	326	2	19	10.28
(Brazilian) Portuguese (PT-BR)	386	348	356	1	16	8.10
All	1,153	1,031	1,051	1	22	9.64

The column "Instances" signifies the total number of instances used in the crowdsourcing lexical simplification experiments, while columns "Tokens" and "Contexts" signify the total number of unique tokens and contexts in those instances, respectively. The columns "Min", "Max", and "Avg" represent the minimal, maximal and average number of unique simpler substitutes suggested by the crowdsourced workers per instance.

substitutes, all measures additionally evaluate the ranking capabilities of the system (SR phase).

Recall, and thus also F1, is additionally influenced by the number of simpler substitutes per instance, and their quality/correctness, in the gold data. Therefore, Recall and F1 are meaningful metrics only in the case of carefully curated gold data, i.e., gold data that is known to be easier for the target population. On the benchmark datasets where gold data consists of candidate replacements suggested by crowdsourced workers as simpler synonyms for the given target word in context, Recall and F1 may not be the right evaluation metrics. Such datasets, instead, are better suited for using Potential and Precision, especially if each instance contains more than just a few simpler synonyms suggested (as it is the case in our dataset, see Table 2 in Section 3.3).

3. New evaluation dataset for lexical simplification

3.1. Data collection

We compiled a new dataset of examples of lexical simplifications across Portuguese, Spanish and English. Crowdsourced workers were presented with instances (sentences) in which a single token is marked as requiring simplification. They were asked to provide simpler synonyms for the marked words, taking into account that the original meaning of the sentence should be preserved.

The following example is taken directly from the English portion of our dataset:

The daily death toll in Syria has declined as the number of **observers** has risen, but few experts expect the U.N. plan to succeed in its entirety. We deliberately present target words in context as the returned words should be grammatically aligned with the context and semantically consistent with the original term. Grammatical alignment is important as the words should be directly replaceable in the original context. If a word is in the wrong tense, or requires a missing preposition, then the resulting sentence would be ungrammatical. Similarly, if a word is a good grammatical fit in the sentence, but is not semantically consistent with the original term, then the resulting sentence will be difficult for readers to understand.

In our selection procedure we only identify single words (as opposed to complex multi-word expressions) for annotation. This simplifies the problem space and makes model input, etc. easier to process. Whilst we acknowledge that the problem of simplifying complex multi-word expressions is important, we leave this to a more thorough handling in dedicated work on the subject. Although we did not select for multi-word expressions, we did allow annotators to return multiple words if they could not think of a relevant single-word simplification. This allowed for the insertion of function words, or for compound terms to be returned in the rare cases where this was necessary.

In each of our datasets we have explicitly chosen to identify simplifications for words which are known to be difficult for a reader to understand. Another option would have been to simplify every token in a context, or select tokens at random or according to some heuristic (such as low lexical frequency). If we had selected random tokens, many would have not required simplification — leading to a waste in annotator effort. If we had used frequency or length, we may have missed words which do not follow these patterns — leading to a biased dataset. By leveraging human annotations of complexity we are able to explicitly produce simplifications for terms which require it.

All instances used in our dataset (for all three languages) were taken from existing corpora that had information about which words need to be simplified. For English and Spanish, we used the respective portions of the 2018 edition of the Complex Word Identification shared task (Yimam et al., 2018). For (Brazilian) Portuguese, we used the PorSimplesSent dataset (Leal et al., 2018). Sentences in all three corpora (English CWI-2018, Spanish CWI-2018, and PorSimplesSent) often contained several words marked for requiring simplification. For the examples in our new dataset, however, we opted for marking only one of those words in each instance. The reason for this was two-fold. First, if we had marked all originally marked complex words in each sentence, the task would be much more complex for crowdsourced workers. Apart from having to propose a simpler synonym that fits well in the context, they would need to pay attention to how all proposed simpler synonyms in a given sentence interact. This would lead to longer cognitive effort by the annotators and higher number of incorrect substitution candidates. It would also make the validation of collected annotations difficult. For example, it could happen that each

suggested simpler synonym in a sentence is correct on its own, but together with other suggested simpler synonyms in the sentence it does not sound natural. Second, the state-of-the-art lexical simplification systems only perform simplification of a single complex word in a given sentence at time. If the given sentence contains several words marked as complex, the stateof-the-art lexical simplification systems would simplify them in several rounds, i.e., first simplifying one of them, then the next one (this time the context would be different as the first complex word is already replaced), and so on.

Specific collection protocols for the source data and complex words in each language are described in the subsections below.

3.1.1. Portuguese

We extracted instances of Portuguese complex words in context from the PorSimplesSent dataset (Leal et al., 2018). Each instance in PorSimplesSent was collected from Brazilian newspapers and was therefore of the Brazilian Portuguese (PT-BR) variety. The PorSimplesSent dataset consists of a collection of original and simplified sentences, whereby a trained linguist manually simplified a complex sentence according to detailed guidelines. To extract complex words from this dataset, we conducted automatic word alignment. A script compared each word within the original and simplified sentence pair and identified potential inconsistencies between the two. A native PT-BR linguist then manually examined these inconsistencies and recognized those instances which contained an accurate simplification for a particular target word. These target words were subsequently considered complex. In total, 386 sentence pairs were found to contain a complex word. Three hundred and forty-eight of these instances contained unique complex words, whereas 39 instances contained a duplicate complex word but in a unique context.

The 386 instances, containing the original complex word in context, were then shared with 25 crowd-sourced MTurk annotators located in Brazil. They were asked to provide the most suitable simplification for each given complex word. This resulted in 9,604 suggested simplifications. Analysis of the provided simplifications found that 70 unique candidate substitutions (10.25% of all suggested simplifications) were either (a) equal to the complex word, (b) not PT-BR, or (c) inappropriate (e.g., words that did not accurately preserve the meaning of the sentence or the original complex word). These candidate substitutions were excluded resulting in a final total of 2,742 unique candidate substitutions, or 8,620 repeated simplifications, for 386 instances.

3.1.2. Spanish

For Spanish, a set of 588 examples were extracted from the CWI Shared Task 2018 dataset⁷ (Yimam et al., 2018). Only the

examples with terms annotated as complex by five or more native language annotators were taken into account. This set was then reduced to 402 examples after a manual judgment process that involved two computational linguistics experts.

The manual judgment process was conducted to decide if the complex word was "simplifiable"⁸ in its context or not. The experts could choose from three options: "simplifiable", "not simplifiable", or "dubious". This resulted in three sets of judgements: (1) a set of 256 examples for which both experts agreed that the complex word is simplifiable, (2) a set of 113 examples for which both experts agreed that the complex word is not simplifiable, and (3) a set of 219 examples for which there is disagreement between the experts, or at least one of the experts indicated that they had doubts about the simplification. This manual judgment process was done using the aid of online dictionaries and thesaurus. Finally, after a joint revision of the 219 previously selected dubious examples, a subset of 146 examples was re-classified as simplifiable, thus leading to a total of 402 simplifiable instances.

After deleting repeated examples, a set of 393 unique simplifiable examples was obtained. An additional filtering step was applied afterwards. It involved removing cases of: (1) very similar examples; (2) complex words from other languages that are not yet commonly used and not (yet) accepted as valid words in Spanish (e.g., *hoax*); (3) complex words that have a sense in the sentence that is used in very specific locations (e.g., *jirón*— when refers to a sense related with a kind of street). This resulted in a final set of 384 instances for crowdsourced annotation of candidate substitutes for lexical simplification.

For every example a simpler substitute was proposed by a set of 25 annotators (3 splits of 128 instances were submitted to 75 different annotators, 25 annotators per split). The demographics data of the 75 annotators is as follows: *Gender*: Female (47), Male (28); *Age Ranges (years)*: 20–30 (54) 31–40 (17) 41–50 (2) 50–59 (1), Unknown (1); *Nationality*: Argentina (1), Greece (1), Italy (2), Venezuela (2), Portugal (4), Spain (6), Chile (13), Mexico (45), Unknown (1).

Once the crowdsourced annotation process was finished, it was decided to exclude three instances: two instances with the complex word repeated two times in its context, and a sentence which has a typographical error. This resulted in 381 instances in the final dataset.

The final dataset contains 356 different target words: 333 words appear once, 21 words appear twice, and two words appear three times. There are a total of 9,524 substitutions in the dataset and after joining the repeated substitutions in each instance we get a total of 3,918 different substitutions.

One of the authors reviewed the crowdsourced annotations and detected that they contain 137 incorrect substitutions

⁷ https://sites.google.com/view/cwisharedtask2018/datasets

⁸ That is if the experts were able to find a simpler substitute for the complex word.

(1.44%) and 93 dubious substitutions⁹ (0.98%), 230 substitutions equal to the complex word (2.41%), and 9,064 correct substitutions (95.17% of the total substitutions). Although these incorrect and dubious substitutions (according to the reviewer) were not excluded from the dataset, we think that this information can be used in further work (in collaboration with linguistic experts) to generate a cleaner set of examples¹⁰.

3.1.3. English

The English data from the 2018 edition of the Complex Word Identification shared task (comprising data from news, Wikinews and Wikipedia articles) (Yimam et al., 2018) was selected as the initial set of instances for annotation. This comprised 34,879 instances, each of which had binary complexity annotations from 10 native and 10 non-native English speakers. We selected all instances where terms had been annotated as difficult by at least five native annotators. We then removed any duplicate tokens and contexts giving 1,949 instances to select from.

We manually selected 400 instances from the set of 1,949 possible candidates. These were selected by a native English speaker and were identified as those instances where a lexical simplification could reasonably be produced, i.e., if the annotator could find at least one single word replacement to simplify the sentence, the instance was kept. If this was not possible, the instance was discarded. Each instance comprised of a token and the context in which that token occurred.

These instances were passed on to annotation using Amazon's Mechanical Turk. We did not record demographic statistics on the annotators, however we requested annotators from English speaking countries to maintain quality. Each instance was annotated by 25 annotators, each of whom were instructed to return a single word to simplify the sentence. The native English annotator reviewed all suggestions (n = 10,000) to determine if they were acceptable in the context of the task and removed unsuitable annotations where the guidelines had not been followed (i.e., not simplifying, returning dictionary definitions, etc.). When suggestions were removed further annotations were requested too ensure each instance had 25 suggestions. 14 instances were removed during this process where it was clear that no good suggestions could be found by the crowd workers, leading to 386 final instances.

3.2. Data annotation

We annotated the data using crowdsourcing platforms (Amazon Mechanical Turk for English and Portuguese, Prolific

for Spanish). The Spanish instances were annotated first, and the guidelines presented in that annotation round were translated into English and Portuguese, with minimal editing to ensure that the task remained the same across languages. One notable difference was that whereas Spanish and Portuguese require gender agreement for replacements, this usually does not apply in English. The text used in the guidelines across the three languages is shown in Appendix I. When using crowdsourcing, it is important to keep the guidelines brief to encourage the annotators to read them, whilst also allowing them to gain enough knowledge to complete the task without being rejected.

We only rejected instances in cases where the data that was returned was clearly abusing the guidelines (e.g., dictionary definitions, whole sentences, nonsense input). All instances were manually verified for correctness and instances were manipulated to fit the context grammatically where necessary. Only affix changes were applied, keeping the original semantics of the simplifications.

3.3. Data description and statistics

The data comprises of 1,153 instances, split across the three languages. Summary statistics for the data are shown in Table 2, and one instance from each language in Table 3. The data is intended as a benchmark test set for Lexical Simplification systems in one of the languages, and for multilingual systems. The gold annotations consists of all simpler substitutes suggested by crowdsourced workers, checked for quality by at least one computational linguist who is native speaker of the respective language. The suggested simpler synonyms are ordered (in descending order) by the number of annotators who suggested them.

Table 2 shows that the dataset is evenly divided between the three languages that we have chosen to annotate. In each language, we started off with 400 instances, but discarded instances that were troublesome for annotators, leaving 386 instances in English and Portuguese and 381 instances in Spanish. The number of unique tokens and contexts is close to the total number of instances, indicating that there is little repetition of tokens across the dataset. Whilst repeated instances may be interesting to explore the effect of context on polysemy and replacement fit, the repeated tokens will likely have similar replacements, in a similar order. For this reason they were kept to a minimum. No two instances are identical (e.g., having the same token and context). The number of suggestions varies similarly across data subsets, with one or two suggestions being offered in the minimal case and up to 22 suggestions being offered for the largest case. On average, there are 9.64 suggestions returned per instance in our dataset. The most frequent suggestion typically far outnumbered the long tail of other, less popular suggestions.

⁹ Dubious substitutions are those which, according to the reviewer, are probably incorrect.

¹⁰ The Spanish part of the dataset has been described in more details in Ferrés and Saggion (2022).

TABLE 3 Examples of instances from the dataset (the number after ":" in the third column represents the number of crowdsourced workers that suggested that replacement).

Lang.	Sentence (target word in bold)	Simpler substitutes suggested by 25 crowdsourced workers
EN	A local witness said a separate group of attackers disguised in burgas—the head-to-toe robes worn by conservative Afghan women—then tried to storm the compound	Concealed:4, dressed:4, hidden:3, camouflaged:2, changed:2, covered:2, disguised:2, masked:2, unrecognizable:2, converted:1, impersonated:1
ES	Conforme avanzaba el debate en el Congreso de Filadelfia, Lee iba asumiento una posición más favorable a la independencia total y no sólo a la autonomía del Imperio Británico, su convicción de la necesidad de la independencia logró convencer a delegados de otras colonias e incluso persuadió a sus propios electores de Virginia, temerosos que Lee pudiera llegar demasiado lejos.	Creencia:5, seguridad:5, certeza:5, convencimiento:3, ideal:2, fé:1, persuación:1, fuerte creencia:1, idea:1
PT-BR	Quem não conseguir esgotar o armazenamento de diesel puro não pode misturar com o b2 porque o produto ficaria fora de especificação.	Acabar:10, esvaziar:7, acabar com:4, gastar:1, consumir:1, diminuir:1, zerar:1

Lexical simplification systems typically employ a natural language engineering approach, leveraging state of the art technology, rather than simply training systems to directly perform simplification. Because of this, our data is intended as a benchmarking test set and not for training. We expect that systems will leverage other resources to improve their performance. Therefore, we have not split our data into training and test subsets.

A number of domains are represented in our dataset, resulting from the diverse corpora that were used to select contexts for each language. The English data was selected from the CWI-2018 dataset, containing Wikitext and news data. The Spanish data was selected also from the CWI-2018 dataset for Spanish, which contains data from the Spanish Wikipedia. The Portuguese data was taken from general and scientific news articles.

We paid annotators at the following rates per instance in each language: \$0.03 for English, \$0.14 for Spanish and \$0.02 for Portuguese. An additional \$0.01 was paid in platform fees per instance for English and Portuguese. In total, this equates to a spend of \$2481.73 to annotate our entire dataset.

We can easily add in further languages at a future point by running further annotations with the same protocol. For languages where a CWI dataset exists, we can use the protocol employed for English and Spanish, whereas for languages without CWI data, we can follow the same protocol as for Portuguese.

3.4. Limitations of the dataset

Although of high quality and being the only multilingual evaluation dataset for lexical simplification, our dataset has some notable limitations.

One limitation of the English and Spanish portion of the dataset is that all instances come from the same source, covering only one genre. Therefore, they only provide a reliable evaluation of lexical simplification systems which focus on those specific genres.

Another limitation of our dataset is that the provided replacements represent simpler synonyms according to the crowdsourced workers, rather than experts in the area. The high number of annotators per instance (25) mitigates this issue to some extent, as it offers a possibility for ranking the replacement candidates according to the number of times they were suggested by different people. The most frequently proposed replacement candidates could thus be considered of a good quality, but some re-ranking may be required to confirm this.

3.5. Intended use

The new dataset is envisioned as the first evaluation dataset that allows for fair comparison of lexical simplification systems across different languages (English, Spanish, and Portuguese), due to comparable procedures for selecting and annotating instances in all three languages. As it contains a high number of simpler replacements suggested for each target/complex word, it is particularly valuable for evaluating substitution generation (SG) and substitution selection (SS) modules of lexical simplification systems. Due to a high number of annotations per instance (25 crowdsourced workers), it can also be used to evaluate substitute ranking (SR) capabilities of lexical simplification systems. However, as gold data was crowdsourced, and only professionally checked for preservation of grammaticality and original meaning, the ranking of substitution candidates can only be used as a proxy for general notion of simplicity, and not as simplicity ranking for any particular target user/population.

Given a high number of simpler replacements offered for each target/complex word, this dataset can also be used as a starting point for building evaluation datasets for lexical simplification intended for some particular audience. In that case, the original set of simpler replacements should be filtered and ranked based on expert annotations [e.g., by carers or expert psycholinguists aware of particular simplification needs of the target user(s)] or user studies with the target users (e.g., comprehension tests, eye-tracking studies, etc.).

4. Experiments

Our newly compiled evaluation dataset for lexical simplification allows us to compare the performance of different lexical simplification approaches across the three languages we have incorporated (English, Spanish, and Portuguese). To demonstrate this, we adapt the state-of-the-art lexical simplification system for Spanish (TUNER) and the state-of-the-art lexical simplification system for English (LSBert) to all three languages (English, Spanish, and Portuguese). As well as representing state-of-the-art systems in two languages, these systems also represent two different approaches to lexical simplification. TUNER relies on static resources such as vocabularies and thesauri, whereas LSBert relies on a large scale BERT-based language model.

We further compare the performances of those systems on our benchmark dataset using several evaluation metrics that aim to capture different aspects of system's performances. The next three subsections describe: (1) the TUNER lexical simplification system (a non-neural system) and its adaptation to the three languages (Section 4.1); (2) the LSBert lexical simplification system (a system that leverages neural language models) and its adaptation to the three languages (Section 4.2); and (3) the evaluation metrics used for comparing performances of TUNER and LSBert across the three languages (Section 4.3).

4.1. TUNER-LS

The TUNER Candidate Ranking System used in this evaluation is an adaptation of the TUNER Lexical Simplification architecture (Ferrés et al., 2017b) to work with Spanish, Portuguese and English. Some components for English were obtained from the YATS Simplifier for English (Ferrés et al., 2016). The TUNER simplifies words (common nouns, verbs, adjectives, and adverbs) in context. The adaptation presented here omits the Complex Word Identification phase and the Context Adaptation phase, returning the lists of ranked candidates with correct inflections instead of returning the complete sentence simplified using the top ranked candidate.

The adapted system has the following phases (executed sequentially): (1) Sentence Analysis, (2) Word Sense Disambiguation (WSD), (3) Synonyms Ranking, and (4) Morphological Generation. The Sentence Analysis phase uses the FreeLing 4.0 system to perform tokenization, sentence

TABLE 4	Statistics of the EuroWordNet thesaurus and the Wikipedia
collectio	ns processed for the TUNER-LS system.

Language	EuroW	ordNet	Wikipedia ¹²		
	#entries	#senses	#documents	#words	
EN	63,649	87,792	99,943	15M	
ES	36,571	50,397	1,061,535	349M	
PT	35,635	45,737	956,553	203M	

splitting, part-of-speech (PoS) tagging, lemmatization, and Named Entity Recognition.

The WSD algorithm used is based on the Vector Space Model approach for lexical semantics. The thesauri used for WSD were extracted from FreeLing 4.0 data which is derived from Multilingual Central Repository (MCR) 3.011 (release 2012). Each thesaurus contains a set of synonyms and its associated set of senses with related synonyms (see the number of entries and senses of the thesaurus for each language used in Table 4). The WSD algorithm uses a word vector model derived from a large text collection from which a vector for each word in the thesaurus is created by collecting co-occurring word lemmas of the word in 11-word window (five content words to each side of the target word) contexts (only nouns, verbs, adjectives, and adverbs). Then, a common vector is computed for each of the word senses of a given target word (lemma and PoS) by adding the vectors of all words in each sense. When a complex word is detected, the WSD algorithm computes the cosine distance between the context vector computed from the words of the complex word context (at sentence level) and the word vectors of each sense from the model. The word sense selected is the one with the lowest cosine distance between its word vector in the model and the context vector of the complex word in the sentence or document to simplify.

The word vector models for each language were extracted from Wikipedia dumps. For Spanish and Portuguese, the Spanish Wikipedia and Portuguese Wikipedia were used, respectively. For English, the word vectors model was extracted from the Simple English Wikipedia¹³. The plain text of the documents was extracted (see in Table 4 the number of documents and words extracted from the Wikipedia dump). The FreeLing 3.1 system was used to extract the lemmas and PoS tags of each word, from an 11-word window (five content words to each side of the target word).

The Synonym Ranking phase ranks synonyms using word form (or lemma) frequency as a simplicity measure. The

¹¹ http://adimen.si.ehu.es/web/MCR/

¹² We used the Simple English Wikipedia for English (EN).

¹³ Simple Spanish Wikipedia and Simple Portuguese Wikipedia do not exist.

frequency list used by each of the languages are: 1) Spanish Wikipedia form counts for Spanish, 2) Portuguese Wikipedia form counts for Portuguese, and 3) Simple English Wikipedia word form counts for English.

The Morphological Generator module generates the correct inflected forms of the final selected synonyms. Given a set of pairs < LEMMA, PoS >, with the lemma corresponding to a substitution candidate and the PoS tag corresponding to the PoS Tag of the complex word in the sentence, this module returns the inflected forms of the candidates. For Spanish and Portuguese this module uses an algorithm that combines lexicon-based generation and predictions from decision-trees (see Ferrés et al., 2017a for a more detailed description of this system). For English, the SimpleNLG Java API¹⁴ was used with its default lexicon to perform this task using the candidate and the PoS tag of the complex word in the original sentence.

4.2. LSBert

The LSBert¹⁵ (Qiang et al., 2020b) state-of-the-art lexical simplifier for English has been adapted to deal also with Spanish and Portuguese. The LSBert system uses a pretrained representation of BERT to propose substitution candidates with high grammatical and semantic similarity to a complex word in a sentence. LSBert uses the masked language model (MLM) of BERT to predict a set of candidate substitution words and their substitution probabilities. BERT is fed with the original sentence concatenated with a copy of the sentence in which the complex word has been masked. LSBert combines five different features for Lexical Simplicity Ranking: BERT prediction order, a BERT-based language model, the PPDB database, word frequency, and word semantic similarity with fasttext.

In the original LSBert system the simplification algorithm selects the top ranked candidate and performs the simplification only if the top candidate has a higher frequency (frequency feature) or lower loss (language model feature) with respect to the original complex word, otherwise returns the same complex word. In our adaptation, we modified the system to retrieve up to K = 5 candidates from a BERT-based pre-trained model in this way: after the ranking procedure the top K = 5 candidates that a have higher frequency (frequency feature) or have a lower loss (language model feature) with respect to the complex word are selected.

For English, we used the same resources described in the original LSBert paper (Qiang et al., 2020b): (1) BERTlarge WWM, (2) Porter Stemmer, (3) Fasttext CBOW model for English¹⁶, (4) SUBTLEX zipf values Brysbaert and New (2009) and (5) PPDB 2.0^{17} . The language specific resources used to adapt the system to Spanish and Portuguese were: (1) BERT-base based models: BETO (Cañete et al., 2020) for Spanish and BERTimbau for Portuguese (Souza et al., 2020), (2) Snowball stemmer for Spanish and Portuguese, (3) Fasttext CBOW model for Spanish¹⁸ and Portuguese¹⁹, and (4) SUBTLEX-ESP (Spanish) and SUBTLEX-PT (Portuguese) word frequencies in zipf format. The configurations parameters used were: probability-mask = 0.5, max-sequence-length = 350. The main features of the BERT-based models used in the LSBert approach for each language are presented in Table 5.

4.3. Evaluation metrics

We evaluate TUNER and LSBert lexical simplification systems for English, Spanish, and Portuguese on the respective portions of our dataset using **Potential**, **Precision**, and **Recall**. As mentioned in Section 2.3, those metrics are the usual metrics for evaluation of lexical simplification systems. For a more detailed comparison of the systems across languages and architecture, we use Potential@1²⁰, Potential@3, Potential@5, Precision@3, Precision@5, and Recall@5. In all cases, Precision and Recall are computed for each instance and then averaged over all instances in the respective portion of the dataset.

We additionally define and calculate Accuracy@1—the percentage of instances for which the best ranked substitution generated by the system is the same as the most frequently suggested simpler synonym in the gold data. The rationale for introducing this metric is to be able to evaluate the possibility of using state-of-the-art lexical simplification systems in a fully automatic lexical simplification setup, where presented with a sentence and a target word in it, the system automatically generates a new sentence with a correct replacement. To evaluate this scenario, we assume that the most frequently proposed simpler replacement in the gold data corresponds to a surely correct simplification in the given context.

5. Results and discussion

Table 6 presents results of LSBert and TUNER models (in all three languages) on the new dataset. Several instances together with their gold annotations and outputs of LSBert and TUNER systems are presented in Tables A1–A3 in Appendix II.

¹⁴ https://github.com/simplenlg/simplenlg

¹⁵ https://github.com/qiang2100/BERT-LS

¹⁶ https://dl.fbaipublicfiles.com/fasttext/vectors-english/crawl-300d-2M-subword.zip

¹⁷ http://nlpgrid.seas.upenn.edu/PPDB/eng/ppdb-2.0-tldr.gz

¹⁸ https://dl.fbaipublicfiles.com/fasttext/vectors-crawl/cc.es.300.vec. gz

¹⁹ https://dl.fbaipublicfiles.com/fasttext/vectors-crawl/cc.pt.300.vec. gz

²⁰ Note that Precision@1 is equal to Potential@1.

Language	Model	Туре	Case	Layers	#parameters
English	BERT-WWM	BERT-large-wwm	uncased	24	336M
Spanish	BETO	BERT-base-wwm	uncased	12	110M
(Brazilian) Portuguese	BERTimbau	BERT-base	cased	12	110M

TABLE 5 Features of the BERT-based models used in the LSBert approach for each language.

TABLE 6 Performances of TUNER and LSBert lexical simplification systems on the new dataset (note that Potential@1 and Precision@1 give the same results as per their definitions).

System	Language	Potential		Precision			Recall		Accuracy@1	
		@1	@3	@5	@1	@3	@5	@3	@5	
LSBert	EN	60.1	82.4	87.8	60.1	47.0	37.0	16.8	21.6	30.8
	ES	28.9	49.3	61.4	28.9	23.3	19.6	8.4	11.7	9.7
	PT-BR	32.4	50.0	58.3	32.4	23.2	18.0	9.6	12.3	15.5
TUNER	EN	24.6	42.0	44.0	24.6	19.4	18.3	6.1	7.2	10.9
	ES	8.9	13.9	14.4	8.9	6.6	6.3	2.5	2.7	5.5
	PT-BR	17.3	26.9	26.9	17.3	12.7	12.0	5.1	5.3	10.6





5.1. Comparison of metrics

We have presented four metrics for the evaluation of lexical simplification and evaluated two systems across three languages using these metrics, as shown in Table 6. These metrics allow us to better understand the nature of the lexical simplification task and how to evaluate it well.

Our first metric, Potential (visualized in Figure 1), is a very relaxed metric that indicates whether any suggested candidate can be found in the gold standard. Systems can do artificially well on this metric by proposing many unrelated candidates (or indeed an entire vocabulary) and so we limit our evaluation to the first 1, 3, and 5 candidates proposed by a system. In this setting, we are able to attain scores that indicate that systems are usually able to suggest at least one relevant suggestion, with the probability being higher if more candidates are allowed.

Unlike Potential, Precision decreases as the number of candidates that are considered increases (Figure 2). A Precision@1 of 60.1 for LSBert in English indicates that the system is usually suggesting a word that is among the gold standard annotations. Precision drops as K increases because each candidate considered must be available in the gold standard. The candidates are returned ranked by likelihood of inclusion, so considering further less likely candidates is bound to lead to fewer inclusions for each instance.



TABLE 7 The statistics for the number of simpler substitutes in gold data and the systems' output.

Source	Language	Substitutes					
		Min	Max	Avg			
Gold	English	2	22	10.55			
	Spanish	2	19	10.28			
	Portuguese	1	16	8.10			
LSBert	English	5	5	5.00			
	Spanish	5	5	5.00			
	Portuguese	5	5	5.00			
TUNER	English	1	12	2.99			
	Spanish	1	12	1.59			
	Portuguese	1	10	1.80			

human editor, and thus lower the editing costs by speeding up manual simplification.

Recall (Figure 3) is consistently the lowest metric across our evaluation. We previously discussed (Section 2.3) the suitability of Recall when the gold data contains a large number of simpler substitutes per instance, as in our dataset. We thus report the values for Recall@3 and Recall@5 here only to have complete benchmark results. Recall is consistently the lowest metric, and is heavily dependent on the number of instances in the gold data for each instance.

We also considered a new metric for simplification which is Accuracy@1. This replicates the simplification scenario of an automated system that is selecting a simplification for replacement in a sentence. This metric requires the system to choose exactly the same candidate as the top-ranked candidate in the gold standard. Our best system (LS-BERT on English) achieved this 30.8% of the time. This only indicates to us that the system replicated the top choice of the gold standard at the given rate. It may have been the case that in the 69.2% of remaining cases another valid, but less likely candidate was chosen. This has some crossover with our potential metric, that shows that as the number of candidates considered increases the likelihood of finding a valid candidate also increases. Our Accuracy@1 metric is helpful to show how well our systems can perform in a very strict setting, replicating the annotators decisions, but should always be taken in the context of other metrics that give a different view of the systems' performance.

For comparing lexical simplification performances of the systems with different architectures and across the three languages, when using our dataset, we suggest using Potential@1, Potential@3, and Accuracy@1. Potential@1 and Accuracy@1 can be seen as the upper and lower measure of the usefulness of the system which automatically replaces the target word with a simpler (best ranked) synonym. Potential@3 is a valuable metric if the envisioned real-world usage of the system is to suggest a few simpler synonyms to the

5.2. Cross-system comparisons

The neural lexical simplification system (LSBert) outperforms the non-neural system (TUNER) in all three languages by all evaluation metrics (Table 6). This is not surprising given that LSBert is the state-of-the-art system for lexical simplification in English and it uses resources with better coverage than TUNER. This results in a higher number of final simplification suggestions generated by LSBert. While LSBert generates more simplification suggestions, in this work, we take into account only its first five (best ranked) simplification suggestions as we keep in mind real-world usage and want to have a fairer comparison between LSBert and TUNER. We allow TUNER to vary the number of substitutes returned according to the probabilities of the candidates words. Nevertheless, although for some instances TUNER was able to generate up to 12 candidates, for the majority of instances TUNER generated fewer candidates than this, with an average number of candidates between 1.59 and 2.99 (Table 7). The gold standard data typically contained more candidates than those proposed by either of the systems.

Appendix II contains Tables A1–A3, which demonstrate sample outputs of both systems in each language. LSBert seems to suggest more generalized contextual fits, whereas TUNER appears to suggest more conservative semantically accurate candidates. For example, the following sentence is presented to both systems:

War **maniacs** of the South Korean puppet military made another grave provocation to the DPRK in the central western sector of the front on Thursday afternoon. In this example, **maniacs** is the complex word requiring simplification. LSBert suggests: *criminals, victims, machines, freaks, people,* whereas TUNER suggests: *lunatics, madmans, maniacs.* The LSBert outputs fit well with the context, but are sometimes semantically incorrect (e.g., "war machines"). Another candidate "victims" is in fact the antonym of the original candidate ("maniacs"). TUNER on the other hand proposes three suggestions of reasonable quality. "madmans" is obviously the result of a pipeline error in the morphological adapter. "madman" has been incorrectly pluralized and the correct form should be "madmen". The relationship of these candidates to the gold standard data is presented in Table A1.

A further manual inspection of the systems' output revealed that TUNER is capable to outperform LSBert in some special cases, where the target word appears in a less commonly used context. In such cases, we find that LSBert suggests more frequently used words, but at the cost of severely changing the original meaning of the sentence.

5.3. Cross-lingual comparisons

We can make two key observations comparing the LS approaches across the three languages. Firstly, both architectures perform significantly better for English than for the other two languages (Table 6). Secondly, the TUNER architecture performs better for Portuguese than for Spanish, by all evaluation metrics (Table 6) and the LSBert architecture performs similar in Portuguese and Spanish with the exception of the Accuracy@1 metric.

These facts may be explained by the following factors: (1) Linguistic differences between English and Romance languages and between Spanish and Portuguese, (2) Dataset specific differences, and (3) The tools and resources used in the simplification algorithms (which all have better performances and coverage for English than for the other two languages). It is reasonable to assume the hypothesis that linguistic differences among Spanish, Portuguese and English can influence the results of our experiments. Spanish and Portuguese are part of the Ibero-Romance sub-family of Romance languages and English is part of the Germanic family of languages. According to Ethnologue²¹ lexical similarity²² between Spanish and Portuguese is about 89%.

On the other hand, although the procedures to collect the subsets for each language were very similar, minor differences in both the level of lexical complexity of complex words and the gold annotated substitutions in each language-specific subset of the trilingual dataset could have influenced the results. The instructions were translated into each language and adapted for language specific concerns (e.g., inflection handling). It is possible that the translated guidelines led to different interpretations of the task, which could affect the internal consistency of each dataset. Further, the annotator pools in each language were selected according to those crowd workers available at annotation time. These groups are mutually exclusive as they were selected for their first language, and so it may be the case that one group returned more or less consistent annotations than the others.

Moreover, the two different algorithms (and associated resources) have a notable influence in the differences of results among the three languages tested. For the LSBert algorithm we have several resources to compare among the three languages: BERT model, word-embeddings, stemmer, PPDB, and frequency files. The BERT model is the most important feature as it is used in the Substitution Generation phase and the Substitution Ranking phase. Regarding the comparison among language-specific BERT models in the experiments reported in Section 4, we have used BERT-large-uncased-WWM for English, BERTimbau BERT-base cased for Portuguese, and BETO BERTbase uncased for Spanish. The size of BERT (BERT-large vs. BERT-base) seems not to be the most important factor for those differences, as in follow-up experiments with LSBERT with BERT-base-uncased for English we obtained a Potential@1 of 0.544 and Accuracy@1 of 0.251, thus still outperforming the results for Portuguese and Spanish. Those results indicate that the BERT model in English has some attributes that greatly outperform the other models in Spanish and Portuguese.

For the TUNER architecture we can compare the following resources: NLP resources, thesaurus, context-vectors, and frequency counts lists. The most important and influential resources are the thesaurus and the context-vectors. The number of thesaurus entries and senses is notably greater in English with respect to Spanish and Portuguese (Table 4). This obviously can have a big impact in the results and this might be additionally reflected in the number of generated simpler substitutes (Table 7). Moreover, the context-vectors are crucial to select the correct set of synonyms. For English we used the Simple Wikipedia, which has less data but uses Simple English words and grammar, and for Spanish and Portuguese we used the original Wikipedia (as Simple Spanish Wikipedia and Simple Portuguese Wikipedia do not exist).

6. Conclusion

In this article, we presented a new evaluation dataset for lexical simplification and benchmarked the state-of-theart lexical simplification systems for English, Spanish, and Portuguese. This dataset has several advantages over the existing evaluation datasets for lexical simplification:

²¹ www.ethnologe.com

²² Ethnologue's definition of Lexical Similarity: the percentage of lexical similarity between two linguistic varieties is determined by comparing a set of standardized wordlists and counting those forms that show similarity in both form and meaning.

- It is the first multilingual evaluation dataset, with instances in Portuguese, Spanish, and English, selected and annotated using comparable procedures. As such, it is the first dataset that offers a reliable comparison of system's performances across the three languages.
- Simpler replacements that preserve the original meaning and grammaticality of the sentence were suggested by 25 people per each target word, resulting in 10 simpler replacements for each target word across all three languages on average.
- The quality of crowdsourced replacement suggestions was checked by at least one native computational linguist in each language.

Due to a large number of simpler replacements per target word (10 on average), the new dataset offers a possibility for further adaptation to evaluation of lexical simplification systems intended for specific target audiences, by ranking the substitutes based on their simplicity for that specific audience.

To demonstrate the usefulness of the new dataset, we adapted the state-of-the-art neural (LSBert) and non-neural (TUNER) lexical simplification systems to all three languages and evaluated them on this dataset. We found that LSBert architecture outperforms TUNER architecture for all three languages (English, Spanish, and Portuguese). We also found that the performance of LSBert significantly drops when the system is adapted to Spanish and Portuguese. Taking into account that neither Spanish or Portuguese are low-resource languages, this finding poses a question about applicability of LSBert method to other languages with even less resources with satisfying performances.

6.1. Future work

6.1.1. Shared task

The new multilingual dataset for lexical simplification presented in this study will be used for the shared task organized as a part of the TSAR workshop at the EMNLP 2022 conference. The data will be released as part of the shared task and is split into trial and testing data. The systems that participate in the shared task will be evaluated with metrics used in this work. After the completion of the shared task, the dataset will be further cleaned and enriched based on the manual error analysis performed on the output of the competing lexical simplification systems.

6.1.2. Dataset extension

Our data is constrained to the genres that are represented in the original source corpora. Whilst these corpora are intended for general audiences (comprising of texts that are not specific to any one domain), they still represent the specific instance of style and form that is found in those corpora. Further work could replicate the process for selecting and annotating instances in specific domains, allowing the creation of LS systems in those domains. For example, medical text could be annotated and selected for medical complex words and then simpler alternatives suggested for these. In domain-specific research, it is important to consult domain-experts for annotation to ensure that the original complex words are properly understood and transformed.

We have incorporated three languages, according to the expertise found in our research team. We would welcome the addition of further languages to our corpus following the same annotation protocol. For languages where no existing CWI resource already exists, researchers can follow the selection protocol used for Portuguese.

6.2. Improved performance

Our results are intended to represent a strong baseline of performance. We expect that through the shared task we will see further systems that improve on these results, advancing the state of the art in lexical simplification. Our dataset is released as a test dataset only, and so we expect the majority of systems to be unsupervised in nature. Further work to produce additional resources for training lexical simplification systems will clearly help to further push the state of the art, although we expect that this task will remain a hybrid, rather than fully unsupervised task due to the complex nature of the pipeline operations required.

A further use of the dataset is in training multilingual models for lexical simplification. A multilingual model configured to work well for the languages in our dataset may also be able to perform simplification in a zero-shot setting for unseen languages that can be incorporated into the multilingual model.

Data availability statement

The original contributions presented in the study are publicly available. This data can be found here: https://github.com/LaSTUS-TALN-UPF/TSAR-2022-Shared-Task.

Author contributions

SŠ: team management, idea generation, and manuscript writing. DF: experiments, data collection (ES), and manuscript writing. MS: data collection (EN) and manuscript writing. KN: data collection (PT) and manuscript writing. MZ: data collection (PT). HS: data collection and team management. All authors contributed to the article and approved the submitted version.

Funding

HS and DF acknowledge support from the project Context-aware Multilingual Text Simplification (ConMuTeS) PID2019-109066GB-I00/AEI/10.13039/501100011033 awarded by Ministerio de Ciencia, Innovación y Universidades (MCIU) and by Agencia Estatal de Investigación (AEI) of Spain.

Acknowledgments

We would like to thank Kim Cheng Sheang (Universitat Pompeu Fabra) for his assistance with the review of the evaluation software.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships

References

Abrahamsson, E., Forni, T., Skeppstedt, M., and Kvist, M. (2014). "Medical text simplification using synonym replacement: adapting assessment of word difficulty to a compounding language," in *Proceedings of the 3rd Workshop on Predicting and Improving Text Readability for Target Reader Populations (PITR)* (Gothenburg: Association for Computational Linguistics), 57–65. doi: 10.3115/v1/W14-1207

Alarcón, R. (2021). Dataset of Sentences AnnotatedWith Complex Words and Their Synonyms to Support Lexical Simplification. doi: 10.17632/ywhmbnzvmx.2

Alarcón, R., Moreno, L., and Martínez, P. (2021). "Exploration of Spanish word embeddings for lexical simplification," in *Proceedings of the First Workshop on Current Trends in Text Simplification CTTS 2021* (Online).

Alarcón, R., Moreno, L., and Martínez, P. (2021). Lexical simplification system to improve web accessibility. *IEEE Access* 9, 58755–58767. doi: 10.1109/ACCESS.2021.3072697

Alva-Manchego, F., Scarton, C., and Specia, L. (2020). Data-driven sentence simplification: Survey and benchmark. *Comput. Linguist.* 46, 135–187. doi: 10.1162/coli_a_00370

Aprosio, A. P., Menini, S., Tonelli, S., Ducceschi, L., and Herzog, L. (2018). "Towards Personalised Simplification Based On L2 learners' native language," in *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLIC-IT)* (Torino), 305–310. doi: 10.4000/books.aaccademia.3559

Baeza-Yates, R., Rello, L., and Dembowski, J. (2015). "CASSA: a context-aware synonym simplification algorithm," in *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Denver, Colorado: ACL), 1380–1385. doi:10.3115/v1/N15-1156

Barthe, K., Juaneda, C., Leseigneur, D., Loquet, J.-C., Morin, C., Escande, J., et al. (1999). GIFAS rationalized French: a controlled language for aerospace documentation in French. *Tech. Commun.* 46, 220–229.

Biran, O., Brody, S., and Elhadad, N. (2011). "Putting it simply: a context-aware approach to lexical simplification," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics* (Portland, OR: ACL), 496–501.

Bott, S., Rello, L., Drndarević, B., and Saggion, H. (2012). "Can Spanish be simpler? LexSiS: lexical simplification for Spanish," in *Proceedings of the 24th International Conference on Computational Linguistics (COLING)* (Mumbai), 357–374.

Brysbaert, M., and New, B. (2009). Moving beyond Kučera and Francis: a critical evaluation of current word frequency norms and the introduction of a new and

that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/ frai.2022.991242/full#supplementary-material

improved word frequency measure for American English. Behav. Res. Methods 41, 977–990. doi: 10.3758/BRM.41.4.977

Cañete, J., Chaperon, G., Fuentes, R., Ho, J.-H., Kang, H., and Pérez, J. (2020). "Spanish pre-trained BERT model and evaluation data," in *PML4DC at ICLR 2020*.

Cooper, M., Reid, L., Vanderheiden, G., and Caldwell, B. (2010). Understanding WCAG 2.0. A Guide to Understanding and Implementing Web Content Accessibility Guidelines 2.0. World Wide Web Consortium (W3C).

Crystal, D. (1987). The Cambridge Encyclopedia of Language. Cambridge University Press.

De Belder, J., and Moens, M.-F. (2012). "A dataset for the evaluation of lexical simplification," in *Computational Linguistics and Intelligent Text Processing*, ed A. Gelbukh (Berlin; Heidelberg: Springer), 426-437. doi: 10.1007/978-3-642-28601-8-36

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). "BERT: pretraining 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* (Minneapolis, MN: Association for Computational Linguistics), 4171–4186.

Ferrés, D., AbuRa'ed, A., and Saggion, H. (2017a). Spanish morphological generation with wide-coverage lexicons and decision trees. *Procesamiento del Lenguaje Natural* 58, 109–116. Available online at: http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/5419/3183

Ferrés, D., Marimon, M., Saggion, H., and AbuRa'ed, A. (2016). "YATS: yet another text simplifier," in *NLDB* (Spain: Springer), 335–342. doi: 10.1007/978-3-319-41754-7_32

Ferrés, D., and Saggion, H. (2022). "ALEXSIS: a dataset for lexical simplification in Spanish," in *Proceedings of the Language Resources and Evaluation Conference* (Marseille. European Language Resources Association), 3582–3594.

Ferrés, D., Saggion, H., and Gómez Guinovart, X. (2017b). "An adaptable lexical simplification architecture for Major Ibero-Romance languages," in *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems* (Copenhagen: Association for Computational Linguistics), 40–47. doi: 10.18653/v1/W17-5406

Freyhoff, G., Hess, G., Kerr, L., Tronbacke, B., and Van Der Veken, K. (1998). Make it Simple, European Guidelines for the Production of Easy-toRead Information for People with Learning Disability. ILSMH European Association, Brussels. Gala, N., Tack, A., Javourey-Drevet, L., François, T., and Ziegler, J. C. (2020). "Alector: a parallel corpus of simplified French texts with alignments of misreadings by poor and dyslexic readers," in *Proceedings of the 12th Language Resources and Evaluation Conference* (Marseille: European Language Resources Association), 1353–1361.

Glavaš, G., and Štajner, S. (2015). "Simplifying lexical simplification: do we need simplified corpora?" in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing* (Beijing: ACL), 63–68. doi: 10.3115/v1/P15-2011

Gooding, S., and Kochmar, E. (2019). "Recursive context-aware lexical simplification," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (Hong Kong: Association for Computational Linguistics), 4853–4863. doi: 10.18653/v1/D19-1491

Hading, M., Matsumoto, Y., and Sakamoto, M. (2016). "Japanese lexical simplification for non-native speakers," in *Proceedings of the 3rd Workshop on Natural Language Processing Techniques for Educational Applications (NLPTEA)* (Osaka), 92–96.

Hartmann, N., Paetzold, G. H., and Aluísio, S. (2020). "SIMPLEX-PB 2.0: a reliable dataset for lexical simplification in Brazilian Portuguese," in *Proceedings of the The Fourth Widening Natural Language Processing Workshop* (Seattle), 18–22. doi: 10.18653/v1/2020.winlp-1.6

Hartmann, N. S., and Aluísio, S. M. (2020). Adaptação lexical automática em textos informativos do Português Brasileiro para o ensino fundamental. *Linguamática* 12, 3–27. doi: 10.21814/lm.12.2.323

Hmida, F., Billami, M. B., François, T., and Gala, N. (2018). "Assisted lexical simplification for French native children with reading difficulties," in *Proceedings of the 1st Workshop on Automatic Text Adaptation (ATA)* (Tilburg: Association for Computational Linguistics), 21–28. doi: 10.18653/v1/W18-7004

Horn, C., Manduca, C., and Kauchak, D. (2014). "Learning a lexical simplifier using wikipedia," in *Proceedings of ACL (Short Papers)* (Baltimore), 458–463. doi: 10.3115/v1/P14-2075

Kajiwara, T., and Yamamoto, K. (2015). "Evaluation dataset and system for Japanese lexical simplification," in *Proceedings of the ACL-IJCNLP 2015 Student Research Workshop* (Beijing: Association for Computational Linguistics), 35–40. doi: 10.3115/v1/P15-3006

Kodaira, T., Kajiwara, T., and Komachi, M. (2016). "Controlled and balanced dataset for Japanese lexical simplification," in *Proceedings of the ACL 2016 Student Research Workshop* (Berlin: Association for Computational Linguistics), 1–7. doi: 10.18653/v1/P16-3001

Leal, S. E., Duran, M. S., and Aluísio, S. M. (2018). "A nontrivial sentence corpus for the task of sentence readability assessment in Portuguese," in *Proceedings of COLING* (Santa Fe), 401–413.

Mencap (2002). Am I Making Myself Clear? Mencap's Guidelines for Accessible Writing.

OECD (2013). OECD Skills Outlook 2013: First Results from the Survey of Adult Skills. Technical report, OECD Publishing. doi: 10.1787/9789264204256-en

Ogden, C. K. (1937). Basic English: A General Introduction with Rules and Grammar. London: Paul Treber.

Orăsan, C., Evans, R., and Mitkov, R. (2018). Intelligent Text Processing to Help Readers with Autism. Cham: Springer International Publishing. doi: 10.1007/978-3-319-67056-0_33

Paetzold, G., and Specia, L. (2015). "LEXenstein: a framework for lexical simplification," in *Proceedings of the ACL-IJCNLP 2015 System Demonstrations* (Beijing), 85–90. doi: 10.3115/v1/P15-4015

Paetzold, G., and Specia, L. (2016a). "Benchmarking lexical simplification systems," in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)* (Portorož), 3074–3080.

Paetzold, G., and Specia, L. (2016b). "SemEval 2016 task 11: complex word identification," in *Proceedings of the 10th International Workshop on Semantic Evaluation, SEMEVAL* (San Diego, CA: Association for Computational Linguistics), 560–569. doi: 10.18653/v1/S16-1085

Paetzold, G., and Specia, L. (2017a). "Lexical simplification with neural ranking," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (Valencia: Association for Computational Linguistics), 34–40. doi: 10.18653/v1/E17-2006

Paetzold, G. H., and Specia, L. (2016c). "Unsupervised lexical simplification for non-native speakers," in *Proceedings of the 30th AAAI* (Phoenix, AZ). doi: 10.1609/aaai.v30i1.9885 Paetzold, G. H., and Specia, L. (2017b). A survey on lexical simplification. J. Artif. Intell. Res. 60, 549–593. doi: 10.1613/jair.5526

PlainLanguage (2011). Federal Plain Language Guidelines.

Przybyła, P., and Shardlow, M. (2020). "Multi-word lexical simplification," in *Proceedings of the 28th International Conference on Computational Linguistics* (Barcelona: International Committee on Computational Linguistics), 1435–1446. doi: 10.18653/v1/2020.coling-main.123

Qiang, J., Li, Y., Yi, Z., Yuan, Y., and Wu, X. (2020a). "Lexical simplification with pretrained encoders," in *Thirty-Fourth AAAI Conference on Artificial Intelligence* (New York, NY), 8649–8656. doi: 10.1609/aaai.v34i05.6389

Qiang, J., Li, Y., Zhu, Y., Yuan, Y., and Wu, X. (2020b). LSBert: a simple framework for lexical simplification. *arXiv preprint arXiv:2006.14939*. doi: 10.48550/arXiv.2006.14939

Qiang, J., Lu, X., Li, Y., Yuan, Y., and Wu, X. (2021). Chinese lexical simplification. *IEEE/ACM Trans. Audio Speech Lang. Process.* 29, 1819–1828. doi: 10.1109/TASLP.2021.3078361

Rello, L. (2014). DysWebxia. A text accessibility model for people with dyslexia (Ph.D. thesis). Universitat Pompeu Fabra, Barcelona, Spain.

Rello, L., Baeza-Yates, R. A., Dempere-Marco, L., and Saggion, H. (2013). "Frequent words improve readability and short words improve understandability for people with dyslexia," in *Proceedings of the International Conference on Human-Computer Interaction (Part IV), INTERACT* (Cape Town), 203–219. doi: 10.1007/978-3-642-40498-6_15

Rolin, E., Langlois, Q., Watrin, P., and François, T. (2021). "FrenLyS: a tool for the automatic simplification of French general language texts," in *Proceedings of the International Conference on Recent Advances in Natural Language Processing* (RANLP 2021), 1196–1205. doi: 10.26615/978-954-452-072-4_135

Saggion, H. (2017). Automatic Text Simplification. Morgan & Claypool Publishers. doi: 10.1007/978-3-031-02166-4

Shardlow, M., Evans, R., Paetzold, G. H., and Zampieri, M. (2021). "SemEval-2021 task 1: lexical complexity prediction," in *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)* (Association for Computational Linguistics), 1–16. doi: 10.18653/v1/2021.semeval-1.1

Shardlow, M., and Nawaz, R. (2019). "Neural text simplification of clinical letters with a domain specific phrase table," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (Florence: Association for Computational Linguistics), 380–389. doi: 10.18653/v1/P19-1037

Souza, F., Nogueira, R., and Lotufo, R. (2020). "BERTimbau: pretrained BERT models for Brazilian Portuguese," in *Intelligent Systems*, R. Cerri and R. C. Prati (Rio Grande: Springer International Publishing), 403–417. doi: 10.1007/978-3-030-61377-8_28

Specia, L. (2010). "Translating from complex to simplified sentences," in Proceedings of the 9th international conference on Computational Processing of the Portuguese Language (PROPOR) (Berlin; Heidelberg: Springer), 30–39. doi: 10.1007/978-3-642-12320-7_5

Specia, L., Jauhar, S. K., and Mihalcea, R. (2012). "SemEval-2012 task 1: English Lexical Simplification," in *Proceedings of the First Joint Conference on Lexical and Computational Semantics, SemEval* (Stroudsburg, PA: Association for Computational Linguistics), 347–355.

Štajner, S. (2021). "Automatic text simplification for social good: progress and challenges," in *Findings of the Association for Computational Linguistics*, 2637–2652.

Temnikova, I. P., Castillo, C., and Vieweg, S. (2015). "Emterms 1.0: a terminological resource for crisis tweets," in *Proceedings of the ISCRAM* (Krystiansand).

Uchida, S., Takada, S., and Arase, Y. (2018). "CEFR-based lexical simplification dataset," in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)* (Miyazaki: European Language Resources Association).

W3C (2008). Web Content Accessibility Guidelines (WCAG) 2.0.

Yimam, S. M., Biemann, C., Malmasi, S., Paetzold, G., Specia, L., Štajner, S., et al. (2018). "A report on the complex word identification shared task 2018," in *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications* (New Orleans, LA: Association for Computational Linguistics), 66–78. doi: 10.18653/v1/W18-0507

Yimam, S. M., Štajner, S., Riedl, M., and Biemann, C. (2017). "CWIG3G2 - complex word identification task across three text genres and two user groups," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing* (Taipei: Asian Federation of Natural Language Processing), 401-407.