

SemEval-2021 Task 1: Lexical Complexity Prediction

Matthew Shardlow¹, Richard Evans², Gustavo Henrique Paetzold³, Marcos Zampieri⁴

¹Manchester Metropolitan University, UK

²University of Wolverhampton, UK

³Universidade Tecnológica Federal do Paraná, Brazil

⁴Rochester Institute of Technology USA

m.shardlow@mmu.ac.uk

Abstract

This paper presents the results and main findings of SemEval-2021 Task 1 - Lexical Complexity Prediction. We provided participants with an augmented version of the CompLex Corpus (Shardlow et al., 2020). CompLex is an English multi-domain corpus in which words and multi-word expressions (MWEs) were annotated with respect to their complexity using a five point Likert scale. SemEval-2021 Task 1 featured two Sub-tasks: Sub-task 1 focused on single words and Sub-task 2 focused on MWEs. The competition attracted 198 teams in total, of which 54 teams submitted official runs on the test data to Sub-task 1 and 37 to Sub-task 2.

1 Introduction

The occurrence of an unknown word in a sentence can adversely affect its comprehension by readers. Either they give up, misinterpret, or plough on without understanding. A committed reader may take the time to look up a word and expand their vocabulary, but even in this case they must leave the text, undermining their concentration. The natural language processing solution is to identify candidate words in a text that may be too difficult for a reader (Shardlow, 2013; Paetzold and Specia, 2016a). Each potential word is assigned a judgment by a system to determine if it was deemed ‘complex’ or not. These scores indicate which words are likely to cause problems for a reader. The words that are identified as problematic can be the subject of numerous types of intervention, such as direct replacement in the setting of lexical simplification (Gooding and Kochmar, 2019), or extra information being given in the context of explanation generation (Rello et al., 2015).

Whereas previous solutions to this task have typically considered the Complex Word Identification (CWI) task (Paetzold and Specia, 2016a; Yimam et al., 2018) in which a binary judgment of a word’s

complexity is given (i.e., is a word complex or not?), we instead focus on the Lexical Complexity Prediction (LCP) task (Shardlow et al., 2020) in which a value is assigned from a continuous scale to identify a word’s complexity (i.e., how complex is this word?). We ask multiple annotators to give a judgment on each instance in our corpus and take the average prediction as our complexity label. The former task (CWI) forces each user to make a subjective judgment about the nature of the word that models their personal vocabulary. Many factors may affect the annotator’s judgment including their education level, first language, specialism or familiarity with the text at hand. The annotators may also disagree on the level of difficulty at which to label a word as complex. One annotator may label every word they feel is above average difficulty, another may label words that they feel unfamiliar with, but understand from the context, whereas another annotator may only label those words that they find totally incomprehensible, even in context. Our introduction of the LCP task seeks to address this annotator confusion by giving annotators a Likert scale to provide their judgments. Whilst annotators must still give a subjective judgment depending on their own understanding, familiarity and vocabulary — they do so in a way that better captures the meaning behind each judgment they have given. By aggregating these judgments we have developed a dataset that contains continuous labels in the range of 0–1 for each instance. This means that rather than a system predicting whether a word is complex or not (0 or 1), instead a system must now predict where, on our continuous scale, a word falls (0–1).

Consider the following sentence taken from a biomedical source, where the target word ‘observation’ has been highlighted:

- (1) The **observation** of unequal expression leads to a number of questions.

In the binary annotation setting of CWI some annotators may rightly consider this term non-complex, whereas others may rightly consider it to be complex. Whilst the meaning of the word is reasonably clear to someone with scientific training, the context in which it is used is unfamiliar for a lay reader and will likely lead to them considering it complex. In our new LCP setting, we are able to ask annotators to mark the word on a scale from very easy to very difficult. Each user can give their subjective interpretation on this scale indicating how difficult they found the word. Whilst annotators will inevitably disagree (some finding it more or less difficult), this is captured and quantified as part of our annotations, with a word of this type likely to lead to a medium complexity value.

LCP is useful as part of the wider task of lexical simplification (Devlin and Tait, 1998), where it can be used to both identify candidate words for simplification (Shardlow, 2013) and rank potential words as replacements (Paetzold and Specia, 2017). LCP is also relevant to the field of readability assessment, where knowing the proportion of complex words in a text helps to identify the overall complexity of the text (Dale and Chall., 1948).

This paper presents SemEval-2021 Task 1: Lexical Complexity Prediction. In this task we developed a new dataset for complexity prediction based on the previously published CompLex dataset. Our dataset covers 10,800 instances spanning 3 genres and containing unigrams and bigrams as targets for complexity prediction. We solicited participants in our task and released a trial, training and test split in accordance with the SemEval schedule. We accepted submissions in two separate Sub-tasks, the first being single words only and the second taking single words and multi-word expressions (modelled by our bigrams). In total 55 teams participated across the two Sub-tasks.

The rest of this paper is structured as follows: In Section 2 we discuss the previous two iterations of the CWI task. In Section 3, we present the CompLex 2.0 dataset that we have used for our task, including the methodology we used to produce trial, test and training splits. In Section 5, we show the results of the participating systems and compare the features that were used by each system. We finally discuss the nature of LCP in Section 7 and give concluding remarks in Section 8

2 Related Tasks

CWI 2016 at SemEval The CWI shared task was organized at SemEval 2016 (Paetzold and Specia, 2016a). The CWI 2016 organizers introduced a new CWI dataset and reported the results of 42 CWI systems developed by 21 teams. Words in their dataset were considered complex if they were difficult to understand for non-native English speakers according to a binary labelling protocol. A word was considered complex if at least one of the annotators found it to be difficult. The training dataset consisted of 2,237 instances, each labelled by 20 annotators and the test dataset had 88,221 instances, each labelled by 1 annotator (Paetzold and Specia, 2016a).

The participating systems leveraged lexical features (Choubey and Pateria, 2016; Bingel et al., 2016; Quijada and Medero, 2016) and word embeddings (Kuru, 2016; S.P et al., 2016; Gillin, 2016), as well as finding that frequency features, such as those taken from Wikipedia (Konkol, 2016; Wróbel, 2016) were useful. Systems used binary classifiers such as SVMs (Kuru, 2016; S.P et al., 2016; Choubey and Pateria, 2016), Decision Trees (Choubey and Pateria, 2016; Quijada and Medero, 2016; Malmasi et al., 2016), Random Forests (Ronzano et al., 2016; Brooke et al., 2016; Zampieri et al., 2016; Mukherjee et al., 2016) and threshold-based metrics (Kauchak, 2016; Wróbel, 2016) to predict the complexity labels. The winning system made use of threshold-based methods and features extracted from Simple Wikipedia (Paetzold and Specia, 2016b).

A post-competition analysis (Zampieri et al., 2017) with oracle and ensemble methods showed that most systems performed poorly due mostly to the way in which the data was annotated and the the small size of the training dataset.

CWI 2018 at BEA The second CWI Shared Task was organized at the BEA workshop 2018 (Yimam et al., 2018). Unlike the first task, this second task had two objectives. The first objective was the binary complex or non-complex classification of target words. The second objective was regression or probabilistic classification in which 13 teams were asked to assign the probability of a target word being considered complex by a set of language learners. A major difference in this second task was that datasets of differing genres: (TEXT GENRES) as well as English, German and Spanish datasets

for monolingual speakers and a French dataset for multilingual speakers were provided (Yimam et al., 2018).

Similar to 2016, systems made use of a variety of lexical features including word length (Wani et al., 2018; De Hertog and Tack, 2018; AbuRa’ed and Saggion, 2018; Hartmann and dos Santos, 2018; Alfter and Pilán, 2018; Kajiwara and Komachi, 2018), frequency (De Hertog and Tack, 2018; Aroyehun et al., 2018; Alfter and Pilán, 2018; Kajiwara and Komachi, 2018), N-gram features (Gooding and Kochmar, 2018; Popović, 2018; Hartmann and dos Santos, 2018; Alfter and Pilán, 2018; Butnaru and Ionescu, 2018) and word embeddings (De Hertog and Tack, 2018; AbuRa’ed and Saggion, 2018; Aroyehun et al., 2018; Butnaru and Ionescu, 2018). A variety of classifiers were used ranging from traditional machine learning classifiers (Gooding and Kochmar, 2018; Popović, 2018; AbuRa’ed and Saggion, 2018), to Neural Networks (De Hertog and Tack, 2018; Aroyehun et al., 2018). The winning system made use of Adaboost with WordNet features, POS tags, dependency parsing relations and psycholinguistic features (Gooding and Kochmar, 2018).

3 Data

We previously reported on the annotation of the CompLex dataset (Shardlow et al., 2020) (hereafter referred to as CompLex 1.0), in which we annotated around 10,000 instances for lexical complexity using the Figure Eight platform. The instances spanned three genres: **Europarl**, taken from the proceedings of the European Parliament (Koehn, 2005); **The Bible**, taken from an electronic distribution of the World English Bible translation (Christodouloupoulos and Steedman, 2015) and **Biomedical** literature, taken from the CRAFT corpus (Bada et al., 2012). We limited our annotations to focus only on nouns and multi-word expressions following a Noun-Noun or Adjective-Noun pattern, using the POS tagger from Stanford CoreNLP (Manning et al., 2014) to identify these patterns.

Whilst these annotations allowed us to report on the dataset and to show some trends, the overall quality of the annotations we received was poor and we ended up discarding a large number of the annotations. For CompLex 1.0 we retained only instances with four or more annotations and the low number of annotations (average number of annotators = 7) led to the overall dataset being less

reliable than initially expected

For the Shared Task we chose to boost the number of annotations on the same data as used for CompLex 1.0 using Amazon’s Mechanical Turk platform. We requested a further 10 annotations on each data instance bringing up the average number of annotators per instance. Annotators were presented with the same task layout as in the annotation of CompLex 1.0 and we defined the Likert Scale points as previously:

Very Easy: Words which were very familiar to an annotator.

Easy: Words with which an annotator was aware of the meaning.

Neutral: A word which was neither difficult nor easy.

Difficult: Words which an annotator was unclear of the meaning, but may have been able to infer the meaning from the sentence.

Very Difficult: Words that an annotator had never seen before, or were very unclear.

These annotations were aggregated with the retained annotations of CompLex 1.0 to give our new dataset, CompLex 2.0, covering 10,800 instances across single and multi-words and across 3 genres.

The features that make our corpus distinct from other corpora which focus on the CWI and LCP tasks are described below:

Continuous Annotations: We have annotated our data using a 5-point Likert Scale. Each instance has been annotated multiple times and we have taken the mean average of these annotations as the label for each data instance. To calculate this average we converted the Likert Scale points to a continuous scale as follows: Very Easy → 0, Easy → 0.25, Neutral → 0.5, Difficult → 0.75, Very Difficult → 1.0.

Contextual Annotations: Each instance in the corpus is presented with its enclosing sentence as context. This ensures that the sense of a word can be identified when assigning it a complexity value. Whereas previous work has reannotated the data from the CWI–2018 shared task with word senses (Strohmaier et al., 2020), we do not make explicit sense distinctions between our tokens, instead leaving this task up to participants.

Repeated Token Instances: We provide more than one context for each token (up to a maximum of five contexts per genre). These words were annotated separately during annotation, with the expectation that tokens in different contexts would receive differing complexity values. This deliberately penalises systems that do not take the context of a word into account.

Multi-word Expressions: In our corpus we have provided 1,800 instances of multi-word expressions (split across our 3 sub-corpora). Each MWE is modelled as a Noun-Noun or Adjective-Noun pattern followed by any POS tag which is not a noun. This avoids selecting the first portion of complex noun phrases. There is no guarantee that these will correspond to true MWEs that take on a meaning beyond the sum of their parts, and further investigation into the types of MWEs present in the corpus would be informative.

Aggregated Annotations: By aggregating the Likert scale labels we have generated crowd-sourced complexity labels for each instance in our corpus. We are assuming that, although there is inevitably some noise in any large annotation project (and especially so in crowd-sourcing), this will even out in the averaging process to give a mean value reflecting the appropriate complexity for each instance. By taking the mean average we are assuming unimodal distributions in our annotations.

Varied Genres: We have selected for diverse genres as mentioned above. Previous CWI datasets have focused on informal text such as Wikipedia and multi-genre text, such as news. By focusing on specific texts we force systems to learn generalised complexity annotations that are appropriate in a cross-genre setting.

We have presented summary statistics for CompLex 2.0 in Table 1. In total, 5,617 unique words are split across 10,800 contexts, with an average complexity across our entire dataset of 0.321. Each genre has 3,600 contexts, with each split between 3,000 single words and 600 multi-word expressions. Whereas single words are slightly below the average complexity of the dataset at 0.302, multi-word expressions are much more complex at 0.419,

indicating that annotators found these more difficult to understand. Similarly Europarl and the Bible were less complex than the corpus average, whereas the Biomedical articles were more complex. The number of unique tokens varies from one genre to another as the tokens were selected at random and discarded if there were already more than 5 occurrences of the given token already in the dataset. This stochastic selection process led to a varied dataset with some tokens only having one context, whereas others have as many as five in a given genre. On average each token has around 2 contexts.

4 Data Splits

In order to run the shared task we partitioned our dataset into Trial, Train and Test splits and distributed these according to the SemEval schedule. A criticism of previous CWI shared tasks is that the training data did not accurately reflect the distribution of instances in the testing data. We sought to avoid this by stratifying our selection process for a number of factors. The first factor we considered was genre. We ensured that an even number of instances from each genre was present in each split. We also stratified for complexity, ensuring that each split had a similar distribution of complexities. Finally we also stratified the splits by token, ensuring that multiple instances containing the same token occurred in only one split. This last criterion ensures that systems do not overfit to the test data by learning the complexities of specific tokens in the training data.

Performing a robust stratification of a dataset according to multiple features is a non-trivial optimisation problem. We solved this by first grouping all instances in a genre by token and sorting these groups by the complexity of the least complex instance in the group. For each genre, we passed through this sorted list and for each set of 20 groups we put the first group in the trial set, the next two groups in the test set and the remaining 17 groups in the training data. This allowed us to get a rough 5-85-10 split between trial, training and test data. The trial and training data were released in this ordered format, however to prevent systems from guessing the labels based on the data ordering we randomised the order of the instances in the test data prior to release. The splits that we used for the Shared Task are available via GitHub¹.

¹<https://github.com/MMU-TDMLab/CompLex>

Subset	Genre	Contexts	Unique Tokens	Average Complexity
All	Total	10,800	5,617	0.321
	Europarl	3,600	2,227	0.303
	Biomed	3,600	1,904	0.353
	Bible	3,600	1,934	0.307
Single	Total	9,000	4,129	0.302
	Europarl	3,000	1,725	0.286
	Biomed	3,000	1,388	0.325
	Bible	3,000	1,462	0.293
MWE	Total	1,800	1,488	0.419
	Europarl	600	502	0.388
	Biomed	600	516	0.491
	Bible	600	472	0.377

Table 1: The statistics for CompLex 2.0.

Table 2 presents statistics on each split in our data, where it can be seen that we were able to achieve a roughly even split between genres across the trial, train and test data.

Subset	Genre	Trial	Train	Test
All	Total	520	9179	1101
	Europarl	180	3010	410
	Biomed	168	3090	342
	Bible	172	3079	349
Single	Total	421	7662	917
	Europarl	143	2512	345
	Biomed	135	2576	289
	Bible	143	2574	283
MWE	Total	99	1517	184
	Europarl	37	498	65
	Biomed	33	514	53
	Bible	29	505	66

Table 2: The Trial, Train and Test splits that were used as part of the shared task.

5 Results

The full results of our task can be seen in Appendix A. We had 55 teams participate in our 2 Sub-tasks, with 19 participating in Sub-task 1 only, 1 participating in Sub-task 2 only and 36 participating in both Sub-tasks. We have used Pearson’s correlation for our final ranking of participants, but we have also included other metrics that are appropriate for evaluating continuous and ranked data and provided secondary rankings of these.

Sub-task 1 asked participants to assign complexity values to each of the single words instances in our corpus. For Sub-task 2, we asked participants to submit results on both single words and MWEs. We did not rank participants on MWE-only submis-

sions due to the relatively small number of MWEs in our corpus (184 in the test set).

The metrics we chose for ranking were as follows:

Pearson’s Correlation: We chose this metric as our primary method of ranking as it is well known and understood, especially in the context of evaluating systems with continuous outputs. Pearson’s correlation is robust to changes in scale and measures how the input variables change with each other.

Spearman’s Rank: This metric does not consider the values output by a system, or in the test labels, only the order of those labels. It was chosen as a secondary metric as it is more robust to outliers than Pearson’s correlation.

Mean Absolute Error (MAE): Typically used for the evaluation of regression tasks, we included MAE as it gives an indication of how close the predicted labels were to the gold labels for our task.

Mean Squared Error (MSE): There is little difference in the calculation of MSE vs. MAE, however we also include this metric for completeness.

R2: This measures the proportion of variance of the original labels captured by the predicted labels. It is possible to do well on all the other metrics, yet do poorly on R2 if a system produces annotations with a different distribution than those in the original labels.

In Table 3 we show the scores of the top 10 systems across our 2 Sub-tasks according to Pearson’s

Team	Task 1	
	Pearson	R2
JUST BLUE	0.7886 (1)	0.6172 (2)
DeepBlueAI	0.7882 (2)	0.6210 (1)
Alejandro Mosquera	0.7790 (3)	0.6062 (3)
Andi	0.7782 (4)	0.6036 (4)
CS-UM6P	0.7779 (5)	0.3813 (47)
tuqa	0.7772 (6)	0.5771 (12)
OCHADAI-KYOTO	0.7772 (7)	0.6015 (5)
BigGreen	0.7749 (8)	0.5983 (6)
CSECU-DSC	0.7716 (9)	0.5909 (8)
IA PUCP	0.7704 (10)	0.5929 (7)
<i>Frequency Baseline</i>	0.5287	0.2779
Task 2		
DeepBlueAI	0.8612 (1)	0.7389 (1)
rg_pa	0.8575 (2)	0.7035 (5)
xiang_wen_tian	0.8571 (3)	0.7012 (7)
andi_gpu	0.8543 (4)	0.7055 (4)
ren_wo_xing	0.8541 (5)	0.6967 (8)
Andi	0.8506 (6)	0.7107 (2)
CS-UM6P	0.8489 (7)	0.6380 (17)
OCHADAI-KYOTO	0.8438 (8)	0.7103 (3)
LAST	0.8417 (9)	0.7030 (6)
KFU	0.8406 (10)	0.6967 (9)
<i>Frequency Baseline</i>	0.6571	0.4030

Table 3: The top 10 systems for each task according to Pearson’s correlation. We have also included R2 score to help interpret the former. For full rankings, see Appendix A

Correlation. We have only reported on Pearson’s correlation and R2 in these tables, but the full results with all metrics are available in Appendix A. We have included a Frequency Baseline produced using log-frequency from the Google Web1T and linear regression, which was beaten by the majority of our systems. From these results we can see that systems were able to attain reasonably high scores on our dataset, with the winning systems reporting Pearson’s Correlation of 0.7886 for Sub-task 1 and 0.8612 for Sub-task 2, as well as high R2 scores of 0.6210 for Sub-task 1 and 0.7389 for Sub-task 2. The rankings remained stable across Spearman’s rank, MAE and MSE, with some small variations. Scores were generally higher on Sub-task 2 than on Sub-task 1, and this is likely to be because of the different groups of token-types (single words and MWEs). MWEs are known to be more complex than single words and so this fact may have implicitly helped systems to better model the variance of complexities between the two groups.

6 Participating Systems

In this section we have analysed the participating systems in our task. System Description papers were submitted by 32 teams. In the subsections below, we have first given brief summaries of some of the top systems according to Pearson’s correlation for each task for which we had a description. We then discuss the features used across different systems, as well as the approaches to the task that different teams chose to take. We have prepared a comprehensive table comparing the features and approaches of all systems for which we have the relevant information in Appendix B.

6.1 System Summaries

DeepBlueAI: This system attained the highest Pearson’s Correlation on Sub-task 2 and the second highest Pearson’s Correlation on Sub-task 1. It also attained the highest R2 score across both tasks. The system used an ensemble of pre-trained language models fine-tuned for the task with Pseudo Labelling, Data Augmentation, Stacked Training Models and Multi-Sample Dropout. The data was encoded for the transformer models using the genre and token as a query string and the given context as a supplementary input.

JUST BLUE: This system attained the highest Pearson’s Correlation for Sub-task 1. The system did not participate in Sub-task 2. This system makes use of an ensemble of BERT and RoBERTa. Separate models are fine-tuned for context and token prediction and these are weighted 20-80 respectively. The average of the BERT models and RoBERTa models is taken to give a final score.

RG PA: This system attained the second highest Pearson’s Correlation for Sub-task 2. The system uses a fine-tuned RoBERTa model and boosts the training data for the second task by identifying similar examples from the single-word portion of the dataset to train the multi-word classifier. They use an ensemble of RoBERTa models in their final classification, averaging the outputs to enhance performance.

Alejandro Mosquera: This system attained the third highest Pearson’s Correlation for Sub-task 1. The system used a feature-based approach, incorporating length, frequency, semantic features from WordNet and sentence level readability features. These were passed through a Gradient Boosted Regression.

Andi: This system attained the fourth highest Pearson’s Correlation for Sub-task 1. They combine a traditional feature based approach with features from pre-trained language models. They use psycholinguistic features, as well as GLoVE and Word2Vec Embeddings. They also take features from an ensemble of Language models: BERT, RoBERTa, ELECTRA, ALBERT, DeBERTa. All features are passed through Gradient Boosted Regression to give the final output score.

CS-UM6P: This system attained the fifth highest Pearson’s Correlation for Sub-task 1 and the seventh highest Pearson’s Correlation for Sub-task 2. The system uses BERT and RoBERTa and encodes the context and token for the language models to learn from. Interestingly, whilst this system scored highly for Pearson’s correlation the R2 metric is much lower on both Sub-tasks. This may indicate the presence of significant outliers in the system’s output.

OCHADAI-KYOTO: This system attained the seventh highest Pearson’s Correlation on Sub-task 1 and the eight highest Pearson’s Correlation on Sub-task 2. The system used a fine-tuned BERT and RoBERTa model with the token and context encoded. They employed multiple training strategies to boost performance.

6.2 Approaches

There are three main types of systems that were submitted to our task. In line with the state of the art in modern NLP, these can be categorised as: Feature-based systems, Deep Learning Systems and Systems which use a hybrid of the former two approaches. Although Deep Learning Based systems have attained the highest Pearson’s Correlation on both Sub-tasks, occupying the first two places in each task, Feature based systems are not far behind, attaining the third and fourth spots on Sub-task 1 with a similar score to the top systems. We have described each approach as applied to our task below.

Feature-based systems use a variety of features known to be useful for lexical complexity. In particular, lexical frequency and word length feature heavily with many different ways of calculating these metrics such as looking at various corpora and investigating syllable or morpheme length. Psycholinguistic features which model people’s perception of words are understandably popular for this task as complexity is a perceived phenomenon.

Semantic features taken from WordNet modelling the sense of the word and it’s ambiguity or abstractness have been used widely, as well as sentence level features aiming to model the context around the target words. Some systems chose to identify named entities, as these may be innately more difficult for a reader. Word inclusion lists were also a popular feature, denoting whether a word was found on a given list of easy to read vocabulary. Finally, word embeddings are a popular feature, coming from static resources such as GLoVE or Word2Vec, but also being derived through the use of Transformer models such as BERT, RoBERTa, XLNet or GPT-2, which provide context dependent embeddings suitable for our task.

These features are passed through a regression system, with Gradient Boosted Regression and Random Forest Regression being two popular approaches amongst participants for this task. Both apply scale invariance meaning that less pre-processing of inputs is necessary.

Deep Learning Based systems invariably rely on a pre-trained language model and fine-tune this using transfer learning to attain strong scores on the task. BERT and RoBERTa were used widely in our task, with some participants also opting for ALBERT, ERNIE, or other such language models. To prepare data for these language models, most participants following this approach concatenated the token with the context, separated by a special token ($\langle SEP \rangle$). The Language Model was then trained and the embedding of the $\langle CLS \rangle$ token extracted and passed through a further fine-tuned network for complexity prediction. Adaptations to this methodology include applying training strategies such as adversarial training, multi-task learning, dummy annotation generation and capsule networks.

Finally, hybrid approaches use a mixture of Deep Learning by fine-tuning a neural network alongside feature-based approaches. The features may be concatenated to the input embeddings, or may be concatenated at the output prior to further training. Whilst this strategy appears to be the best of both worlds, uniting linguistic knowledge with the power of pre-trained language models, the hybrid systems do not tend to perform as well as either feature based or deep learning systems.

6.3 MWEs

For Sub-task 2 we asked participants to submit both predictions for single words and multi-words

from our corpus. We hoped this would encourage participants to consider models that adapted single word lexical complexity to multi-word lexical complexity. We observed a number of strategies that participants employed to create the annotations for this secondary portion of our data.

For systems that employed a deep learning approach, it was relatively simple to incorporate MWEs as part of their training procedure. These systems encoded the input using a query and context, separated by a $\langle SEP \rangle$ token. The number of tokens prior to the $\langle SEP \rangle$ token did not matter and either one or two tokens could be placed there to handle single and multi-word instances simultaneously.

However, feature based systems could not employ this trick and needed to devise more imaginative strategies for handling MWEs. Some systems handled them by averaging the features of both tokens in the MWE, or by predicting scores for each token and then averaging these scores. Other systems doubled their feature space for MWEs and trained a new model which took the features of both words into account.

7 Discussion

In this paper we have posited the new task of Lexical Complexity Prediction. This builds on previous work on Complex Word Identification, specifically by providing annotations which are continuous rather than binary or probabilistic as in previous tasks. Additionally, we provided a dataset with annotations in context, covering three diverse genres and incorporating MWEs, as well as single tokens. We have moved towards this task, rather than rerunning another CWI task as the outputs of the models are more useful for a diverse range of follow-on tasks. For example, whereas CWI is particularly useful as a preprocessing step for Lexical simplification (identifying which words should be transformed), LCP may also be useful for readability assessment or as a rich feature in other downstream NLP tasks. A continuous annotation allows a ranking to be given over words, rather than binary categories, meaning that we can not only tell whether a word is likely to be difficult for a reader, but also how difficult that word is likely to be. If a system requires binary complexity (as in the case of lexical simplification) it is easy to transform our continuous complexity values into a binary value by placing a threshold on the complex-

ity scale. The value of the threshold to be selected will likely depend on the target audience, with more competent speakers requiring a higher threshold. When selecting a threshold, the categories we used for annotation should be taken into account, so for example a threshold of 0.5 would indicate all words that were rated as neutral or above.

To create our annotated dataset, we employed crowdsourcing with a Likert scale and aggregated the categorical judgments on this scale to give a continuous annotation. It should be noted that this is not the same as giving a truly continuous judgment (i.e., asking each annotator to give a value between 0 and 1). We selected this protocol as the Likert Scale is familiar to annotators and allows them to select according to defined points (we provided the definitions given earlier at annotation time). The annotation points that we gave were not intended to give an even distribution of annotations and it was our expectation that most words would be familiar to some degree, falling in the very easy or easy categories. We pre-selected for harder words to ensure that there were also words in the difficult and very difficult categories. As such, the corpus we have presented is not designed to be representative of the distribution of words across the English language. To create such a corpus, one would need to annotate all words according to our scale with no filtering. The general distribution of annotations in our corpus is towards the easier end of the Likert scale.

A criticism of the approach we have employed is that it allows for subjectivity in the annotation process. Certainly one annotator’s perception of complexity will be different to another’s. Giving fixed values of complexity for every word will not reflect the specific difficulties that one reader, or one reader group will face. The annotations we have provided are averaged values of the annotations given by our annotators, we chose to keep all instances, rather than filtering out those where annotators gave a wide spread of complexity annotations. Further work may be undertaken to give interesting insights into the nature of subjectivity in annotations. For example, some words may be rated as easy or difficult by all annotators, whereas others may receive both easy and difficult annotations, indicating that the perceived complexity of the instance is more subjective. We did not make the individual annotations available as part of the shared task data, to encourage systems to focus

primarily on the prediction of complexity.

An issue with the previous shared tasks is that scores were typically low and that systems tended to struggle to beat reasonable baselines, such as those based on lexical frequency. We were pleased to see that systems participating in our task returned scores that indicated that they had learnt to model the problem well (Pearson’s Correlation of 0.7886 on Task 1 and 0.8612 on Task 2). MWEs are typically more complex than single words and it may be the case that these exhibited a lower variance, and were thus easier to predict for the systems. The strong Pearson’s Correlation is backed up by a high R2 score (0.6172 for Task 1 and 0.7389 for Task 2), which indicates that the variance in the data is captured accurately by the models’ predictions. These models strongly outperformed a reasonable baseline based on word frequency as shown in Table 3.

Whilst we have chosen in this report to rank systems based on their score on Pearson’s correlation, giving a final ranking over all systems, it should be noted that there is very little variation in score between the top systems and all other systems. For Task 1 there are 0.0182 points of Pearson’s Correlation separating the systems at ranks 1 and 10. For Task 2 a similar difference of 0.021 points of Pearson’s Correlation separates the systems at ranks 1 and 10. These are small differences and it may be the case that had we selected a different random split in our dataset this would have led to a different ordering in our results (Gorman and Bedrick, 2019; Søgaard et al., 2020). This is not unique to our task and is something for the SemEval community to ruminate on as the focus of NLP tasks continues to move towards better evaluation rather than better systems.

An analysis of the systems that participated in our task showed that there was little variation between Deep Learning approaches and Feature Based approaches, although Deep Learning approaches ultimately attained the highest scores on our data. Generally the Deep Learning and Feature Based approaches are interleaved in our results table, showing that both approaches are still relevant for LCP. One factor that did appear to affect system output was the inclusion of context, whether that was in a deep learning setting or a feature based setting. Systems which reported using no context appeared to perform worse in the overall rankings. Another feature that may have helped performance

is the inclusion of previous CWI datasets (Yimam et al., 2017; Maddela and Xu, 2018). We were aware of these when developing the corpus and attempted to make our data sufficiently distinct in style to prevent direct reuse of these resources.

A limitation of our task is that it focuses solely on LCP for the English Language. Previous CWI shared tasks (Yimam et al., 2018) and simplification efforts (Saggion et al., 2015; Aluísio and Gasperin, 2010) have focused on languages other than English and we hope to extend this task in the future to other languages.

8 Conclusion

We have presented the SemEval-2021 Task 1 on Lexical Complexity Prediction. We developed a new dataset focusing on continuous annotations in context across three genres. We solicited participants via SemEval and 55 teams submitted results across our two Sub-tasks. We have shown the results of these systems and discussed the factors that helped systems to perform well. We have analysed all the systems that participated and categorised their findings to help future researchers understand which approaches are suitable for LCP.

References

- Ahmed AbuRa’ed and Horacio Saggion. 2018. LaS-TUS/TALN at Complex Word Identification (CWI) 2018 Shared Task. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- David Alfter and Ildikó Pilán. 2018. SB@GU at the Complex Word Identification 2018 Shared Task. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Sandra Aluísio and Caroline Gasperin. 2010. *Fostering digital inclusion and accessibility: The PorSimple project for simplification of Portuguese texts*. In *Proceedings of the NAACL HLT 2010 Young Investigators Workshop on Computational Approaches to Languages of the Americas*, pages 46–53, Los Angeles, California. Association for Computational Linguistics.
- Segun Taofeek Aroyehun, Jason Angel, Daniel Alejandro Pérez Alvarez, and Alexander Gelbukh. 2018. Complex Word Identification: Convolutional Neural Network vs. Feature Engineering . In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans,

- United States. Association for Computational Linguistics.
- Abdul Aziz, MD, Akram Hossain, and Abu Nowshed Chy. 2021. CSECU-DSG at SemEval-2021 Task 1: Fusion of Transformer Models for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Michael Bada, Miriam Eckert, Donald Evans, Kristin Garcia, Krista Shipley, Dmitry Sitnikov, William A Baumgartner, K Bretonnel Cohen, Karin Verspoor, Judith A Blake, et al. 2012. Concept annotation in the craft corpus. *BMC bioinformatics*, 13(1):1–20.
- Yves Bestgen. 2021. LAST at SemEval-2021 Task 1: Improving Multi-Word Complexity Prediction Using Bigram Association Measures. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Joachim Bingel, Natalie Schluter, and Héctor Martínez Alonso. 2016. CoastalCPH at SemEval-2016 Task 11: The importance of designing your Neural Networks right. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1028–1033, San Diego, California. Association for Computational Linguistics.
- Julian Brooke, Alexandra Uitdenbogerd, and Timothy Baldwin. 2016. Melbourne at SemEval 2016 Task 11: Classifying Type-level Word Complexity using Random Forests with Corpus and Word List Features. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 975–981, San Diego, California. Association for Computational Linguistics.
- Andrei Butnaru and Radu Tudor Ionescu. 2018. UnibucKernel: A kernel-based learning method for complex word identification. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Prafulla Choubey and Shubham Pateria. 2016. Garuda & Bhasha at SemEval-2016 Task 11: Complex Word Identification Using Aggregated Learning Models. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1006–1010, San Diego, California. Association for Computational Linguistics.
- Christos Christodouloupoulos and Mark Steedman. 2015. A massively parallel corpus: the bible in 100 languages. *Language resources and evaluation*, 49(2):375–395.
- E. Dale and J. S. Chall. 1948. A formula for predicting readability. *Educational research bulletin*, 27.
- Dirk De Hertog and Anaïs Tack. 2018. Deep Learning Architecture for Complex Word Identification. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Abhinandan Desai, Kai North, Marcos Zampieri, and Christopher Homan. 2021. LCP-RIT at SemEval-2021 Task 1: Exploring Linguistic Features for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- S. Devlin and J. Tait. 1998. The use of a psycholinguistic database in the simplification of text for aphasic readers. In *Linguistic Databases*. Stanford, USA: CSLI Publications.
- Robert Flynn and Matthew Shardlow. 2021. Manchester Metropolitan at SemEval-2021 Task 1: Convolutional Networks for Complex Word Identification. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Nat Gillin. 2016. Sensible at SemEval-2016 Task 11: Neural Nonsense Mangled in Ensemble Mess. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 963–968, San Diego, California. Association for Computational Linguistics.
- Sebastian Gombert and Sabine Bartsch. 2021. TUDA-CCL at SemEval-2021 Task 1: Using Gradient-boosted Regression Tree Ensembles Trained on a Heterogeneous Feature Set for Predicting Lexical Complexity. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Sian Gooding and Ekaterina Kochmar. 2018. CAMB at CWI Shared Task 2018: Complex Word Identification with Ensemble-Based Voting. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Sian Gooding and Ekaterina Kochmar. 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)*, pages 4855–4865.
- Kyle Gorman and Steven Bedrick. 2019. We need to talk about standard splits. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.
- Nathan Hartmann and Leandro Borges dos Santos. 2018. NILC at CWI 2018: Exploring Feature Engineering and Feature Learning. In *Proceedings of the*

- 13th Workshop on Innovative Use of NLP for Building Educational Applications, New Orleans, United States. Association for Computational Linguistics.
- Bo Huang, Yang Bai, and Xiaobing Zhou. 2021. hub at SemEval-2021 Task 1: Fusion of Sentence and Word Frequency to Predict Lexical Complexity. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Aadil Islam, Weicheng Ma, and Soroush Vosoughi. 2021. BigGreen at SemEval-2021 Task 1: Lexical Complexity Prediction with Assembly Models. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Tomoyuki Kajiwara and Mamoru Komachi. 2018. Complex Word Identification Based on Frequency in a Learner Corpus. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- David Kauchak. 2016. Pomona at SemEval-2016 Task 11: Predicting Word Complexity Based on Corpus Frequency. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1047–1051, San Diego, California. Association for Computational Linguistics.
- Milton King, Ali Hakimi Parizi, Samin Fakharian, and Paul Cook. 2021. UNBNLP at SemEval-2021 Task 1: Predicting lexical complexity with masked language models and character-level encoders. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86. Citeseer.
- Michal Konkol. 2016. UWB at SemEval-2016 Task 11: Exploring Features for Complex Word Identification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1038–1041, San Diego, California. Association for Computational Linguistics.
- Onur Kuru. 2016. AI-KU at SemEval-2016 Task 11: Word Embeddings and Substring Features for Complex Word Identification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1042–1046, San Diego, California. Association for Computational Linguistics.
- Chaya Liebeskind, Otniel Elkayam, and Shmuel Liebeskind. 2021. JCT at SemEval-2021 Task 1: Context-aware Representation for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Mounica Maddela and Wei Xu. 2018. A word-complexity lexicon and a neural readability ranking model for lexical simplification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3749–3760.
- Shervin Malmasi, Mark Dras, and Marcos Zampieri. 2016. LTG at SemEval-2016 Task 11: Complex Word Identification with Classifier Ensembles. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 996–1000, San Diego, California. Association for Computational Linguistics.
- Nabil El Mamoun, Abdelkader El Mahdaouy, Abdelah El Mekki, Kabil Essefar, and Ismail Berrada. 2021. CS-UM6P at SemEval-2021 Task 1: A Deep Learning Model-based Pre-trained Transformer Encoder for Lexical Complexity. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Alejandro Mosquera. 2021. Alejandro Mosquera at SemEval-2021 Task 1: Exploring Sentence and Word Features for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Niloy Mukherjee, Braja Gopal Patra, Dipankar Das, and Sivaji Bandyopadhyay. 2016. JU_NLP at SemEval-2016 Task 11: Identifying Complex Words in a Sentence. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 986–990, San Diego, California. Association for Computational Linguistics.
- Gustavo Paetzold and Lucia Specia. 2016a. SemEval 2016 Task 11: Complex Word Identification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 560–569, San Diego, California. Association for Computational Linguistics.
- Gustavo Paetzold and Lucia Specia. 2016b. SV000gg at SemEval-2016 Task 11: Heavy Gauge Complex Word Identification with System Voting. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 969–974, San Diego, California. Association for Computational Linguistics.
- Gustavo Paetzold and Lucia Specia. 2017. 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*, pages 34–40, Valencia, Spain. Association for Computational Linguistics.
- Gustavo Henrique Paetzold. 2021. UTFPR at SemEval-2021 Task 1: Complexity Prediction by Combining BERT Vectors and Classic Features. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Chunguang Pan, Bingyan Song, Shengguang Wang, and Zhipeng Luo. 2021. DeepBlueAI at SemEval-2021 Task 1: Lexical Complexity Prediction with A Deep Ensemble Approach. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Maja Popović. 2018. Complex Word Identification using Character n-grams. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Maury Quijada and Julie Medero. 2016. [HMC at SemEval-2016 Task 11: Identifying Complex Words Using Depth-limited Decision Trees](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1034–1037, San Diego, California. Association for Computational Linguistics.
- Gang Rao, Maochang Li, Xiaolong Hou, Lianxin Jiang, Yang Mo, and Jianping Shen. 2021. RG PA at SemEval-2021 Task 1: A Contextual Attention-based Model with RoBERTa for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Luz Rello, Roberto Carlini, Ricardo Baeza-Yates, and Jeffrey P Bigham. 2015. A plug-in to aid online reading in spanish. In *Proceedings of the 12th Web for All Conference*, pages 1–4.
- Kervy Rivas Rojas and Fernando Alva-Manchego. 2021. IAPUCP at SemEval-2021 Task 1: Stacking Fine-Tuned Transformers is Almost All You Need for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Francesco Ronzano, Ahmed Abura’ed, Luis Espinosa Anke, and Horacio Saggion. 2016. [TALN at SemEval-2016 Task 11: Modelling Complex Words by Contextual, Lexical and Semantic Features](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1011–1016, San Diego, California. Association for Computational Linguistics.
- Armand Rotaru. 2021. ANDI at SemEval-2021 Task 1: Predicting complexity in context using distributional models, behavioural norms, and lexical resources. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Erik Rozi, Niveditha Iyer, Gordon Chi, Enok Choe, Kathy Lee, Kevin Liu, Patrick Liu, Zander Lack, Jillian Tang, and Ethan A. Chi. 2021. Stanford MLab at SemEval-2021 Task 1: Tree-Based Modelling of Lexical Complexity using Word Embeddings. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Irene Russo. 2021. archer at SemEval-2021 Task 1: Contextualising Lexical Complexity. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Horacio Saggion, Sanja Štajner, Stefan Bott, Simon Mille, Luz Rello, and Biljana Drndarevic. 2015. Making it simplext: Implementation and evaluation of a text simplification system for spanish. *ACM Transactions on Accessible Computing (TACCESS)*, 6(4):1–36.
- Matthew Shardlow. 2013. [A comparison of techniques to automatically identify complex words](#). In *51st Annual Meeting of the Association for Computational Linguistics Proceedings of the Student Research Workshop*, pages 103–109, Sofia, Bulgaria. Association for Computational Linguistics.
- Matthew Shardlow, Michael Cooper, and Marcos Zampieri. 2020. [CompLex — a new corpus for lexical complexity prediction from Likert Scale data](#). In *Proceedings of the 1st Workshop on Tools and Resources to Empower People with READING Difficulties (READI)*, pages 57–62, Marseille, France. European Language Resources Association.
- Neil Shirude, Sagnik Mukherjee, Tushar Shandhilya, Ananta Mukherjee, and Ashutosh Modi. 2021. IITK@LCP at SemEval-2021 Task 1: Classification for Lexical Complexity Regression Task. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Greta Smolenska, Peter Kolb, Sinan Tang, Mironas Bitinis, Héctor Hernández, and Elin Asklöv. 2021. CLULEX at SemEval-2021 Task 1: A Simple System Goes a Long Way. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Anders Søgaard, Sebastian Ebert, Joost Bastings, and Katja Filippova. 2020. We need to talk about random splits. *arXiv preprint arXiv:2005.00636*.
- Sanjay S.P, Anand Kumar M, and Soman K P. 2016. [AmritaCEN at SemEval-2016 Task 11: Complex Word Identification using Word Embedding](#). In

- Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1022–1027, San Diego, California. Association for Computational Linguistics.
- Regina Stodden and Gayatri Venugopal. 2021. RS.GV at SemEval-2021 Task 1: Sense Relative Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- David Strohmaier, Sian Gooding, Shiva Taslimipoor, and Ekaterina Kochmar. 2020. **SeCoDa: Sense complexity dataset**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5962–5967, Marseille, France. European Language Resources Association.
- Yuki Taya, Lis Kanashiro Pereira, Fei Cheng, and Ichiro Kobayashi. 2021. OCHADAI-KYODAI at SemEval-2021 Task 1: Enhancing Model Generalization and Robustness for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Giuseppe Vettigli and Antonio Sorgente. 2021. CompNA at SemEval-2021 Task 1: Prediction of lexical complexity analyzing heterogeneous features. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Katja Voskoboinik. 2021. katildakat at SemEval-2021 Task 1: Lexical Complexity Prediction of Single Words and Multi-Word Expressions in English. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Nikhil Wani, Sandeep Mathias, Jayashree Aanand Gajjam, and Pushpak Bhattacharyya. 2018. The Whole is Greater than the Sum of its Parts: Towards the Effectiveness of Voting Ensemble Classifiers for Complex Word Identification. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Krzysztof Wróbel. 2016. **PLUJAGH at SemEval-2016 Task 11: Simple System for Complex Word Identification**. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 953–957, San Diego, California. Association for Computational Linguistics.
- Rong Xiang, Jinghang Gu, Emmanuele Chersoni, Wenjie Li, Qin Lu, and Chu-Ren Huang. 2021. PolyU CBS-Comp at SemEval-2021 Task 1: Lexical Complexity Prediction (LCP). In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Tuqa Bani Yaseen, Qusai Ismail, Sarah Al-Omari, Eslam Al-Sobh, and Malak Abdullah. 2021. JUSTBLUE at SemEval-2021 Task 1: Predicting Lexical Complexity using BERT and RoBERTa Pre-trained Language Models. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Seid Muhie Yimam, Chris Biemann, Shervin Malmasi, Gustavo Paetzold, Luci Specia, Sanja Štajner, Anaïs Tack, and Marcos Zampieri. 2018. A Report on the Complex Word Identification Shared Task 2018. In *Proceedings of the 13th Workshop on Innovative Use of NLP for Building Educational Applications*, New Orleans, United States. Association for Computational Linguistics.
- Seid Muhie Yimam, Sanja Štajner, Martin Riedl, and Chris Biemann. 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 (Volume 2: Short Papers)*, pages 401–407, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zheng Yuan, Gladys Tyen, and David Strohmaier. 2021. Cambridge at SemEval-2021 Task 1: An Ensemble of Feature-Based and Neural Models for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- George-Eduard Zaharia, Dumitru-Clementin Cercel, and Mihai Dascalu. 2021. UPB at SemEval-2021 Task 1: Combining Deep Learning and Hand-Crafted Features for Lexical Complexity Prediction. In *Proceedings of the Fifteenth Workshop on Semantic Evaluation*, Bangkok, Thailand (online). Association for Computational Linguistics.
- Marcos Zampieri, Shervin Malmasi, Gustavo Paetzold, and Lucia Specia. 2017. **Complex Word Identification: Challenges in Data Annotation and System Performance**. In *Proceedings of the 4th Workshop on Natural Language Processing Techniques for Educational Applications (NLPTEA 2017)*, pages 59–63, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Marcos Zampieri, Liling Tan, and Josef van Genabith. 2016. **MacSaar at SemEval-2016 Task 11: Zipfian and Character Features for Complex Word Identification**. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1001–1005, San Diego, California. Association for Computational Linguistics.

A Full Results

Rank	Team	Pearson	Spearman	MAE	MSE	R2
1	JUST BLUE	0.7886	0.7369	0.0609	0.0062	0.6172
2	DeepBlueAI	0.7882	0.7425	0.0610	0.0061	0.6210
3	Alejandro Mosquera	0.7790	0.7355	0.0619	0.0064	0.6062
4	Andi	0.7782	0.7287	0.0637	0.0064	0.6036
5	CS-UM6P	0.7779	0.7366	0.0803	0.0100	0.3813
6	tuqa	0.7772	0.7344	0.0635	0.0068	0.5771
7	OCHADAI-KYOTO	0.7772	0.7313	0.0617	0.0065	0.6015
8	BigGreen	0.7749	0.7294	0.0629	0.0065	0.5983
9	CSECU-DSG	0.7716	0.7326	0.0632	0.0066	0.5909
10	ia pucp	0.7704	0.7361	0.0618	0.0066	0.5929
11	CLP	0.7692	0.7336	0.0631	0.0067	0.5854
12	ess	0.7656	0.7308	0.0635	0.0069	0.5747
13	ismail2022	0.7653	0.7245	0.0641	0.0069	0.5766
14	andi_gpu	0.7651	0.7275	0.0629	0.0068	0.5810
15	TUDA-CCL	0.7649	0.7164	0.0643	0.0067	0.5846
16	rg_pa	0.7628	0.7251	0.0634	0.0069	0.5749
17	ren_wo_xing	0.7618	0.7229	0.0639	0.0069	0.5715
18	CLULEX	0.7588	0.7089	0.0649	0.0069	0.5753
19	acccb	0.7586	0.7207	0.0635	0.0069	0.5730
20	jiu_mo_zhi	0.7584	0.7175	0.0635	0.0070	0.5691
21	Eslam93	0.7577	0.7224	0.0640	0.0070	0.5648
22	archer	0.7561	0.7067	0.0641	0.0069	0.5707
23	Cambridge	0.7556	0.7105	0.0646	0.0070	0.5705
24	eee	0.7553	0.7203	0.0673	0.0078	0.5181
25	CompNA	0.7552	0.7153	0.0641	0.0070	0.5701
26	LAST	0.7534	0.6988	0.0652	0.0070	0.5652
27	Stanford MLab	0.7533	0.7044	0.0653	0.0071	0.5615
28	mau_lih	0.7513	0.7263	0.0645	0.0071	0.5587
29	IITK@LCP	0.7511	0.7167	0.0654	0.0071	0.5598
30	cognience	0.7510	0.7193	0.0652	0.0071	0.5625
31	qnamqj	0.7509	0.7086	0.0649	0.0072	0.5536
32	feras1515	0.7503	0.7180	0.0652	0.0073	0.5477
33	eslam	0.7482	0.7237	0.0649	0.0072	0.5525
34	RS_GV	0.7478	0.7077	0.0698	0.0079	0.5144
35	LucasHub	0.7434	0.6995	0.0658	0.0073	0.5486
36	LRL_NC	0.7402	0.7013	0.0661	0.0074	0.5440
37	Manchester Metropolitan	0.7389	0.7135	0.0656	0.0074	0.5398
38	UPB	0.7340	0.6785	0.0699	0.0079	0.5098
39	KFU	0.7201	0.6899	0.0687	0.0079	0.5109
40	PolyU CBS-Comp	0.7188	0.6935	0.0682	0.0078	0.5162
41	LCP_RIT	0.7086	0.6535	0.0716	0.0086	0.4695
42	UNBNLP	0.6953	0.6544	0.0716	0.0089	0.4495
43	chenshi	0.6951	0.6532	0.0740	0.0091	0.4366
44	UTFPR	0.6875	0.6588	0.0735	0.0088	0.4577
45	Katildakat	0.6715	0.6454	0.0756	0.0096	0.4060
46	jct	0.6663	0.6457	0.0736	0.0091	0.4402
47	LECCE	0.6452	0.6405	0.0772	0.0096	0.4046
48	S3003183	0.5834	0.5437	0.0804	0.0110	0.3182
–	<i>Frequency Baseline</i>	0.5287	0.5263	0.0870	0.0136	0.2779
49	C3SL	0.4598	0.3983	0.0866	0.0130	0.1989
50	SINAI	0.4428	0.3961	0.0875	0.0131	0.1930
51	ProjectLIN513	0.3884	0.4316	0.1019	0.0159	0.0198
52	glitterosu	0.1807	0.1516	0.1024	0.0194	-0.2016
53	PyGuajo	0.0971	0.1440	0.1166	0.0338	-1.0861
54	RACAI	-0.0272	-0.0268	0.2777	0.1270	-6.8449

Table 4: Sub-task 1: Results and rank (in brackets) in terms of Pearson, Spearman, MAE, MSE, and R2. The rank corresponds to Pearson.

Rank	Team	Pearson	Spearman	MAE	MSE	R2
1	DeepBlueAI	0.8612	0.8526	0.0616	0.0063	0.7389
2	rg_pa	0.8575	0.8529	0.0672	0.0072	0.7035
3	xiang_wen_tian	0.8571	0.8548	0.0675	0.0072	0.7012
4	andi_gpu	0.8543	0.8448	0.0664	0.0071	0.7055
5	ren_wo_xing	0.8541	0.8473	0.0677	0.0073	0.6967
6	Andi	0.8506	0.8381	0.0667	0.0070	0.7107
7	CS-UM6P	0.8489	0.8406	0.0760	0.0087	0.6380
8	OCHADAI-KYOTO	0.8438	0.8285	0.0660	0.0070	0.7103
9	LAST	0.8417	0.8299	0.0677	0.0072	0.7030
10	KFU	0.8406	0.8337	0.0686	0.0073	0.6967
11	jiu_mo_zhi	0.8355	0.8277	0.0710	0.0083	0.6560
12	CSECU-DSG	0.8311	0.8153	0.0678	0.0077	0.6825
13	accceb	0.8310	0.8157	0.0697	0.0076	0.6850
14	Stanford MLab	0.8280	0.8124	0.0711	0.0080	0.6671
15	IITK@LCP	0.8277	0.8228	0.0811	0.0098	0.5949
16	qnamqj	0.8246	0.8227	0.0787	0.0094	0.6097
17	LRL_NC	0.8244	0.8156	0.0702	0.0079	0.6737
18	mau_lih	0.8234	0.8211	0.0790	0.0096	0.6042
19	TUDA-CCL	0.8190	0.8091	0.0711	0.0080	0.6677
20	Alejandro Mosquera	0.8093	0.8017	0.0731	0.0084	0.6519
21	LucasHub	0.8000	0.7797	0.0754	0.0089	0.6323
22	UPB	0.7962	0.7988	0.0788	0.0099	0.5917
23	CompNA	0.7931	0.7800	0.0783	0.0093	0.6160
24	justglowing	0.7902	0.7851	0.0786	0.0092	0.6169
25	BigGreen	0.7898	0.7769	0.0903	0.0124	0.4858
26	Katildakat	0.7848	0.7869	0.0807	0.0101	0.5816
27	Manchester Metropolitan	0.7611	0.7711	0.0806	0.0102	0.5770
28	UTFPR	0.7601	0.7504	0.0817	0.0102	0.5754
29	UNBNLP	0.7515	0.7420	0.0802	0.0106	0.5623
30	chenshi	0.7500	0.7497	0.0867	0.0112	0.5365
31	PolyU CBS-Comp	0.7416	0.7222	0.0839	0.0109	0.5473
32	cognience	0.7232	0.7301	0.0851	0.0117	0.5144
–	<i>Frequency Baseline</i>	0.6571	0.6345	0.0924	0.0140	0.4030
33	C3SL	0.3941	0.3675	0.1145	0.0206	0.1470
34	PyGuajo	0.3931	0.3902	0.1132	0.0205	0.1488
35	SINAI	0.3197	0.3508	0.1217	0.0243	-0.0062
36	LECCE	0.2821	0.3138	0.1202	0.0226	0.0624
37	glitterosu	0.1860	0.1316	0.1332	0.0255	-0.0564

Table 5: Sub-task 2: Results and rank (in brackets) in terms of Pearson, Spearman, MAE, MSE, and R2. The Rank Corresponds to Pearson.

B System Features

Team	Features	Classification Approach	System Paper
Alejandro Mosquera	Length, Frequency, Semantic, Sentence	Gradient Boosted Regression	(Mosquera, 2021)
Andi	Psycholinguistic, Glove, Word2Vec, ConceptNet NumberBatch, BERT, RoBERTa, ELECTRA, ALBERT, DeBERTa	Ridge Regression, Gradient Boosted Regression	(Rotaru, 2021)
Archer	Length, Frequency, Psycholinguistic, Scrabble Score, Word Inclusion, Semantic	Random Forest Regression, Gradient Boosted Regression	(Russo, 2021)
BigGreen	Length, Semantic, Glove, Elmo, InferSent, Phonetic, Frequency, POS	Gradient Boosted Regression, BERT	(Islam et al., 2021)
C3SL	Sent2Vec	Multi-layer Perceptron	
Cambridge	Frequency, Syntactic, Length	BERT, Random Forest Regression	(Yuan et al., 2021)
CLULEX	Frequency, POS, Named Entities, Word Inclusion, Sentence, Bert	Decision Tree	(Smolenska et al., 2021)
CompNA	Length, Semantic, Glove, Word Inclusion, Token and Context Encoded	Decision Tree Ensemble	(Vettigli and Sorgente, 2021)
CS-UM6P	Token and Context Encoded	BERT, RoBERTa	(Mamoun et al., 2021)
CSECU-DSG	Token and Context Encoded	BERT, RoBERTa	(Aziz et al., 2021)
DeepBlueAI	Token and Context Encoded	BERT, ALBERT, RoBERTa, ERNIE	(Pan et al., 2021)
Hub	TF-IDF, Context Encoded	RoBERTa, Inception	(Huang et al., 2021)
IA PUCP	Sentence, POS, N-gram Frequency, RoBERTa, XLNet, BERT	Gradient Boosted Regression	(Rojas and Alva-Manchego, 2021)
IITK@LCP	Electra + Glove	Linear Regression, Support Vector Machine	(Shirude et al., 2021)
JCT	POS, Frequency, BERT, Cluster Features	Gradient Boosted Regression	(Liebeskind et al., 2021)
JUST BLUE	Token Encoded and Context Encoded	Average of Weighted Bert and Roberta	(Yaseen et al., 2021)
Katildakat	BERT, Length, BERT-score, Frequency, Semantic,	Linear Regression, Multi-layer Perceptron	(Voskoboinik, 2021)
LAST	Frequency, Psycholinguistic, Sentence, Bigram Association	Gradient Boosted Regression	(Bestgen, 2021)
LCP-RIT	Length, Frequency, Character N-Grams, Psycholinguistic, POS	Random forest Regressor	(Desai et al., 2021)
LRL_NC	Frequency, Semantic, Laanguage Model, Psycholinguistic, Word Inclusion	Random forest regressor	
Manchester Metropolitan	Frequency, Psycholinguistic, Length, Embeddings	CNN	(Flynn and Shardlow, 2021)
OCHADAI-KYOTO	Token and Context Encoded	BERT, RoBERTa	(Taya et al., 2021)
PolyU CBS-Comp	Frequency, Length, Capitalisation, POS, Embeddings, BERT, GPT-2	Gradient Boosted Regression	(Xiang et al., 2021)
RG_PA	Context Encoded	RoBERTa	(Rao et al., 2021)
RS.GV	GLoVe, ELMo, BERT, Flair, Readability, Length, Frequency, Semantic, Psycholinguistic, Morphological, Word Inclusion, Named Entity	Feed-Forward Neural Network	(Stodden and Venugopal, 2021)
Stanford MLab	Glove, Length, POS, Named Entity	Gradient Boosted Regression	(Rozi et al., 2021)
TUDA-CCL	Linguistic, Semantic, Embeddings, Psycholinguistic, Frequencies, Word Inclusion	Gradient Boosted Regression	(Gombert and Bartsch, 2021)
UNBNLP	Length, Frequency, Character-Level-Encoder, BERT	Neural Network, Support Vector Machine	(King et al., 2021)
UPB	Transformers, Word Embeddings, Character Embeddings, Length, Psycholinguistic	BERT, RoBERTa, Regression	(Zaharia et al., 2021)
UTFPR	Frequency, Length, Semantic, Bert Embedding	Support Vector Machine	(Paetzold, 2021)

Table 6: Systems that participated and submitted a paper, the features and classification approaches they employed.