

# Pseudo-Labels Are All You Need

**Bogdan Kostić** and **Mathis Lucka** and **Julian Risch**

deepset

{bogdan.kostic, mathis.lucka, julian.risch}@deepset.ai

## Abstract

Automatically estimating the complexity of texts for readers has a variety of applications, such as recommending texts with an appropriate complexity level to language learners or supporting the evaluation of text simplification approaches. In this paper, we present our submission to the Text Complexity DE Challenge 2022, a regression task where the goal is to predict the complexity of a German sentence for German learners at level B. Our approach relies on more than 220,000 pseudo-labels created from the German Wikipedia and other corpora to train Transformer-based models, and refrains from any feature engineering or any additional, labeled data. We find that the pseudo-label-based approach gives impressive results yet requires little to no adjustment to the specific task and therefore could be easily adapted to other domains and tasks.

## 1 Introduction

What makes some texts more difficult to read for learners of a foreign language than others? How does a complicated sentence construction or the use of rare vocabulary increase complexity? The prediction of text complexity with machine learning methods addresses these questions. In contrast to last years' shared tasks at KONVENS, which focused on the disambiguation of German verbal idioms (Ehren et al., 2021), the identification of toxic, engaging, and fact-claiming comments (Risch et al., 2021), and scene segmentation in narrative texts (Zehe et al., 2021), the task of 2022 is about text complexity. In this paper, we present our submission to this Text Complexity DE Challenge 2022. It is a shared task addressing the automatic estimation of the complexity of German sentences for readers, in particular, German learners at level B. The provided training dataset contains about 1000 sentences and the test dataset

about 300 sentences in German. Figure 1 shows an exemplary sentence from the shared task dataset in German, an English translation, and the arithmetic mean of ratings from all annotators. With a seven-level Likert-scale with values ranging from very easy (1) to very complex (7), this task is a regression task and it is evaluated using the Root Mean Squared Error (RMSE). A third-order mapping is applied before the error is measured so that the impact of any systematic bias in the predictions on the metrics is reduced. Thereby, the focus of the evaluation is shifted towards ranking sentences correctly with regards to their complexity rather than assigning the correct absolute complexity score. We refer to the overview paper of the shared task for more details about the dataset and the overall results (Mohtaj et al., 2022).

The remainder of this paper is structured as follows. Section 2 summarizes related work on text complexity estimation and on pseudo-labeling techniques for machine learning. We describe our approach in Section 3 and its evaluation in Section 4, with experiments on the validation dataset provided by the shared task organizers. We conclude in Section 5 and provide an outlook on future work.

## 2 Related Work

Research on reading complexity of German texts is so far relatively scarce with several papers introducing datasets of annotated German sentences or longer texts and mostly feature-based approaches for text complexity prediction. First of all, there is a dataset with sentence-level annotations, which is the basis of this shared task (Naderi et al., 2019). Rios et al. (2021) introduce a dataset for document-level text complexity with the application focus of text simplification and there are two other document-level text complexity datasets by Battisti et al. (2020) and by Hewett and Stede

**German Sentence:** Als Versauerung der Meere wird die Abnahme des pH-Wertes des Meerwassers bezeichnet.

**English Translation:** Ocean acidification is the term used to describe the decrease in the pH of seawater.

**Complexity Score:** 2.13

**German Sentence:** Nach chemischer Härtung des Rußes war er in der Lage, auf galvanoplastischem Wege ein Zink-Positiv und von diesem ein Negativ der Platte anzufertigen, das als Stempel zur Pressung beliebig vieler Positive genutzt werden konnte – die Schallplatte war erfunden.

**English Translation:** After chemical hardening of the carbon black, he was able to produce a zinc positive by galvanoplastic means and from this a negative of the record, which could be used as a stamp for pressing any number of positives - the record was invented.

**Complexity Score:** 4.70

Figure 1: Two sentences from the training dataset.

(2021). The latter follows the format of a similar study (Hulpuş et al., 2019) based on a dataset of English newspaper articles (Xu et al., 2015). Another dataset is from a Kaggle challenge called CommonLit Readability Prize, where the task is to rate the complexity of literary passages for school grades 3-12.<sup>1</sup> Last but not least, there are unlabeled datasets of German texts with simple language, such as the Tagesschau/Logo corpus and the Geo/Geolino corpus (Weiß and Meurers, 2018) or Klexikon (Aumiller and Gertz, 2022). These datasets cannot be used directly for fine-tuning models on the text complexity prediction task due to the lack of annotations. However, we show in our approach that they can be used in combination with pseudo-labeling.

Similar to the pseudo-labeling approach that we use, there is a data augmentation technique where a slow but more accurate cross-encoder model is used to label a large set of otherwise unlabeled data samples (Thakur et al., 2021). This technique augments the training data for a faster, less complex

<sup>1</sup><https://www.kaggle.com/competitions/commonlitreadabilityprize/>

bi-encoder model to address a pairwise sentence ranking task. Du et al. (2021) present a data augmentation method where given a small, labeled training dataset, they retrieve additional training samples from a large unlabeled dataset and then label these samples automatically with a model trained on the original, smaller training dataset. The resulting augmented, synthetic dataset can then be used to train another model that generalizes better to unseen data. Of the related work presented, this approach, also referred to as self-training, is the most similar to the approach we present in this paper. The main difference is that Du et al. (2021) tailor their approach mainly to domain-specific pre-training, whereas we focus on task-specific fine-tuning. Xie et al. (2019) extend the self-training method by intentionally adding noise to the training process to foster better generalization of the trained models. Further, they repeat the self-training process several times, so that the model trained on pseudo-labels is again used to create another set of pseudo-labels, which are in turn used to train another model and so on. As this iterative approach is very resource-intensive in terms of training time, we limit our approach to only one iteration. However, no inherent limitation prevents our approach from more training iterations. To the best of our knowledge, there are no published approaches that use neural language models for the particular task of complexity prediction of German texts, but only of English texts (Martinc et al., 2021).

### 3 A Semi-Supervised Learning Approach for Text Complexity Prediction

Our semi-supervised learning approach uses neural language models based on the Transformer architecture (Vaswani et al., 2017). As pre-trained models that are not fine-tuned to a specific natural language processing task yet, we use GBERT and GELECTRA models by Chan et al. (2020) and an XLM-RoBERTa model by Conneau et al. (2020). Given that the training dataset provided by the shared task organizers is relatively small for fine-tuning these pre-trained models, the core idea of our approach is to increase the number of training samples by automatically generating pseudo-labels. Figure 2 visualizes the different steps of the entire approach with its three main steps: pseudo-labeling, fine-tuning, and ensembling. For the implementation of these steps, we use the two open



Figure 2: Overview of the different steps that comprise the pseudo-labeling, fine-tuning, and ensembling approach.

source frameworks FARM<sup>2</sup> and Haystack<sup>3</sup>.

The first step is to create a large corpus of German sentences with varying text complexity to serve as a source for the pseudo-label sentences. This corpus comprises the following resources:

- eight random subparts of a German Wikipedia dump (about 8 percent of all 2.3 million German Wikipedia articles as of 2019),<sup>4</sup>
- 130,000 news articles from the German news platform Zeit Online,<sup>5</sup>
- three million sentences from German newspaper texts as part of the Leipzig Corpus Collection (Goldhahn et al., 2012),
- the Geo/Geolino/Tagesschau/Logo corpus (Weiß and Meurers, 2018),
- the Corpus Simple German (all subsets except for Klexikon),<sup>6</sup>
- the Klexikon (Aumiller and Gertz, 2022), and
- the Hurraki dictionary for plain language.<sup>7</sup>

Combining these datasets results in a total of 12,955,913 sentences. As some of them appear more than once, this corresponds to 12,562,164 distinct sentences. Each of them is embedded using a SentenceTransformers msmarco-distilbert model<sup>8</sup> (Reimers and Gurevych, 2019) and added to an OpenSearch index. Further, we fine-tune a deepset/gbert-large model on the task of sentence complexity on all of the provided training labels as a baseline. Subsequently, to get our

<sup>2</sup><https://github.com/deepset-ai/farm>

<sup>3</sup><https://github.com/deepset-ai/haystack>

<sup>4</sup>The most recent dump is available online: <https://dumps.wikimedia.org/dewiki/20220720/>

<sup>5</sup><https://www.zeit.de>

<sup>6</sup><https://daniel-jach.github.io/simple-german/simple-german.html>

<sup>7</sup><https://hurraki.de>

<sup>8</sup>[sentence-transformers/msmarco-distilbertmultilingual-en-de-v2-tmp-ling-aligned](https://sentence-transformers.com/multilingual-en-de-v2-tmp-ling-aligned)

set of pseudo-labels, we embed each of the sentences in the provided training set with the same SentenceTransformers msmarco-distilbert model and retrieve the 500 most similar sentences from our large corpus of German sentences as potential pseudo-labels. The baseline complexity scorer model produces a complexity score for each potential pseudo-label. To keep roughly the same distribution as in the original training dataset, we filter the generated pseudo-labels in the following way: we keep only those sentences whose predicted score does not deviate more than the standard deviation of the ratings of the original sentence used to retrieve the 500 potential pseudo-labels. This filtering results in a total of 228,796 pseudo-labels. Table 1 lists the number of pseudo-labels originating from the different data sources.

The pseudo-labels are used to fine-tune different Transformer-based models on the task of complexity scoring. We fine-tune deepset/gelectra-large, deepset/gbert-large and xlm-roberta-large using three different seeds for each model, resulting in a total of nine models. Subsequently, we fine-tune each of these models using five-fold cross-validation with the original training set, resulting in a total of 45 models. Finally, to combine these 45 models into an ensemble providing a single prediction score per data sample, we train linear regression models on the out-of-fold predictions from the previous cross-validations.

## 4 Experiments

We evaluate four different settings using five-fold cross-validation:

- a baseline deepset/gbert-large model fine-tuned on the provided training set,
- an ensemble of nine models fine-tuned only on pseudo-labels and with three different random seeds, scores aggregated by mean (deepset/gbert-large, deepset/gelectra-large, xlm-roberta-large)

Table 1: Number of pseudo-labeled sentences with the average length in characters, and the average mean opinion score per data source. The text complexity of the sources differs with Wikipedia and Hurraki being the most, respectively least difficult.

Data Source	#Sentences	ØLength	ØMOS
GERMAN WIKIPEDIA	137,228	133	3.0
ZEIT ONLINE	47,613	108	2.5
3 MILLION NEWS SENTENCES	25,928	110	2.6
GEO/GEOLINO/TAGESSCHAU/LOGO	7,971	98	2.5
CORPUS SIMPLE GERMAN	4,896	93	2.3
KLEXIKON	3,600	75	2.0
HURRAKI	1,559	43	1.5

- an ensemble of 45 models fine-tuned on pseudo-labels and the provided training set, with scores aggregated by mean, and
- an ensemble of 45 models fine-tuned on pseudo-labels and the provided training set, with scores aggregated by a linear model.

We submitted the predictions of each of the last three settings to the shared task competition.<sup>9</sup>

The first setting serves as our baseline with a `deepset/gbert-large` model fine-tuned on the provided training data. The model is trained on each of the cross-validation folds using early stopping for a maximum of four epochs. Each training run was tracked using MLflow and can be found [here](#).

The second setting is an ensemble of the language models `deepset/gbert-large`, `deepset/gelectra-large` and `xlm-roberta-large`. Each of these models is fine-tuned for two epochs on the pseudo-labels described in Section 3 using three different random seeds. This results in an ensemble of nine models. One training run takes approximately three hours on an NVIDIA Tesla V100 GPU with 16 GB of RAM.

For the third and fourth setting, we fine-tune the resulting models of the previous step on the provided training dataset. To further increase the number of models in the ensemble, we perform five-fold cross-validation on each of the nine models, resulting in a total of 45 models. Again, each training run was tracked using MLflow and can be found [here](#). To ensemble these 45 models, we use two different techniques. The third setting aggregates each individual score into a single score

by simply taking the arithmetic mean of all scores. The fourth setting trains a linear ridge regression model on the out-of-fold predictions for each model that we trained, resulting in five linear regression models. Applying these linear regression models decreases the number of scores from 45 to 5. To get a single score out of these five scores, we calculate their arithmetic mean. Table 2 summarizes the hyperparameters that are used to train the models for the different described settings.

Table 3 lists the cross-validation RMSE on the provided training set. As expected, the approach of using pseudo-labels in combination with ensembling outperforms the simple baseline. We observe that fine-tuning the models only on the pseudo-labels already outperforms the baseline that uses only the original training data. Performance improves further if the models that were fine-tuned on the pseudo-labels are additionally fine-tuned on the original training data. Moreover, using a linear model to aggregate the individual scores instead of using the plain average does not further improve the final score. The best setting, consisting of an ensemble of 45 Transformer models fine-tuned on both the pseudo-labels and the provided training data, with results aggregated using a linear regression model, yields an RMSE of 0.433.

## 5 Conclusion

In this paper, we presented our submission to the Text Complexity DE Challenge 2022. We leveraged pseudo-labeled sentences from Wikipedia and several other publicly available, unlabeled corpora. Based on the labeled training dataset from the shared task and the additional pseudo-labeled data, we fine-tuned Transformer-based neural language models. Our best ensemble model achieved an

<sup>9</sup><https://codalab.lisn.upsaclay.fr/competitions/4964>

Table 2: Hyperparameters for fine-tuning the language models on the pseudo-labels and the provided training data.

Hyperparameter	Fine-Tuning on Pseudo-Labels	Fine-Tuning on Training Set
Learning rate	1e-5	1e-6
LR schedule	linear	linear
Warm-up steps	10%	10%
Batch size	20 for xlm-roberta-large, 32 otherwise	20 for xlm-roberta-large, 32 otherwise
Early stopping (Max.) epochs	✘ 2	✓ 4
Optimizer	Adam	Adam
Max sequence length	128	128

Table 3: Cross-validation RMSE.

Model	1	2	3	4	5	∅
Baseline	0.512	0.460	0.440	0.398	0.488	0.460
Ensemble pseudo-labels only	0.500	0.462	0.381	0.450	0.442	0.447
Ensemble simple mean aggregation	0.491	0.443	0.374	0.443	0.426	0.435
Ensemble linear model aggregation	0.445	0.455	0.405	0.443	0.418	<b>0.433</b>

RMSE of 0.433 in cross-validation on the public dataset without third-order mapping and an RMSE of 0.454 on the private test dataset with third-order mapping (0.484 without third-order mapping). For future work, our trained model could be used to create more pseudo-labels for another iteration of the entire approach, presumably resulting in a model that generalizes even better to unseen test data.

## References

- Dennis Aumiller and Michael Gertz. 2022. [Klexikon: A German Dataset for Joint Summarization and Simplification](#). *arXiv preprint arXiv:2201.07198*.
- Alessia Battisti, Dominik Pfütze, Andreas Säuberli, Marek Kostrzewa, and Sarah Ebling. 2020. [A corpus for automatic readability assessment and text simplification of German](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3302–3311, Marseille, France. European Language Resources Association.
- Branden Chan, Stefan Schweter, and Timo Möller. 2020. [German’s next language model](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6788–6796, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Veselin Stoyanov, and Alexis Conneau. 2021. [Self-training improves pre-training for natural language understanding](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5408–5418, Online. Association for Computational Linguistics.
- Rafael Ehren, Timm Lichte, Jakub Waszczuk, and Laura Kallmeyer. 2021. [Shared Task on the Disambiguation of German Verbal Idioms at KONVENS 2021](#). In *Proceedings of the Shared Task on the Disambiguation of German Verbal Idioms at KONVENS*.
- Dirk Goldhahn, Thomas Eckart, and Uwe Quasthoff. 2012. [Building large monolingual dictionaries at the Leipzig corpora collection: From 100 to 200 languages](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 759–765, Istanbul, Turkey. European Language Resources Association (ELRA).
- Freya Hewett and Manfred Stede. 2021. [Automatically evaluating the conceptual complexity of German texts](#). In *Proceedings of the 17th Conference on Natural Language Processing (KONVENS 2021)*, pages 228–234, Düsseldorf, Germany. KONVENS 2021 Organizers.

- Ioana Hulpuş, Sanja Štajner, and Heiner Stuckenschmidt. 2019. [A spreading activation framework for tracking conceptual complexity of texts](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3878–3887, Florence, Italy. Association for Computational Linguistics.
- Matej Martinc, Senja Pollak, and Marko Robnik-Šikonja. 2021. [Supervised and unsupervised neural approaches to text readability](#). *Computational Linguistics*, 47(1):141–179.
- Salar Mohtaj, Babak Naderi, and Sebastian Möller. 2022. Overview of the GermEval 2022 shared task on text complexity assessment of german text. In *Proceedings of the GermEval 2022 Shared Task on Text Complexity Assessment of German Text*, Potsdam, Germany. Association for Computational Linguistics.
- Babak Naderi, Salar Mohtaj, Kaspar Ensikat, and Sebastian Möller. 2019. [Subjective assessment of text complexity: A dataset for german language](#).
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Annette Rios, Nicolas Spring, Tannon Kew, Marek Kostrzewa, Andreas Säuberli, Mathias Müller, and Sarah Ebling. 2021. [A new dataset and efficient baselines for document-level text simplification in German](#). In *Proceedings of the Third Workshop on New Frontiers in Summarization*, pages 152–161, Online and in Dominican Republic. Association for Computational Linguistics.
- Julian Risch, Anke Stoll, Lena Wilms, and Michael Wiegand. 2021. [Overview of the GermEval 2021 shared task on the identification of toxic, engaging, and fact-claiming comments](#). In *Proceedings of the GermEval 2021 Shared Task on the Identification of Toxic, Engaging, and Fact-Claiming Comments*, pages 1–12, Duesseldorf, Germany. Association for Computational Linguistics.
- Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. 2021. [Augmented SBERT: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 296–310, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zarah Weiß and Detmar Meurers. 2018. [Modeling the readability of German targeting adults and children: An empirically broad analysis and its cross-corpus validation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 303–317, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. 2019. [Self-training with noisy student improves imagenet classification](#).
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. [Problems in Current Text Simplification Research: New Data Can Help](#). *Transactions of the Association for Computational Linguistics*, 3:283–297.
- Albin Zehe, Leonard Konle, Svenja Guhr, Lea Katharina Dümpelmann, Evelyn Gius, Andreas Hotho, Fotis Jannidis, Lucas Kaufmann, Markus Krug, Frank Puppe, et al. 2021. [Shared task on scene segmentation@ konvens 2021](#). In *Proceedings of the Shared Task on Scene Segmentation at KONVENS*, pages 1–21.