

SwissBERT: The Multilingual Language Model for Switzerland

Jannis Vamvas¹ Johannes Graen² Rico Sennrich¹

¹Department of Computational Linguistics, University of Zurich

²Linguistic Research Infrastructure, University of Zurich

johannes.graen@linguistik.uzh.ch,

{vamvas, sennrich}@cl.uzh.ch

Abstract

We present SwissBERT, a masked language model created specifically for processing Switzerland-related text. SwissBERT is a pre-trained model that we adapted to news articles written in the national languages of Switzerland – German, French, Italian, and Romansh. We evaluate SwissBERT on natural language understanding tasks related to Switzerland and find that it tends to outperform previous models on these tasks, especially when processing contemporary news and/or Romansh Grischun. Since SwissBERT uses language adapters, it may be extended to Swiss German dialects in future work. The model and our open-source code are publicly released at <https://github.com/ZurichNLP/swissbert>.

1 Introduction

Self-supervised learning for natural language processing (NLP) has inspired the release of numerous language models, like BERT (Devlin et al., 2019). However, NLP researchers in Switzerland, a country with four national languages, are confronted by a unique language situation. Individual models for German, French or Italian [Chan et al., 2020; Martin et al., 2020; Polignano et al., 2019 etc.] are difficult to combine for multilingual tasks, and massively multilingual models such as XLM-R (Conneau et al., 2020) do not focus on the multilingualism that is particular to Switzerland. The fourth national language, Romansh, is not represented in a neural language model so far.

In this paper, we describe SwissBERT, a model trained on more than 21 million Swiss news articles with a total of 12 billion tokens. By combining articles in Swiss Standard German, French, Italian, and Romansh Grischun, we aim to create multilingual representations by implicitly exploiting common entities and events in the news.

The SwissBERT model is adapted from a *Cross-lingual Modular* (X-MOD) transformer that was

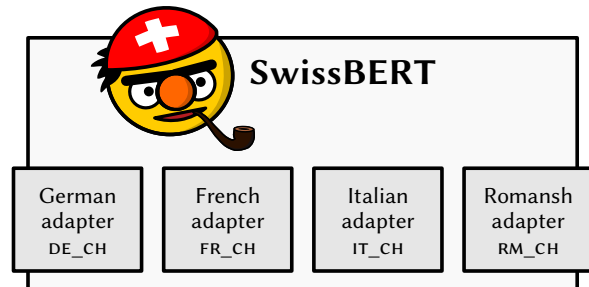


Figure 1: SwissBERT is a transformer encoder with language adapters (Pfeiffer et al., 2022) in each layer. There is an adapter for each national language of Switzerland. The other parameters in the model are shared among the four languages.

pre-trained jointly in 81 languages (Pfeiffer et al., 2022). We adapt X-MOD to our corpus by training custom language adapters. We also create a Switzerland-specific subword vocabulary for SwissBERT. The resulting model has 153M parameters.

Because SwissBERT inherits X-MOD’s modularity, future work may extend it beyond the four national languages. In particular, Swiss German dialects are absent in our training corpus of written news articles but might have other resources that could be used for adding a fifth language adapter to SwissBERT.

In order to evaluate our model, we create a test set for named entity recognition on contemporary news (SwissNER) and find that our model improves over common baselines. When probing our model’s capabilities on Romansh, we find that it strongly outperforms models that have not been trained on the language, both in terms of zero-shot cross-lingual transfer, and German–Romansh alignment of words and sentences (Dolev, 2023).

Since SwissBERT has been adapted to news articles only, we make sure to also gauge its out-of-domain performance. We observe a moderate but systematic improvement over XLM-R when detect-

ing stance in user-generated comments on Swiss politics (Vamvas and Sennrich, 2020) but do not observe state-of-the-art accuracy when recognizing named entities in historical, OCR-processed news (Ehrmann et al., 2022).

We release the SwissBERT model to the research community.¹ Our code repository² includes examples for fine-tuning on downstream tasks based on the *transformers* library (Wolf et al., 2020). Due to the nature of the pre-training corpus, the SwissBERT model may currently not be used for commercial purposes. However, our model may be used in any non-commercial setting, including academic research.³

2 Background and Related Work

Masked Language Models Masked language modeling is a standard approach for learning computational representations from raw text. Masked language models for various languages and domains have been released in the wake of the BERT model (Devlin et al., 2019), a Transformer (Vaswani et al., 2017) that has been trained on English text. For German, such monolingual models have been released by Chan et al. (2020) and Scheible et al. (2020), among others. Similarly, monolingual masked language models have been created for French (Martin et al., 2020; Le et al., 2020), for Italian (Polignano et al., 2019; Muffo and Bertino, 2020) and many other languages. BERT-style models have also been trained on digitized historical newspapers (Schweter, 2020; Schweter et al., 2022).

Multilingual Models Some masked language models have been trained jointly on multiple languages, which allows for transfer learning across languages (Devlin et al., 2019; Conneau and Lample, 2019). While massively multilingual models such as XLM-R enable transfer to languages that have fewer pre-training resources, their overall performance tends to decline compared to monolingual models (Conneau et al., 2020). This trade-off extends to multilingual subword vocabularies that are created jointly for many languages and scripts (Rust et al., 2021).

¹<https://huggingface.co/ZurichNLP/swissbert>

²<https://github.com/ZurichNLP/swissbert>

³A more permissive license might be applied at a later time. We kindly advise you to refer to the code repository for the most recent information.

Cross-lingual Modular Transformers Pfeiffer et al. (2022) have proposed X-MOD, a multilingual model that is similar to XLM-R but has monolingual components. These components are included in each Transformer layer during pre-training. In this paper, we refer to them as *language adapters*, as they are reminiscent of adapters that are added post-hoc to a pre-trained model (Houlsby et al., 2019; Pfeiffer et al., 2020). When fine-tuning X-MOD on a downstream task, the language adapters may be frozen in order to facilitate cross-lingual transfer. Pfeiffer et al. (2022) have shown that their approach better preserves monolingual performance. They have also demonstrated that additional language adapters can be trained after the initial pre-training.

Multilingual Adaptive Pre-training The latter can be seen as an instance of adaptive pre-training, i.e., continuing masked language modeling on a corpus of interest. Alabi et al. (2022) have shown that such adaptation may be performed simultaneously in many languages. In addition to adaptation to new languages, downstream tasks can benefit from adaptive pre-training on specific language varieties (Han and Eisenstein, 2019) or domains (Gururangan et al., 2020). Domain adaptation may be performed with data in multiple languages in order to maintain or improve the multilinguality of the model (Kær Jørgensen et al., 2021).

3 Pre-training Approach

To create a model that is specialized on the Swiss national languages, we build on a massively multilingual X-MOD model.⁴ This model has been pre-trained by Pfeiffer et al. (2022) on filtered web text in 81 languages, including German, French and Italian. Our approach combines three ideas from previous work:

- **Domain adaptation:** We continue training the existing language adapters on a large amount of Swiss news articles.
- **Language adaptation:** We train an adapter for the Romansh language.
- **Multilinguality:** We promote transfer between the four languages by using a joint vocabulary and shared embeddings.

⁴<https://huggingface.co/facebook/xmod-base>

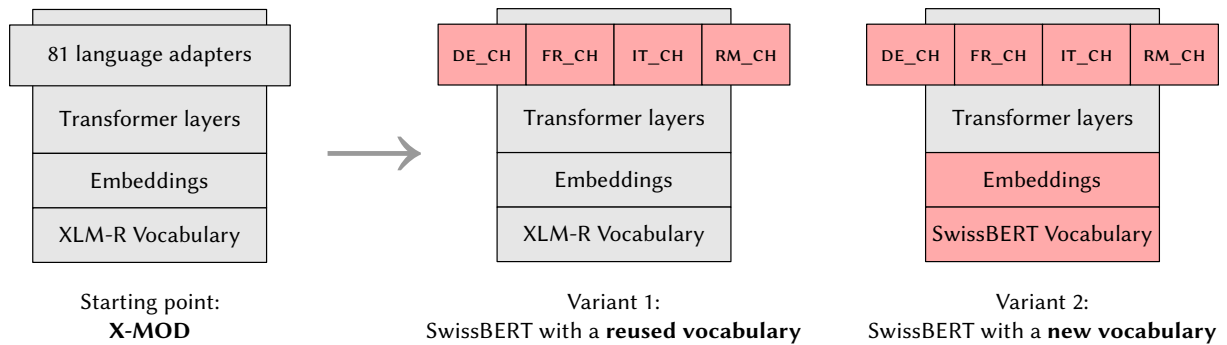


Figure 2: We train two variants of SwissBERT: Variant 1 reuses the vocabulary and embeddings of the pre-trained model, and only language adapters are trained. Variant 2 uses a custom SwissBERT vocabulary based on our pre-training corpus, and multilingual embeddings are trained in addition to the adapters.

3.1 Pre-training Corpus

Our pre-training corpus is composed of media items that have appeared until the end of 2022 and are collected in the `Swissdox@LiRI` database⁵. The large majority of the items are news articles published in print or in online news portals. A small part of the items are related types of documents, such as letters to the editor or transcripts of TV news broadcasts.

We retrieve the items directly from the database, which distinguishes our corpus from web-crawled corpora such as the CC100 dataset (Conneau et al., 2020), on which XLM-R and X-MOD have been trained. Another difference to CC100 is that our corpus extends to 2022, while the former has been created in or before 2019. Previous work shows that adaptation to more recent data can improve performance on present-time downstream tasks (Lazaridou et al., 2021).

We rely on the metadata provided by `Swissdox@LiRI` to select the articles in the respective languages. For each language, we hold out articles of the most recent days in the dataset (at least 200 articles) as a validation set. Like previous work (Conneau and Lample, 2019; Conneau et al., 2020), we use exponential smoothing to upsample languages with fewer documents, setting $\alpha = 0.3$.

3.2 Modularity

We follow recommendations by Pfeiffer et al. (2022) for ensuring the modularity of SwissBERT. When pre-training our language adapters, we freeze the shared parameters of the transformer layers. Conversely, when fine-tuning on downstream tasks, we freeze the language adapters and train the shared parameters. Pfeiffer et al. (2022) freeze the embed-

ding layer as well, in order to demonstrate transfer learning across languages with different subword vocabularies. In this paper, we do not perform experiments of this kind and do not freeze the embedding layer.

3.3 Vocabulary

The X-MOD model reuses the vocabulary of XLM-R, which has 250k tokens and has been created based on text in 100 languages (Conneau et al., 2020). This presents an interesting trade-off. On the one hand, X-MOD already has useful pre-trained multilingual embeddings. On the other hand, creating a new vocabulary could allow us to represent Switzerland-related words with a smaller degree of segmentation. This is especially relevant for the Romansh language, which did not contribute to the XLM-R vocabulary and as a consequence, is split into many subwords by XLM-R:

Co din ins quai per rumantsch?

Co din in|s quai per
rum|ants|ch ?

To further explore this trade-off, we train two variants of SwissBERT (Figure 2):

Variant 1: reused vocabulary We reuse the XLM-R vocabulary of X-MOD and freeze the pre-trained embeddings. As a consequence, we only train the language adapters. The other parameters remain identical to X-MOD.

Variant 2: new vocabulary We create a new multilingual vocabulary based on our pre-training corpus. We follow the procedure of XLM-R (Conneau et al., 2020) but restrict the vocabulary size to 50k words. Specifically, we use SentencePiece (Kudo and Richardson, 2018) to create a cased unigram language model (Kudo, 2018) with

⁵<https://swissdox.linguistik.uzh.ch/>

default settings, again smoothing the languages with $\alpha = 0.3$. We then train a new embedding matrix, including new positional embeddings. Following the recommendation by Pfeiffer et al. (2022), we initialize subwords that occur in the original vocabulary with the original embeddings.

Analyzing the new vocabulary, we find that 18k of the 50k subwords occur in the original XLM-R vocabulary, and the other 32k are new subwords. Appendix H lists the new subwords that occur most frequently in the corpus. Most are Romansh words, orthographic variants, media titles, toponyms or political entities of Switzerland.

3.4 Preprocessing

We preprocess the news articles by removing any markup and separating the layout elements, such as headlines, crossheadings, image captions and sidebars, with the special token `</s>`. We also remove bylines with author names, photographer names etc., wherever they are marked up as such.

Since previous work has shown that metadata can benefit language modeling (Dhingra et al., 2022), we prefix the articles with their medium and date, for example:

```
<medium> rtr.ch <year> 2019 <month> July
</s> ...
```

where `<medium>`, `<year>` and `<month>` are special tokens. When training Variant 1, we use the separator symbol instead of custom special tokens:

```
</s> rtr.ch </s> 2019 </s> July </s> ...
```

3.5 Data Analysis

Additional analysis of the pre-training corpus is provided in the appendices. Appendix C shows that there is no relevant overlap with the datasets we use for downstream evaluation. Appendix G breaks down the number of tokens for each pre-training language, news medium and year of publication.

3.6 Pre-training Setup

We generally use the same pre-training setup, implemented in Fairseq (Ott et al., 2019), as was used for X-MOD. We make some changes to optimize the efficiency of our pre-training. Namely, we do not split the articles into sentences but instead train on random contiguous spans of 512 tokens. In addition, we use a peak learning rate of $7e-4$ throughout. We train with an effective batch size of 768 across 8 RTX 2080 Ti GPUs. Both variants of SwissBERT were trained for 10 epochs.

Initialization strategy	Validation ppl.
Italian (IT_IT)	<u>2.53</u>
Random initialization	2.95±.13

Table 1: Preliminary experiments for choosing the best initialization of the Italian (IT_CH) language adapter. We report the standard deviation across three random initializations.

Initialization strategy	Validation ppl.
Italian (IT_IT)	1.85
French (FR_XX)	1.85
German (DE_DE)	1.87
Average of all Romance languages	1.90
Random initialization	<u>1.82±.02</u>

Table 2: Preliminary experiments for choosing the best initialization of the Romansh language adapter. The overall perplexity is lower than in Table 1 due to the high degree of segmentation when segmenting Romansh text with the XLM-R vocabulary.

3.7 Initialization of Language Adapters

In order to choose a strategy for initializing the language adapters, we perform some preliminary experiments based on Variant 1. Our goal is to train adapters for four language varieties: DE_CH, FR_CH, IT_CH and RM_CH. Three languages already have adapters in X-MOD – DE_DE, FR_XX and IT_IT – and so we expect that the best result can be achieved by continuing training these adapters.

We verify this hypothesis on the example of Italian. Table 1 shows the validation perplexity of the model after pre-training on the Italian part of our corpus for 2k steps. An adapter initialized with X-MOD’s Italian adapter yields a lower perplexity than a randomly initialized adapter. Thus, domain-adaptive (and variety-adaptive) pre-training seems more efficient than training an adapter from scratch.

In the case of Romansh, we similarly hypothesize that initializing from Italian or another Romance language will outperform a randomly initialized adapter, given the relatedness of these languages. However, Table 2 shows that random initialization yields a lower perplexity for Romansh. In addition, averaging multiple language adapters – e.g., the adapters for all the Romance languages in X-MOD – is clearly not a viable strategy. Given

these findings, we opt for the following initialization strategy:

- DE_CH from DE_DE;
- FR_CH from FR_XX;
- IT_CH from IT_IT;
- RM_CH from scratch.

4 Evaluation

For evaluating SwissBERT, we focus on Switzerland-related natural language understanding tasks, and especially multilingual and cross-lingual tasks on the token or sequence level.

4.1 Tasks

Named Entity Recognition (NER) Our main question is whether SwissBERT has improved natural language understanding capabilities in the domain it has been adapted to. To evaluate this, we annotate named entities in contemporary news articles and test whether a SwissBERT model fine-tuned on NER can detect the entities with higher accuracy than baseline models.

We name our test set SwissNER.⁶ Specifically, we annotate 200 paragraphs per language that we extracted from publicly accessible articles by the Swiss Broadcasting Corporation (SRG SSR). The annotated articles have been published in February 2023 and are thus not contained in the pre-training corpus. Appendices B and E describe the dataset in detail.

For fine-tuning on the NER task we use WikiNEuRal, an automatically labeled dataset in nine languages (Tedeschi et al., 2021). Only the data in German, French and Italian are relevant to SwissBERT, and so we train the model jointly on these three parts of WikiNEuRal. As a consequence, when training baselines on WikiNEuRal, we report separate results for training only on German, French and Italian, and for training on all the nine languages.

Since WikiNEuRal does not contain training data in Romansh, we evaluate zero-shot transfer to this language. In the case of X-MOD, we activate the Italian adapter when performing inference on Romansh.

⁶<https://huggingface.co/datasets/ZurichNLP/SwissNER>

NER on Historical News In addition to contemporary news, we report results for two datasets from the HIPE-2022 shared task (Ehrmann et al., 2022). Other than SwissNER, this task involves NER on mostly historical, OCR-processed news articles:

- *hipe2020*: We fine-tune and evaluate on annotated Swiss and Luxembourgish newspaper articles from the *Impresso* collection (Ehrmann et al., 2020) that are written in French or German, ranging between the years 1798 and 2018.
- *letemps*: We fine-tune and evaluate on annotated newspaper articles from two Swiss newspapers in French (Ehrmann et al., 2016), ranging between 1804 and 1981.

Stance Detection Another source of domain shift, apart from historical text, could be user-generated text. We evaluate our models on multilingual stance detection with the *X-stance* dataset (Vamvas and Sennrich, 2020), which is based on comments written by Swiss political candidates. The dataset contains 67k comments on various political issues in either German, French or Italian. Given a question and a comment, the task is to judge whether the candidate has taken a stance in favor or against the issue at hand. We follow Vamvas and Sennrich (2020) and use the concatenation of the two sequences as an input to SwissBERT:

```
<s> [question] </s></s>
[comment] </s>
```

The model is then trained to predict a binary label for the sequence pair based on the hidden state for <s>.

Sentence Retrieval To further investigate SwissBERT’s ability to align text in the Romansh language to the other languages, we construct a sentence retrieval task out of a German–Romansh parallel corpus of 597 unique sentence pairs (Dolev, 2023). This task is inspired by parallel corpus mining tasks (Zweigenbaum et al., 2017) and the Tatoeba test set used by Artetxe and Schwenk (2019).

Specifically, we use the German sentences as queries and report top-1 accuracy when retrieving the corresponding Romansh sentences. As similarity metric we use BERTScore (Zhang et al., 2020), which allows us to use the pre-trained models di-

	Supervised DE_CH	Supervised FR_CH	Supervised IT_CH	Zero-shot RM_CH
XLM-R (Conneau et al., 2020)				
– fine-tuned on 9 languages	70.7±1.0	70.9±0.6	76.6±1.2	63.8±0.7
– fine-tuned on DE, FR, IT	71.7±0.7	70.5±0.2	76.7±0.7	64.6±0.7
X-MOD (Pfeiffer et al., 2022)				
– fine-tuned on 9 languages	71.2±0.7	70.4±0.3	75.9±0.9	61.5±0.7
– fine-tuned on DE, FR, IT	72.2±0.5	71.8±1.1	76.7±0.8	61.4±1.8
SwissBERT (fine-tuned on DE, FR, IT)				
– reused vocabulary	74.5±0.8	74.2±0.9	78.6±0.1	81.8±0.9
– new vocabulary	<u>74.8±1.2</u>	<u>75.9±0.8</u>	<u>79.2±0.5</u>	<u>83.7±0.9</u>

Table 3: Named entity recognition results on the SwissNER test set. The last column reports zero-shot results for Romansh. Since X-MOD does not have a Romansh adapter, we use the Italian adapter when applying X-MOD to the Romansh test set. The best results are underlined.

	Coarse		Fine		Fine	
	hipe2020 FR	hipe2020 DE	letemps FR	hipe2020 FR	hipe2020 DE	letemps FR
French Europeana BERT (Schweter, 2020)	<u>81.2±0.4</u>	-	<u>68.3±1.7</u>	<u>75.9±0.6</u>	-	<u>63.0±1.2</u>
German Europeana BERT (Schweter, 2020)	-	<u>76.1±0.7</u>	-	-	<u>68.2±1.0</u>	-
XLM-R (Conneau et al., 2020)	79.3±1.1	72.7±1.5	66.1±1.2	73.6±1.3	64.4±0.8	60.6±1.0
X-MOD (Pfeiffer et al., 2022)	77.2±1.1	69.0±2.1	63.5±1.1	70.2±1.1	58.9±2.4	58.1±1.1
SwissBERT						
– reused vocabulary	77.7±1.3	69.2±1.9	64.3±1.1	71.7±1.1	58.8±1.1	57.6±1.1
– new vocabulary	80.0±1.4	71.6±1.9	66.2±1.1	73.4±1.0	62.2±1.7	60.4±1.4

Table 4: Named entity recognition on historical newspapers (HIPE-2022, (Ehrmann et al., 2022)). We report a strict micro-averaged F1-score for the coarse tag set (left) and the fine-grained tag set (right).

rectly without any fine-tuning.⁷ While Zhang et al. (2020) recommend using a validation set to determine the best transformer layer for BERTScore, we opt for a simpler approach and use the average hidden states across all layers.

The German–Romansh sentence pairs have been sampled by Dolev (2023) from press releases published by the Canton of the Grisons between 1997 and 2022.⁸ Most of the releases were originally written in German and then manually translated into Romansh Grischun (Scherrer and Cartoni, 2012). The gold sentence alignment is based on an automatic alignment that has been manually verified by a trained linguist.

Word Alignment Finally, we evaluate SwissBERT on German–Romansh word alignment using the unsupervised SimAlign technique (Jalili Sabet et al., 2020). For testing we use the same parallel sentences as above, which have been manually annotated with word alignments by a trained linguist (Dolev, 2023). We predict word alignments

⁷Note that calculating BERTScore for all pairs of sentences is viable in the context of this experiment but would not be efficient for large-scale parallel corpus mining.

⁸<https://github.com/eyldlv/DERMIT-Corpus>

using the “Match” variant of SimAlign and report the F1-score with regard to the gold annotations. We do not perform a grid search to find the optimal layer but instead average the hidden states across all transformer layers.

4.2 Baseline Models

General-purpose models

- XLM-R, a model trained jointly on 100 languages (Conneau et al., 2020)
- X-MOD, a model trained with language adapters on 81 languages, which is the basis of SwissBERT (Pfeiffer et al., 2022)

Specialized models

- Europeana BERT models pre-trained on historical newspapers in the German or French language (Schweter, 2020)

	Supervised DE	Supervised FR	Cross-topic DE	Cross-topic FR	Cross-lingual IT
XLM-R (Conneau et al., 2020)	76.9±0.8	78.6±0.8	73.0±1.3	75.3±1.9	74.4±0.7
X-MOD (Pfeiffer et al., 2022)	77.5±0.7	78.5±0.7	73.6±0.7	74.5±0.8	74.7±0.8
SwissBERT					
– reused vocabulary	77.9±0.4	79.2±0.3	73.8±0.5	74.5±0.8	74.8±0.6
– new vocabulary	78.3±0.4	80.1±0.5	74.0±0.6	75.8±0.5	74.9±0.7

Table 5: Stance detection on political comments in the X-stance dataset (Vamvas and Sennrich, 2020). We report the F1-score for different test sets of X-stance.

	Sentence retrieval	Word alignment
XLM-R (Conneau et al., 2020)	25.3	62.6
X-MOD (Pfeiffer et al., 2022)	31.8	65.1
SwissBERT		
– reused vocabulary	92.0	85.9
– new vocabulary	95.6	86.4

Table 6: German–Romansh parallel corpus alignment: sentence retrieval accuracy and word alignment F1-score across 597 sentence pairs.

4.3 Fine-tuning

We try to avoid hyperparameter optimization and instead use settings from previous work that are known to work well for XLM-R and similar models.

- For fine-tuning on WikiNEuRal, we train the models for 3 epochs with a learning rate of $2e-5$ and a batch size of 16.⁹ We report the average and standard deviation across 5 random seeds.
- For fine-tuning on the HIPE-2022 datasets, we use a learning rate of $5e-5$ and a batch size of 8 (Ehrmann et al., 2022). However, we train for up to 25 epochs to ensure that all models converge. We report the average and standard deviation across 10 random seeds.
- For fine-tuning on X-stance, we train with a learning rate of $1e-5$ and a batch size of 16 for 3 epochs, with a maximum sequence length of 256 tokens (Schick and Schütze, 2021). We report the average and standard deviation across 10 random seeds.

We implement fine-tuning with the *transformers* library (Wolf et al., 2020) and otherwise use the default settings of the library.

⁹<https://huggingface.co/Babelscape/wikineural-multilingual-ner>

4.4 Results

Evaluation results are presented in Tables 3–6. Overall, we find that SwissBERT outperforms the baselines on Switzerland-related tasks, and especially on Romansh.

Furthermore, the results show that using a custom vocabulary when adapting X-MOD is beneficial, not only for Romansh but also for the three languages that are represented in the original XLM-R vocabulary. One reason could be that the custom vocabulary better matches the evaluation domain. Another reason could be that the model has more capacity to adapt to the target domain if the embedding layer is trained in addition to the language adapters, irrespective of the vocabulary.

An interesting comparison is NER on contemporary news (Table 3) and historical news (Table 4). While SwissBERT outperforms the baselines on contemporary news, the model is not consistently better than XLM-R on historical news. On the latter task, SwissBERT strongly improves over the non-adapted model, X-MOD, but inherits the low baseline performance of X-MOD compared to XLM-R. One explanation why XLM-R outperforms X-MOD on historical NER is that it was trained for more steps with a larger batch size. Secondly, XLM-R does not depend on language identification, which might be beneficial when training on historical or OCR-processed text. We find that monolingual models trained on historical news surpass general-purpose multilingual models, confirming previous findings (Ehrmann et al., 2022; Ryser et al., 2022).

The X-stance task is informative because it is based on user-generated text, as opposed to newspaper articles. SwissBERT moderately but systematically outperforms the baselines on this task (Table 5), which indicates that it could be a useful model for processing not only news, but Switzerland-related text in general.

Finally, the German–Romansh alignment exper-

iment (Table 6) demonstrates that self-supervised training is sufficient to enable multilingual representations for Romansh Grischun, despite the rather small pre-training corpus. SwissBERT strongly outperforms multilingual encoders that have not been specifically trained on Romansh. We expect that SwissBERT could be a valuable resource for future Romansh NLP applications, such as classification, retrieval, or parallel corpus alignment.

5 Conclusion

We release a language model that supports the four national languages of Switzerland. Specific challenges of the Swiss language situation are addressed using methods from the recent literature, including multilingual masked language modeling, language adapters, and adaptive pre-training. We evaluate the resulting model, which we call SwissBERT, on a range of Switzerland-related natural language understanding tasks and mostly see an improved accuracy. In addition, SwissBERT excels in tasks involving Romansh, compared to models that do not cover this language.

Limitations

The SwissBERT model and our evaluation experiments have a limited scope. First of all, the training objective of SwissBERT limits the range of direct applications. SwissBERT is mainly intended for tagging tokens in written text (e.g., named entity recognition, part-of-speech tagging), text classification, and the encoding of words, sentences or documents into fixed-size embeddings. SwissBERT is not designed for generating text.

Secondly, we expect SwissBERT to perform best on input that is similar to our pre-training corpus of written news. Switzerland also has language varieties that are rarely found in newspapers, e.g., Swiss German and dialects of Romansh. While these are currently not covered by SwissBERT, the model is designed to be extensible.

Finally, the main goal of our evaluation experiments is to verify that the adaptation of SwissBERT has been effective, i.e., that SwissBERT has a higher accuracy on Switzerland-related tasks than non-adapted baselines. We do not methodically compare different approaches. In this paper, we present one approach that we have found to work well, but further ablation experiments would be required to verify that it is the optimal approach.

Acknowledgements

This work was funded by the Swiss National Science Foundation (project MUTAMUR; no. 176727). It makes use of media data made available via Swissdiox@LiRI by the Linguistic Research Infrastructure of the University of Zurich (see <https://t.uzh.ch/lhI> for more information). We thank Pedro Ortiz Suarez for early feedback and Eyal Dolev, Maud Ehrmann, Sven Najem-Meyer and Jonas Pfeiffer for help with downstream evaluation.

References

- Jesujoba O. Alabi, David Ifeoluwa Adelani, Marius Mosbach, and Dietrich Klakow. 2022. [Adapting pre-trained language models to African languages via multilingual adaptive fine-tuning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4336–4349, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- 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.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Bhuwan Dhingra, Jeremy R. Cole, Julian Martin Eisenschlos, Daniel Gillick, Jacob Eisenstein, and William W. Cohen. 2022. [Time-aware language models as temporal knowledge bases](#). *Transactions of the Association for Computational Linguistics*, 10:257–273.
- Eyal Liron Dolev. 2023. Does mBERT understand romansh? evaluating word embeddings using word alignment. In *Proceedings of the 8th Swiss Text Analytics Conference (SwissText)*.
- Maud Ehrmann, Giovanni Colavizza, Yannick Rochat, and Frédéric Kaplan. 2016. [Diachronic evaluation of ner systems on old newspapers](#). pages 97–107, Bochum, Germany. Bochumer Linguistische Arbeitsberichte.
- Maud Ehrmann, Matteo Romanello, Simon Clemenide, Phillip Benjamin Ströbel, and Raphaël Barman. 2020. [Language resources for historical newspapers: the impresso collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 958–968, Marseille, France. European Language Resources Association.
- Maud Ehrmann, Matteo Romanello, Sven Najem-Meyer, Antoine Doucet, and Simon Clemenide. 2022. [Overview of HIPE-2022: named entity recognition and linking in multilingual historical documents](#). In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 423–446. Springer.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Xiaochuang Han and Jacob Eisenstein. 2019. [Unsupervised domain adaptation of contextualized embeddings for sequence labeling](#). 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 4238–4248, Hong Kong, China. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. [SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Rasmus Kær Jørgensen, Mareike Hartmann, Xiang Dai, and Desmond Elliott. 2021. [mDAPT: Multilingual domain adaptive pretraining in a single model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3404–3418, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. [Quantifying the carbon emissions of machine learning](#). *arXiv preprint arXiv:1910.09700*.
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, Dani Yogatama, Kris Cao, Susannah Young, and Phil Blunsom. 2021. [Mind the gap: Assessing temporal generalization in neural language models](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 29348–29363. Curran Associates, Inc.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Al-lauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. 2020. [FlauBERT: Unsupervised language model pre-training for French](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France. European Language Resources Association.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Laura Mascarell, Tatyana Ruzsics, Christian Schneebeli, Philippe Schlattner, Luca Campanella, Severin Klingler, and Cristina Kadar. 2021. [Stance detection in German news articles](#). In *Proceedings of the*

- Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 66–77, Dominican Republic. Association for Computational Linguistics.
- Matteo Muffo and Enrico Bertino. 2020. **BERTino: An italian distilBERT model**. In *Proceedings of the Seventh Italian Conference on Computational Linguistics (CLiC-it 2020)*, volume 2769. CEUR.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. **fairseq: A fast, extensible toolkit for sequence modeling**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022. **Lifting the curse of multilinguality by pre-training modular transformers**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3479–3495, Seattle, United States. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. **MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Marco Polignano, Pierpaolo Basile, Marco de Gemmis, Giovanni Semeraro, and Valerio Basile. 2019. **ALBERTo: Italian BERT language understanding model for NLP challenging tasks based on tweets**. In *Proceedings of the Sixth Italian Conference on Computational Linguistics (CLiC-it 2019)*, volume 2481. CEUR.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. **How good is your tokenizer? on the monolingual performance of multilingual language models**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, Online. Association for Computational Linguistics.
- Anja Ryser, Quynh Anh Nguyen, Niclas Bodenmann, and Shih-Yun Chen. 2022. **Exploring transformers for multilingual historical named entity recognition**. In *Proceedings of the Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum*, pages 1090–1108, Bologna, Italy. CEUR-WS.
- Raphael Scheible, Fabian Thomczyk, Patric Tippmann, Victor Jaravine, and Martin Boeker. 2020. **Gotbert: a pure german language model**. *CoRR*, abs/2012.02110.
- Yves Scherrer and Bruno Cartoni. 2012. **The trilingual ALLEGRA corpus: Presentation and possible use for lexicon induction**. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2890–2896, Istanbul, Turkey. European Language Resources Association (ELRA).
- Timo Schick and Hinrich Schütze. 2021. **Exploiting cloze-questions for few-shot text classification and natural language inference**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Stefan Schweter. 2020. **Europeana bert and electra models**.
- Stefan Schweter, Luisa März, Katharina Schmid, and Erion Çano. 2022. **hmBERT: Historical multilingual language models for named entity recognition**. In *Proceedings of the Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum*, pages 1109–1129, Bologna, Italy. CEUR-WS.
- Simone Tedeschi, Valentino Maiorca, Niccolò Campolungo, Francesco Cecconi, and Roberto Navigli. 2021. **WikiNEuRal: Combined neural and knowledge-based silver data creation for multilingual NER**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2521–2533, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang. 2002. **Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition**. In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. **Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition**. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Jannis Vamvas and Rico Sennrich. 2020. **X-Stance: A multilingual multi-target dataset for stance detection**. In *Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS)*, Zurich, Switzerland.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen,

Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Pierre Zweigenbaum, Serge Sharoff, and Reinhard Rapp. 2017. [Overview of the second BUCC shared task: Spotting parallel sentences in comparable corpora](#). In *Proceedings of the 10th Workshop on Building and Using Comparable Corpora*, pages 60–67, Vancouver, Canada. Association for Computational Linguistics.

A Model Card

A.1 Model Details

A.1.1 Model Description

- Model type: X-MOD ([Pfeiffer et al., 2022](#)).
- Languages: German, French, Italian, Romansh.
- License: Attribution-NonCommercial 4.0 International (CC BY-NC 4.0).
- Fine-tuned from model: `xmod-base`

A.1.2 Model Sources

- Source code:
<https://github.com/ZurichNLP/swissbert>
- Model weights:
<https://huggingface.co/ZurichNLP/swissbert>
- Backup:
<https://doi.org/10.5281/zenodo.8016844>

A.2 Bias, Risks, and Limitations

- The model was adapted on written news articles and might perform worse on other domains or language varieties.

- While we have removed many author bylines, we did not anonymize the pre-training corpus. The model might have memorized information that has been described in the news but is no longer in the public interest.

A.3 Training Details

A.3.1 Training Data

German, French, Italian and Romansh documents in the `Swissdox@LiRI` database, until 2022 (Section 3.1).

A.3.2 Training Procedure

Masked language modeling ([Devlin et al., 2019](#); [Conneau et al., 2020](#)).

A.4 Environmental Impact

- Hardware type: RTX 2080 Ti
- Hours used: 2 models \times 10 epochs \times 18 hours \times 8 devices = 2880 hours
- Site: Zurich, Switzerland
- Energy source: 100% hydropower¹⁰
- Carbon efficiency: 0.0016 kg CO₂e/kWh¹⁰
- Carbon emitted: 1.15 kg CO₂e ([Lacoste et al., 2019](#))

¹⁰Source: <https://t.uzh.ch/lrU>

B SwissNER Annotation Process

SwissNER is a dataset for named entity recognition based on manually annotated news articles in Swiss Standard German, French, Italian, and Romansh Grischun. We annotate a selection of articles that have been published in February 2023 on the following online news portals:

- German: <https://www.srf.ch/>
- French: <https://www.rts.ch/>
- Italian: <https://www.rsi.ch/>
- Romansh: <https://www.rtr.ch/>

The four portals belong to the Swiss Broadcasting Corporation (SRG SSR). We select news articles in the categories “Switzerland” or “Regional”. The articles in the individual languages are not translations of each other and tend to cover different regions of Switzerland, but the editing style and the overall topics are coherent.

For each article we extract the first two paragraphs after the lead paragraph. We follow the guidelines of the CoNLL-2002 and 2003 shared tasks (Tjong Kim Sang, 2002; Tjong Kim Sang and De Meulder, 2003) and annotate the names of persons, organizations, locations and miscellaneous entities. The annotation was performed by a single annotator.

C Discussion of Data Overlap

Below we analyze the data overlap between pre-training and downstream evaluation:

- SwissNER dataset: None of the articles are in the pre-training corpus, which does not contain articles from 2023.
- hipe2020: No overlap.
- letemps: No overlap.
- X-stance: The dataset does not contain news.
- German–Romansh parallel corpus:
 - German: 36 out of 597 sentences appear verbatim in the pre-training corpus.
 - Romansh: 23 out of 597 sentences appear verbatim in the pre-training corpus.

Note that the German sentences and Romansh sentences never appear together in the pre-training corpus, making it unlikely that overlap gives SwissBERT an advantage in the alignment task.

We also make sure to exclude articles that occur in the CHeeSE dataset (Mascarell et al., 2021) to facilitate future evaluation on this dataset.

D Model Sizes

	Adapters	Vocabulary	Parameters	Trained parameters (adaptation)
XLM-R (Conneau et al., 2020)	-	250 002	278 043 648	-
X-MOD (Pfeiffer et al., 2022)	81	250 002	852 472 320	-
SwissBERT				
– reused vocabulary	4	250 002	306 410 496	28 366 848
– new vocabulary	4	50 262	153 010 176	67 163 136

Table 7: Sizes of the models used in the experiments. The second variant of SwissBERT has fewer parameters due to the smaller vocabulary, but has more trained parameters because we train the embedding layer.

E SwissNER Data Statistics

	DE_CH	FR_CH	IT_CH	RM_CH	Total
Number of paragraphs	200	200	200	200	800
Number of tokens	9 498	11 434	12 423	13 356	46 711
Number of entities	479	475	556	591	2 101
– PER	104	92	93	118	407
– ORG	193	216	266	227	902
– LOC	182	167	197	246	792
– MISC	113	79	88	39	319

Table 8: Statistics for the SwissNER test sets.

F Additional Baselines for SwissNER

	Supervised DE_CH	Supervised FR_CH	Supervised IT_CH	Zero-shot RM_CH
wikineural-multilingual-ner (Tedeschi et al., 2021)	71.4	71.4	75.2	66.5
German Europeana BERT (Schweter, 2020)	67.6±0.5	64.4±1.2	68.6±1.1	58.8±0.6
French Europeana BERT (Schweter, 2020)	57.0±1.2	69.0±0.7	66.8±0.7	61.0±1.1
SwissBERT (new vocabulary)	74.8±1.2	75.9±0.8	79.2±0.5	83.7±0.9

Table 9: Results for additional baselines on the SwissNER test set. wikineural-multilingual-ner is an mBERT model fine-tuned by Tedeschi et al. (2021) on the WikiNEuRal dataset. The other models in the table have been fine-tuned on the German, French and Italian parts of WikiNEuRal.

G Pre-training Data Statistics

	DE_CH	FR_CH	IT_CH	RM_CH	Total
<i>Training set</i>					
Number of articles	17 832 421	3 681 679	48 238	32 750	21 595 088
Number of tokens:					
– in terms of XLM-R vocabulary	11 611 859 339	2 651 272 875	27 504 679	16 977 167	14 307 614 060
– in terms of the new SwissBERT vocabulary	9 857 117 034	2 384 955 915	26 825 471	13 286 172	12 282 184 592
<i>Validation set</i>					
Number of articles	1 401	263	214	211	2 089
Number of tokens:					
– in terms of XLM-R vocabulary	1 648 604	256 794	95 088	348 166	2 348 652
– in terms of the new SwissBERT vocabulary	1 416 928	234 498	93 450	267 904	2 012 780

Table 10: Number of articles and number of subword tokens in our pre-training data.




Figure 3: Number of tokens (in terms of XLM-R vocabulary) per year in the training set.

Medium	Articles	Tokens	Lang.	Medium	Articles	Tokens	Lang.
Neue Zürcher Zeitung	1 189 914	891 991 094	DE	Facts	31 044	45 951 513	DE
St. Galler Tagblatt	1 297 896	660 251 490	DE	Limmattaler Zeit. / MLZ	73 310	44 858 755	DE
Tages-Anzeiger	932 263	598 824 374	DE	luzernerzeitung.ch	49 806	44 347 311	DE
Berner Zeitung	999 254	539 180 275	DE	Berner Rundschau / MLZ	78 420	43 671 698	DE
Neue Luzerner Zeitung	950 663	494 536 537	DE	aargauerzeitung.ch	31 223	43 405 747	DE
Der Bund	720 637	492 383 793	DE	RTS.ch	76 210	39 938 988	FR
nzz.ch	495 217	472 250 494	DE	handelszeitung.ch	52 064	39 558 669	DE
Aargauer Zeitung / MLZ	697 225	422 696 456	DE	Handelszeitung	34 270	36 153 858	DE
Basler Zeitung	784 388	408 674 904	DE	Zentralschweiz am Sonntag	46 753	35 613 607	DE
Le Temps	400 971	398 369 886	FR	Zuger Zeitung	52 369	35 471 227	DE
Tribune de Genève	508 101	344 341 310	FR	Schweiz am Sonntag / MLZ	45 374	33 968 154	DE
cash.ch	536 877	312 495 330	DE	AZ-Tabloid / MLZ	79 120	33 694 480	DE
Blick	584 893	236 515 780	DE	landbote.ch	31 026	33 333 601	DE
Schweizer Illustrierte	239 447	230 676 599	DE	rts.ch	41 202	33 213 658	FR
tagesanzeiger.ch	249 453	226 827 823	DE	Limmattaler Tagblatt / MLZ	65 911	33 025 066	DE
Zürichsee-Zeitung	453 383	224 321 469	DE	Das Magazin	17 676	32 874 311	DE
tdg.ch	273 969	208 280 959	FR	rts Vidéo	3 460	31 243 548	FR
srf.ch	307 414	203 538 182	DE	Schweiz am Wochenende	40 152	30 710 690	DE
Le Matin	316 890	195 390 209	FR	Beobachter	33 005	30 659 313	DE
Der Landbote	371 054	195 316 528	DE	Urner Zeitung	40 906	29 990 931	DE
bernerzeitung.ch	224 923	179 355 942	DE	Bilan	25 174	29 777 904	DE
24 heures	251 709	166 651 833	FR	Tele	40 704	29 728 436	DE
20 minuten online	262 474	165 019 108	DE	Sonntag / MLZ	41 560	29 016 500	DE
20 minutes online	242 166	159 347 598	FR	langenthalertagblatt.ch	23 436	28 866 148	DE
Thurgauer Zeitung	328 194	154 509 265	DE	www.sf.tv	60 662	28 385 473	DE
SonntagsZeitung	181 129	153 405 631	DE	zsz.ch	25 174	28 250 302	DE
24 Heures	190 122	151 106 511	FR	zuonline.ch	24 308	27 908 168	DE
Finanz und Wirtschaft	174 529	149 271 153	DE	berneroberlaender.ch	23 401	27 692 709	DE
24heures.ch	179 011	146 908 146	FR	thunertagblatt.ch	23 525	27 621 041	DE
Soloth. Zeitung / MLZ	242 015	144 028 122	DE	Le Nouveau Quotidien	31 310	27 192 253	FR
lematin.ch	222 071	139 517 644	FR	Berner Oberländer	41 331	27 179 816	DE
blick.ch	234 082	138 524 249	DE	Wiler Zeitung	43 265	26 514 799	DE
Oltner Tagblatt / MLZ	209 641	132 205 560	DE	Appenzeller Zeitung	42 908	26 419 047	DE
derbund.ch	140 174	129 723 367	DE	Toggenburger Tagblatt	41 470	25 560 534	DE
Die Weltwoche	92 430	129 170 717	DE	20 Minuten	123 750	25 065 828	DE
Zofinger Tagblatt / MLZ	236 835	127 415 159	DE	bzbasel.ch	19 618	24 377 477	DE
Blick.ch	242 403	126 481 730	DE	bz - Zeit. f.d. Region Basel	32 265	24 078 982	DE
NZZ am Sonntag	149 052	120 989 713	DE	Obwaldner Zeitung	31 427	23 589 402	DE
tagblatt.ch	59 245	111 287 915	DE	Nidwaldner Zeitung	31 273	23 002 223	DE
Sonntagsblick	163 860	106 270 536	DE	schweizer-illustrierte.ch	29 438	22 230 522	DE
srf Video	15 981	105 429 243	DE	SWI swissinfo.ch	16 978	21 915 894	DE
bazonline.ch	105 805	98 151 941	DE	Bilan	13 993	20 086 032	FR
Le Matin Dimanche	98 321	88 226 207	FR	züritipp (Tages-Anzeiger)	39 076	19 358 626	DE
Solothurner Zeitung	170 381	87 820 588	DE	20 Minutes	90 470	18 526 178	FR
HandelsZeitung	83 784	81 119 111	DE	Badener Tagblatt	24 518	18 210 357	DE
Basellandsch. Zeit. / MLZ	129 362	78 326 422	DE	TV 8	35 231	16 991 076	FR
Aargauer Zeitung	156 674	75 069 605	DE	Ostschweiz am Sonntag	25 472	16 521 054	DE
fuw.ch	84 644	73 232 927	DE	badenertagblatt.ch	18 079	16 382 179	DE
20 minuten	297 464	70 811 524	DE	Der Sonntag / MLZ	20 582	15 563 949	DE
Luzerner Zeitung	107 555	68 626 196	DE	swissinfo.ch	9 484	15 276 786	DE
Zürcher Unterländer	126 689	68 400 789	DE	www.swissinfo.ch	11 373	14 884 339	FR
L'Hebdo	57 388	67 428 263	FR	rtr.ch	32 622	14 746 036	RM
Newsnetz	108 330	66 200 248	DE	solothurnerzeitung.ch	15 475	14 063 626	DE
Cash	67 878	64 433 800	DE	Thuner Tagblatt	19 570	13 719 511	DE
Die Wochenzeitung	49 112	60 875 629	DE	rsi.ch	36 741	12 871 992	IT
Werdenberger & Obertogg.	99 981	56 790 265	DE	dimanche.ch	16 770	12 617 761	FR
24 heures Région La Côte	90 197	56 212 161	FR	PME Magazine	10 371	12 481 323	FR
20 minutes	240 986	55 868 424	FR	BZ - Langenthaler Tagblatt	15 538	12 326 791	DE
L'Illustré	53 216	55 536 046	FR	Grenchner Tagblatt / MLZ	22 587	12 194 558	DE
letemps.ch	37 911	50 789 072	FR	24 h. Région Nord Vaudois	23 573	12 012 095	FR
Blick am Abend	152 891	49 614 784	DE	Grenchner Tagblatt	16 022	11 822 579	DE
Schweizer Familie	56 387	49 099 500	DE	24 h. Région Riviera Chablais	22 803	11 703 359	FR
Glückspost	83 685	48 476 534	DE	24 h. Région Lausannoise	20 208	10 093 139	FR
Mittelland Zeitung	91 380	46 214 129	DE	grenchnertagblatt.ch	11 506	10 062 085	DE

Table 11: Number of articles and number of tokens (in terms of XLM-R vocabulary) per news medium in the training set. Note that some media statistics are distributed over multiple variants of the title. We report the majority language for each medium and highlight the two media that have majority language Italian and Romansh, respectively. The table does not include media with fewer than 10M tokens in the training corpus.

H Vocabulary Analysis

Rank	Subword	Majority language	Rank	Subword	Majority language
126	.»	DE	1239	_Fussball	DE
163	_Franken	DE	1244	_Spital	DE
302	_francs	FR	1268	_vegnir	RM
335	_betg	RM	1271	_cunter	RM
359	_ins	DE	1274	_Covid	FR
387	_Zürcher	DE	1279	_Keystone	DE
403	_rsi	IT	1280	_Gallen	DE
405	_rtr	RM	1300	_Aktien	DE
428	_Berner	DE	1303	_Cussegl	RM
497	_Tagblatt	DE	1323	_liess	DE
516	_quai	RM	1327	_WM	DE
545	MLZ	DE	1335	_uschia	RM
589	_èn	RM	1340	_kommenden	DE
628	_Galler	DE	1363	_Kantons	DE
639	_Lausanne	FR	1380	_Federer	DE
654	_Gemeinderat	DE	1426	_chantun	RM
698	_Luzerner	DE	1434	_sajan	RM
699	_grossen	DE	1438	_erstmals	DE
701	strasse	DE	1439	_Thun	DE
706	_heisst	DE	1442	_dapli	RM
710	_Basler	DE	1445	_Stimmen	DE
732	_onns	RM	1452	_tranter	RM
741	_SVP	DE	1454	_dix	FR
748	_suisse	FR	1459	_fatg	RM
778	_Tribune	FR	1465	_suenten	RM
779	_anc	RM	1491	_milliards	FR
785	_Svizra	RM	1496	_Strassen	DE
807	_Matin	FR	1520	_Mitteilung	DE
816	_persunas	RM	1543	_Entscheid	DE
824	_Quai	RM	1555	_Urs	DE
840	_dentant	RM	1559	_Massnahmen	DE
918	_vegn	RM	1563	_Zurich	FR
924	_Aargauer	DE	1576	_müsse	DE
964	_Lugano	IT	1577	_Behörden	DE
971	_Bundesrat	DE	1580	_Stadtrat	DE
991	Anzeiger	DE	1590	_zeigte	DE
1011	onn	RM	1594	_Regierungsrat	DE
1026	_Grischun	RM	1603	_Kantonspolizei	DE
1033	_Luzern	DE	1604	_machte	DE
1049	sda	DE	1615	_Mrd	DE
1098	_RTR	RM	1644	_gemäss	DE
1102	_canton	FR	1646	_schliesslich	DE
1106	_könne	DE	1648	_Thurgauer	DE
1112	_Sieg	DE	1659	_Amt	DE
1146	Nous	FR	1686	_tdg	FR
1163	_Gemeinden	DE	1690	_Solothurner	DE
1186	_Temps	FR	1697	_duai	RM
1199	_erklärte	DE	1705	_UBS	DE
1208	_FDP	DE	1716	_Cun	RM
1218	_Dass	DE	1719	_CVP	DE

Table 12: The 100 most frequent subwords that appear in our custom SwissBERT vocabulary but not in the XLM-R vocabulary. The symbol _ (U+2581) is used by SentencePiece to denote preceding whitespace. For each subword we report the majority language, i.e., the language that contributes the subword more often than the other languages, after exponential smoothing. Out of the 100 top subwords, 26 originate from Romansh and usually have a functional meaning, e.g., ‘not’, ‘one’, ‘this’, or ‘in’. Other subwords can be explained by Swiss orthographic conventions, such as the use of ‘ss’ in place of ‘ß’ in Swiss Standard German or the use of outward-pointing guillemets without surrounding whitespace (‘.»’). Most remaining subwords are (parts of) media titles, toponyms or political entities. Given that XLM-R was created in or before 2019, the neologism _Covid belongs to the words that only occur in the SwissBERT vocabulary.