

# TartuNLP @ SIGTYP 2024 Shared Task: Adapting XLM-RoBERTa for Ancient and Historical Languages

**Aleksei Dorkin**

Institute of Computer Science  
University of Tartu  
aleksei.dorkin@ut.ee

**Kairit Sirts**

Institute of Computer Science  
University of Tartu  
kairit.sirts@ut.ee

## Abstract

We present our submission to the unconstrained subtask of the SIGTYP 2024 Shared Task on Word Embedding Evaluation for Ancient and Historical Languages for morphological annotation, POS-tagging, lemmatization, character- and word-level gap-filling. We developed a simple, uniform, and computationally lightweight approach based on the adapters framework using parameter-efficient fine-tuning. We applied the same adapter-based approach uniformly to all tasks and 16 languages by fine-tuning stacked language- and task-specific adapters. Our submission obtained an overall second place out of three submissions, with the first place in word-level gap-filling. Our results show the feasibility of adapting language models pre-trained on modern languages to historical and ancient languages via adapter training.

## 1 Introduction

The application of natural language processing techniques and pre-trained language models to analysis of ancient and historical languages is a compelling subject of research that has been so far overlooked. While there exist a number of benchmarks, such as GLUE (Wang et al., 2018), SuperGLUE (Wang et al., 2019), or XGLUE (Liang et al., 2020), for evaluating the quality of embeddings and language models for modern languages, such benchmarks are lacking for ancient and historical languages. Thus, the SIGTYP 2024 Shared Task on Word Embedding Evaluation for Ancient and Historical Languages contributes to filling this gap.

In the current transformer-based language models paradigm, one of the common approaches to solving the tasks present in such benchmarks is to use the task data to fine-tune an encoder transformer model. The approach was first introduced in Devlin et al. (2019) and was shown to yield superior results compared to the alternatives on the GLUE benchmark. Additionally, the pro-

posed pre-training method is very similar to the word-level gap-filling problem in the shared task. Consequently, the model can be applied to solving the problem directly. This motivates our choice to use a transformer model in the shared task.

Large pre-trained language models, however, are predominantly trained on corpora of modern languages, with few exceptions such as LatinBERT (Bamman and Burns, 2020). Ancient and historical languages generally lack sufficient data to perform full pre-training of large language models or to continue training from a checkpoint trained on some different language. Full fine-tuning of modern language models on a relatively small amount of data in an ancient/historical language might lead to overfitting and catastrophic forgetting. These considerations are also common in context of other low resource tasks or domains.

Several approaches have been proposed to alleviate these issues. For instance, the supplementary training approach proposed by Phang et al. (2018) involves first fine-tuning a pre-trained model on an intermediary labeled task with abundant data, and then on the target task which may have limited data. This approach showed gains over simply fine-tuning on the target task. Pfeiffer et al. (2020b) developed a cross-lingual transfer-learning approach based on the adapters framework for parameter efficient fine-tuning of language models (Houlsby et al., 2019; Bapna and Firat, 2019). Their approach involves fine-tuning a language adapter and a task adapter stacked on top of each other. This adapter-based method is expanded by Pfeiffer et al. (2021) by adopting a custom tokenizer and an embedding layer. Several ways of initializing the new embedding layers are compared on underrepresented modern languages. Since ancient and historical languages are underrepresented, the same techniques should be applicable in this context as well.

A useful feature of both the supplementary train-

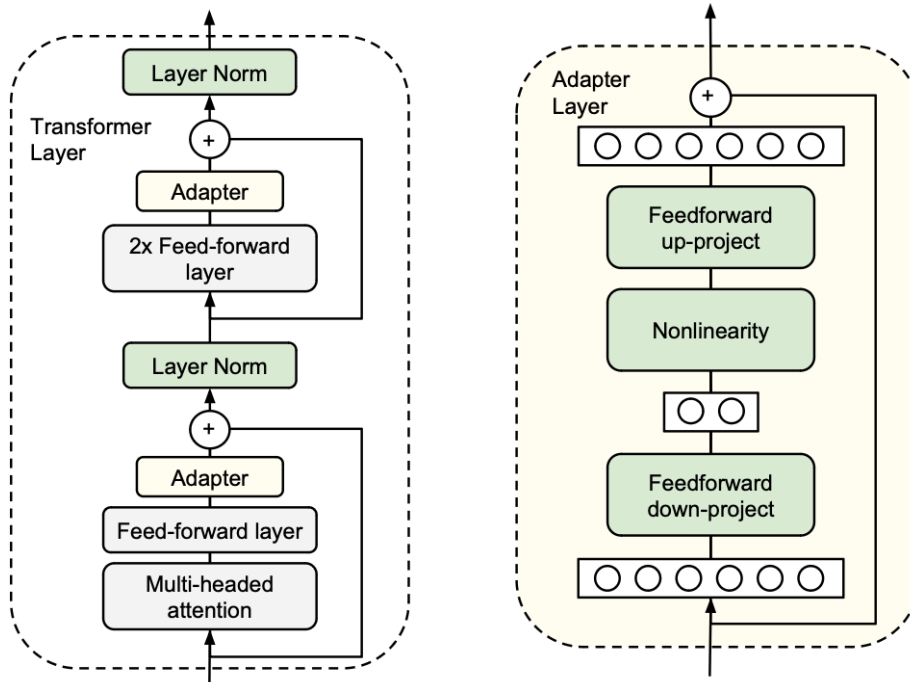


Figure 1: An illustration of the Bottleneck Adapter from (Houlsby et al., 2019). The left side demonstrates how a bottleneck adapter is added to a single transformer layer, while the structure of an individual adapter layer is on the right. Only elements in green are trained, while the rest remains frozen.

ing approach of Phang et al. (2018) and the adapter-based training of Pfeiffer et al. (2020b) is that they provide a uniform framework that can be applied to different languages and tasks in a similar manner. While previous related works mainly focused on modern languages, we aim to assess the feasibility of this unified approach for ancient and historical languages.

Our submission to the SIGTYP 2024 Shared Task on Ancient on Word Embedding Evaluation for Ancient and Historical Languages adopts the methods described by Pfeiffer et al. (2020b, 2021) by stacking fine-tuned language and task adapters, and customizing the tokenizer and the embedding layers. Our system is implemented as a unified framework, where various pre-trained models, languages, and tasks can be plugged in. The system was evaluated on POS tagging, morphological annotation, lemmatization, and filling-in both word-level and character-level gaps for 16 ancient and historical languages. We participated in the unconstrained subtask, which allowed using any additional resources, such as pre-trained language models. Out of 3 participants on the leaderboard we took a close second place, with the first place on the word gap-filling task. Our main contribution is showing that by adopting the parameter-efficient

adapter training methodology, the large language models pre-trained on modern languages are applicable also to low-resource ancient and historical languages.

## 2 Adapters

There exist a number of approaches to parameter efficient fine-tuning. In our submission to the shared task the focus is on adapters exclusively.

Houlsby et al. (2019) propose a parameter efficient fine-tuning strategy that involves injection of a number of additional trainable layers into the original architecture. The architecture of the proposed strategy is illustrated on Figure 1. In each transformer layer an adapter layer is added twice: after the attention sub-layer and after the feedforward sub-layer. To limit the number of trainable parameters, the authors propose a bottleneck architecture of adapter layers: the adapter first projects the input into a smaller dimension, applies non-linearity, then projects back to the original dimension. This is known as the bottleneck adapter.

Pfeiffer et al. (2020b) extend on the strategy in context of cross-lingual transfer learning in two ways. First, the adapter stack technique is introduced (illustrated on Figure 2), which is primarily used to separate language adaptation training and

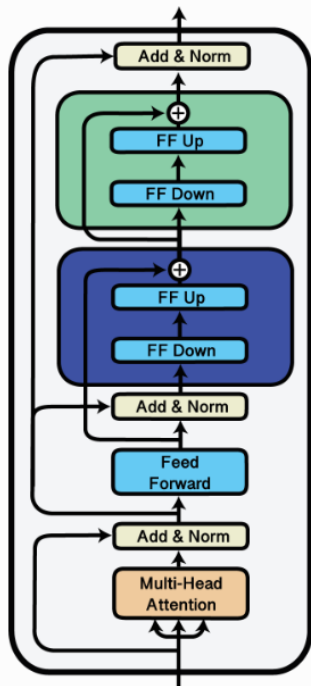


Figure 2: An illustration of an adapter stack as presented on the AdapterHub documentation page<sup>1</sup>. Blue and green blocks represent different adapter layers stacked on top of each other.

task-specific training. In the technique a language adapter is first trained, then a task-specific adapter is stacked on top of it and trained (while the language adapter is frozen). Secondly, the bottleneck adapter is expanded upon to include the embedding layer as well—the invertible adapter is introduced that transforms both input and output token representations to further improve language adaptation.

### 3 Methodology

Our approach in this work is based on training language adapters and task specific adapters for XLM-RoBERTa (Conneau et al., 2020)—the multilingual variant of RoBERTa (Zhuang et al., 2021) trained on 100 different languages. Since many of the languages in the shared task have related modern languages in the training data of XLM-RoBERTa, or are even included themselves (such as Latin), we expect to benefit from knowledge transfer.

#### 3.1 Data

The dataset provided by the organizers is a compilation of various resources (Bauer et al., 2017; Doyle, 2018; Ó Corráin et al., 1997; HAS Research Institute for Linguistics, 2018; Zeman et al., 2023; Acadamh Ríoga na hÉireann, 2017; Simon, 2014)

and comprises 16 languages in total spanning several historical epochs and upper bounded by 1700 CE. The information on the languages in the dataset is presented in Table 1.

#### 3.2 Adapter Training

Full fine-tuning of language models has several difficulties such as the possibility of catastrophic forgetting and the necessity to train and maintain a full copy of the model weights for each individual task. When applying the technique first proposed in (Phang et al., 2018) in a multi-lingual and multi-task setting, the scale of the mentioned disadvantages is magnified. Training and maintaining a separate copy of the model for each combination of language, target task, and intermediate task can become computationally expensive. Meanwhile, each copy of the model is narrowly specialized in a single task, with generalization capabilities being potentially limited. This highlights the practicality of parameter efficient fine-tuning in comparison.

The overall approach is, in general, the same for every language. First, we train a language adapter for every language individually. A language adapter is comprised of a bottleneck adapter and masked language modeling prediction head. Accordingly, the training objective is masked language modeling. This is based on the assumption that a significant portion of the model’s parameters is not actually language-specific, and thus does not need changing. We simply need to adapt the model to a new language. The language adapter is trained for 10 epochs regardless of the size of the data. The unmasked part of the training data for the word-filling task is used, while the masked part is not utilized for training. Secondly, we train task-specific adapters for each language. In this setup we use the adapter stack: the language adapter is loaded, but frozen, and we add a task-specific adapter on top of it, and train said adapter. We have two tasks per language: morphological tagging, which combines both POS-tagging and morphological annotation, and lemmatization, which is also implemented as sequence tagging problem. Similarly to language modeling, in both tasks the adapters are trained for 10 epochs.

#### 3.3 Custom tokenizers and embeddings

Some of the languages in the shared task are not covered by the model’s tokenizer. In other

<sup>1</sup>[https://docs.adapterhub.ml/adapter\\_composition.html](https://docs.adapterhub.ml/adapter_composition.html)

Language	Code	Train Sentences	Valid Sentences	Test Sentences
Ancient Greek	grc	24,800	3,100	3,101
Ancient Hebrew	hbo	1,263	158	158
Classical Chinese	lzh	68,991	8,624	8,624
Coptic	cop	1,730	216	217
Gothic	got	4,320	540	541
Medieval Icelandic	isl	21,820	2,728	2,728
Classical and Late Latin	lat	16,769	2,096	2,097
Medieval Latin	latm	30,176	3,772	3,773
Old Church Slavonic	chu	18,102	2,263	2,263
Old East Slavic	orv	24,788	3,098	3,099
Old French	fro	3,113	389	390
Vedic Sanskrit	san	3,197	400	400
Old Hungarian	ohu	21,346	2,668	2,669
Old Irish	sga	8,748	1,093	1,094
Middle Irish	mga	14,308	1,789	1,789
Early Modern Irish	ghc	24,440	3,055	3,056

Table 1: Language statistics in the Shared Task.

words, a significant portion of the tokenizer’s output includes <unk> tokens. This is possible because XLM-RoBERTa’s employs subword tokenization (Sennrich et al., 2016). A subword tokenizer is a trained model that learns a vocabulary of a certain size. The trained model transforms text input into individual tokens by looking for the longest character sequences in the text that are present in its vocabulary. When attempting to tokenize a text that includes symbols that did not appear in the training data (and consequently, are not in the tokenizer’s vocabulary), each unrecognized symbol is replaced by the <unk> token.

For each language not covered by the model’s tokenizer a new tokenizer is initialized. In addition to that, a new embedding layer is also initialized. The initial weights are copied from the original embeddings for tokens that overlap with the multilingual tokenizer.

To decide whether the model requires a custom tokenizer, two checks are made:

1. The percentage of unknown tokens > 5%
2. <unk> token is in top 10 by frequency

If either of these conditions is true, a custom tokenizer is created. Refer to Table 2 for detailed statistics. Based on these conditions, 5 languages need a separate tokenizer: Old Church Slavonic, Coptic, Classical Chinese, Old Hungarian, and Old East Slavic. When selecting the vocabulary size,

the aim was to have as many full words as possible in the vocabulary, but at the same time allow for morphological variation. Thus, the number of items in the vocabulary should be neither too high nor too low. For that reason, we selected the size to be 3000 for all languages, except Classical Chinese. For Classical Chinese the size was insufficient, so it was increased to 10000.

Pfeiffer et al. (2021) suggest several possible options of initializing the weights of the custom embedding layer corresponding to the new tokenizer. These include random initialization, copying the weights of tokens that overlap between multilingual and language-specific tokenizers, as well as a matrix factorization-based approach that aims to identify latent semantic concepts in the original embedding matrix that are useful for transfer. It is noted, however, that the latter approach shows less gains when used with languages with underrepresented scripts. For this reason and for the sake of simplicity, with settled on using the lexical overlap approach for all affected languages.

### 3.4 Tasks

**Filling Word-level Gaps** Having trained a masked language modeling adapter for each language, we already have a suitable model to perform word-level gap filling. However, the fact XLM-RoBERTa, as well as most modern language models, is based on subword tokens rather than full words becomes a significant hurdle. The reason is

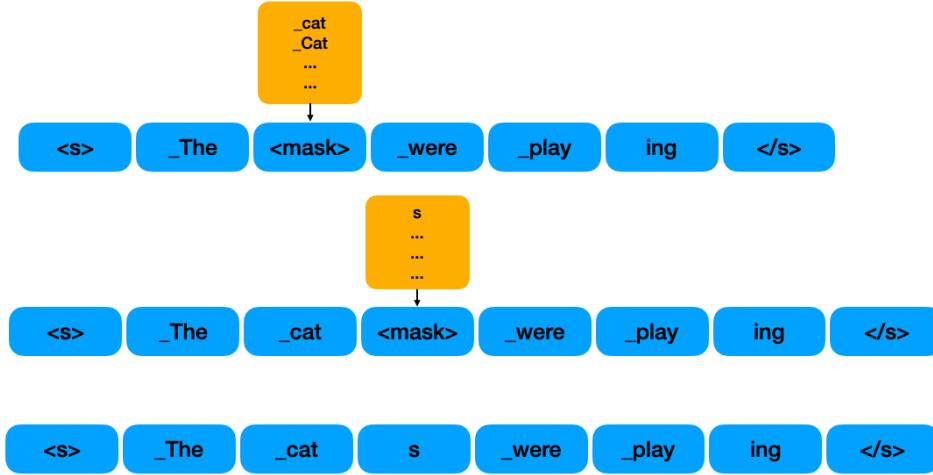


Figure 3: A schematic illustration of the decoding process for word-level mask filling. Blue boxes represent current tokens in the sentence, while orange boxes represent the probability distribution of tokens at the position of the mask token. The upper half represents the first step of decoding. We start with predicting the most likely replacement for the leftmost masked token that starts with `_` symbol representing the beginning of a new word. Then, we replace the mask with that token and append a new mask token to the right of it, as represented in the middle part. We predict the most likely replacement for the new mask token. If it starts with `_` symbol, we discard the mask token, consider the word predicted, and move to the next masked word if it’s present. Conversely, if there’s no `_` in the predicted token, we append it to the previously predicted token. We repeat this process  $k$  times, or until we encounter a token starting with `_`.  $k$  is a hyperparameter that may be tuned, however increasing  $k$  increases the decoding time significantly. For this reason we set  $k$  to 1 for all languages. At the bottom of the figure the final result with no mask tokens is demonstrated. Note that this description is specific to XLM-RoBERTa tokenizer, other model’s tokenizers may have different behaviour.

lang	# tokens	% unknown	<unk> rank
chu	495612	15.31%	1
cop	39771	45.61%	2
fro	55448	0.00%	-
ghc	1282852	0.00%	-
got	98874	1.05%	-
grc	1079457	0.47%	-
hbo	124063	0.36%	-
isl	688580	0.00%	-
lat	296770	0.00%	-
latm	897209	0.00%	-
lzh	408259	1.40%	4
mga	492894	0.00%	-
ohu	352811	5.05%	1
orv	592890	5.00%	2
san	59349	0.01%	-
sga	190711	0.00%	-

Table 2: Tokenizer coverage statistics. <unk> ranks not present in top 10 are not reported.

the we can not know in advance whether the target word we want to predict is comprised of one or more subword units, and that makes the decoding

more complicated. To address this issue, for each sentence with masked words in it, we follow the steps (see Figure 3):

1. Locate the first <mask> token
2. Use the model’s prediction to determine the most likely token that starts with the whitespace prefix
3. Replace the <mask> token with the prediction, and look ahead up to  $k$  steps: does the next token start with the whitespace token? If it does, we break the cycle, and continue processing the remaining <mask> tokens in the sentence, otherwise we append the next predicted token to our previous prediction, and repeat the look ahead.

For Classical Chinese, the process is simplified, and we simply predicted the most likely token.

**Filling Character-level Gaps** The approach to character-level mask filling is purely algorithmic. We start with building a vocabulary of masked word candidates. We consider a sequence of characters with no whitespaces to be a suitable candidate.

The vocabulary is populated with such candidates. For each masked word a look up among all the words of the same length as the masked word is performed: a regular expression built from the masked word by replacing "[\_]" with "." is matched against these words. If there are matches, up to the first 3 matches are returned. For each match character replacements are extracted from each candidate. If there are no matches, and the masked word contains only one "[\_]", the masked word is split in two parts, and each part is matched against the dictionary. If there is at least one match, the whitespace is returned as the replacement.

The described algorithmic method could be extended in two ways:

1. Multiple candidates could be reranked using the trained language model
2. If there are no candidates, a suitable word could be generated using the trained model, similarly to the word-level gap-filling task

However, the extensions were not developed for the shared task for the following reasons. Each sentence may contain multiple masked words, and in addition the whitespace symbol may be masked as well. This significantly increases the solution's complexity, yet the magnitude of the benefit is unclear.

Finally, for Classical Chinese the following approach was applied. First, the symbols in the original dataset were decomposed into the lowest possible forms which are mostly strokes and small indivisible units using the Hanzipy library<sup>2</sup>, then a masked language model adapter was trained on that decomposed data. The model achieved relatively high masked language modeling evaluation accuracy—about 40%—on the validation set, as measured by the Trainer class in the transformers library. The trained model was then used to predict the most likely replacement for each masked symbol. The ground truth data provided by the organizers, however, remained in the original composed form. This proved to be a significant problem. An attempt was made to compose the symbols back into their original form algorithmically, since the library used for decomposition doesn't have the option to reconstruct the original symbol from its constituents, to our best knowledge. However, this attempt was unsuccessful given that our submission attained score of 0 on the test set.

<sup>2</sup><https://github.com/Synkied/hanzipy>

### POS-tagging & Full Morphological Annotation

We frame both POS-tagging and full morphological annotation as a single token classification task. For each token we collect all available morphological features and the POS tag from the annotation. Then we concatenate them as a single string, which is then used as the class to predict. Then we employ the adapter stack technique: we load the corresponding language adapter, and create a task specific adapter with token classification prediction head. We only train the task specific adapter. For inference we decompose the predicted class strings back into individual tags. The approach is the same for every language.

**Lemmatization** Similarly to the previous task, we frame lemmatization as token classification task. To achieve that we generate transformation rules for each lemma/form pair based on the technique first introduced in (Straka, 2018) and expanded upon in (Dorkin and Sirts, 2023). The rules are comprised of individual edits that need to be made to transform the given form into its lemma. The edits are represented as a single string. We use that string as the label to predict. Consequently, inference is done in two steps: first the rule for each token is predicted, then that rule is applied to the form and the resulting lemma is returned. The approach is the same for all language except for Classical Chinese. For Classical Chinese we do not train a model, but rather use a simple dictionary look up which results in nearly 100% accuracy.

### 3.5 Technical Implementation Details

The implementation is primarily based on the Adapters<sup>3</sup> (Poth et al., 2023) library and the provided training script examples. The most significant change in the training scripts is the addition of custom embedding layer initialization. Some aspects of lemmatization as a token classification task were adopted from our previous work. For languages not utilizing custom tokenizers we employ the invertible bottleneck adapters (Pfeiffer et al., 2020b), while for the languages that do require custom tokenizers we employ the regular bottleneck adapters (Houlsby et al., 2019). The motivation is that when we have a custom embedding layer, there is no need to adapt it additionally. In all experiments we use the base version of XLM-RoBERTa due to time constraints, however we can well expect that the large version would show improved perfor-

<sup>3</sup><https://github.com/adaptor-hub/adapters>

	Overall results	POS-tagging	Lemmatization	Morphological analysis	Character gap-filling	Word gap-filling
Ancient Greek	0.70	0.96	0.94	0.97	0.61	0.03
Ancient Hebrew	0.61	0.94	0.97	0.95	0.19	0.00
Classical Chinese	0.57	0.83	1.00	0.89	0.00	0.10
Coptic	0.52	0.61	0.75	0.75	0.45	0.02
Gothic	0.70	0.93	0.93	0.92	0.67	0.03
Medieval Icelandic	0.73	0.97	0.98	0.96	0.57	0.17
Classical and Late Latin	0.73	0.96	0.97	0.96	0.66	0.11
Medieval Latin	0.76	0.99	0.99	0.99	0.70	0.14
Old Church Slavonic	0.50	0.66	0.60	0.67	0.54	0.02
Old East Slavic	0.56	0.76	0.69	0.80	0.48	0.06
Old French	0.69	0.95	0.92	0.98	0.52	0.07
Vedic Sanskrit	0.66	0.84	0.88	0.86	0.65	0.05
Old Hungarian	0.52	0.75	0.63	0.76	0.46	0.00
Old Irish	0.19	-	-	-	0.35	0.03
Middle Irish	0.22	-	-	-	0.39	0.04
Early Modern Irish	0.28	-	-	-	0.50	0.06

Table 3: Results of our submission on the leaderboard. For the variants of Irish, training and test data was only provided for gap-filling tasks.

mance. The code and the instructions to reproduce are available on the project’s GitHub repository<sup>4</sup>.

XLM-RoBERTa uses SentencePiece (Kudo and Richardson, 2018) for tokenization. SentencePiece is a trainable tokenization model limited by a token vocabulary of a specific size. To be able to process languages with scripts which are under-represented by XLM-RoBERTa’s tokenizer, we train a new tokenizer for each language in question. The new tokenizer has to remain compatible with original model—it has to retain all the special tokens and their indices. To achieve that we use the `train_new_from_iterator` method of the `Tokenizer` class instance in HuggingFace transformers associated with model. We supply it with the same data we use for the masked language modeling training.

For training we used a single NVIDIA Tesla V100 GPU on the University’s High Performance Cluster (University of Tartu, 2018). The training time for a single task (including masked language modeling) on average was about 10-20 minutes, ranging from 2 to 40 minutes depending on the amount of data, which is notably less than it would take for full fine-tuning.

<sup>4</sup><https://github.com/slowwavesleep/ancient-lang-adapters/tree/sigtyp2024>

Initially, an error was present in our code that resulted in our models not having trainable language modeling prediction heads. In other words, during the language adaptation stage the models retained the original masked language prediction head of XLM-RoBERTa, rather than training a new language specific prediction head. To our surprise, after fixing the error, the difference turned out to be quite negligible, except for languages that relied on custom embeddings. For these languages, the effect of the fixing the error was most noticeable on the word-level gap-filling task—before that the models were using the prediction head which was mismatched in size with the custom embeddings.

## 4 Results

According to the results on Table 3, our approach seems to generally underperform on languages that require custom tokenizers and embeddings: Classical Chinese, Coptic, Old Church Slavonic, Old East Slavic, Old Hungarian. This can be explained by the amount of data being too limited to be able to train meaningful input representations, while the lexical overlap technique failed to provide benefit due to significant differences in their scripts. We hypothesize that a more sophisticated approach

to embedding initialization and tokenizer training could result in considerable benefits.

We note a surprisingly strong performance of our algorithmic approach to character-level gap-filling with the exception of Classical Chinese where the approach is simply not applicable. We expect the extensions to the approach outlined in Section 3.4 could improve the results further. However, in context of the shared task, we do not believe that the possible improvements would be worth the effort. Due to the whitespaces appearing as masked characters there is no way to know in advance how many individual words in the sentence in total have masked characters in them. In addition to that, a single word may contain several masked characters. The combination of these two factors makes improvements based on language model rescoring quite computationally expensive. For Classical Chinese, however, we hypothesise that the score would be considerably higher if the ground truth was also in the decomposed form, because to the best of our understanding, reconstructing the original characters from individual strokes is not a straightforward task.

Despite taking the first place in the word-level gap filling task, the absolute scores are still very low. This is to be expected, because the difficulty of the task is compounded by the presence of multiple masked words per sentence, as well the reliance on subword tokenization of our approach. Additionally, the implementation may not be suitable at all for certain languages due to differences in the writing systems. Finally, we surmise that due to the amount of training being quite small for the masked language modeling problem, we cannot expect to be able to correctly predict completely new words, however we can, in theory, predict previously unseen word forms due to subword tokenization. This means that experimenting with the  $k$  parameter used to limit the number of times the look ahead for word continuation is performed in the word-level gap filling problem could improve the results somewhat.

The pattern-based approach to lemmatization also performs reasonably well in most languages. The approach is restricted by the set of patterns derived based on the training set. If the evaluation set contains many words for which a suitable pattern is missing, the system simply cannot make a correct prediction. However, we believe that this is not the problem here. In particular, we note that those languages which have lower scores on

lemmatization, also perform consistently lower on other tasks. Thus, we suggest that the lower performance is rather related to insufficient data for training meaningful representations. In addition to underperforming on languages with underrepresented writing systems, we observe somewhat low performance on Vedic Sanskrit as well. One explanation could be that the usage of diacritics may negatively affect the character-based transformation rule generation; however this hypothesis requires further investigation.

Finally, our combined approach to POS-tagging and morphological analysis performs quite well on most of the languages. This is somewhat unexpected, because the model is limited in predicting only the combinations of part-of-speech and morphological features that are present in training data.

## 5 Conclusion

This paper described our solution to the SIGTYP 2024 Shared Task developed based on the adapters framework. The system is simple and uniform, and can be easily extended to other tasks and languages. For each language we trained a language adapter, and two task adapters, one for POS and morphological tagging and one for lemmatization. For languages with scripts underrepresented in the XLM-RoBERTa vocabulary, we additionally created custom tokenizers and embeddings. One notable advantage of the adopted approach is its resource efficiency—adapting the model to a new language and task can be done in less than an hour. As a future work, the system could be improved with more advanced adapter-based techniques such as adapter fusion (Pfeiffer et al., 2020a) to leverage typological relatedness of languages.

## References

- Acadamh Ríoga na hÉireann. 2017. *Corpas Stairiúil na Gaeilge 1600-1926*. Retrieved: June 10, 2022.
- David Bamman and Patrick J. Burns. 2020. *Latin Bert: A Contextual Language Model for Classical Philology*.
- Ankur Bapna and Orhan Firat. 2019. Simple, Scalable Adaptation for Neural Machine Translation. 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 1538–1548.
- Bernhard Bauer, Rijcklof Hofman, and Pádraic Moran.



2017. [St. Gall Priscian Glosses, version 2.0](#). Accessed: February 14, 2023.
- 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.
- 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.
- Aleksei Dorkin and Kairit Sirts. 2023. Comparison of Current Approaches to Lemmatization: A Case Study in Estonian. In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 280–285.
- Adrian Doyle. 2018. [Würzburg Irish Glosses](#). Accessed: February 14, 2023.
- HAS Research Institute for Linguistics. 2018. [Old Hungarian Codices](#).
- 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 *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- 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.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A New Benchmark Dataset for Cross-lingual Pre-training, Understanding and Generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2020a. [AdapterFusion: Non-Destructive Task Composition for Transfer Learning](#). *CoRR*, abs/2005.00247.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. [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.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2021. [UNKs Everywhere: Adapting Multilingual Language Models to New Scripts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10186–10203, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Clifton Poth, Hannah Sterz, Indraneil Paul, Sukannya Purkayastha, Leon Engländer, Timo Imhof, Ivan Vulić, Sebastian Ruder, Iryna Gurevych, and Jonas Pfeiffer. 2023. [Adapters: A Unified Library for Parameter-Efficient and Modular Transfer Learning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 149–160, Singapore. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Eszter Simon. 2014. Corpus Building from Old Hungarian Codices. In Katalin É. Kiss, editor, *The Evolution of Functional Left Peripheries in Hungarian Syntax*. Oxford University Press, Oxford.
- Milan Straka. 2018. [UDPipe 2.0 prototype at CoNLL 2018 UD Shared Task](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 197–207, Brussels, Belgium. Association for Computational Linguistics.
- University of Tartu. 2018. [UT Rocket](#).
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*,

pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Daniel Zeman, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aepli, Hamid Aghaei, Željko Agić, Amir Ahmadi, Lars Ahrenberg, Chika Kennedy Ajede, Salih Furkan Akkurt, Gabrielè Aleksandravičiūtė, Ika Alfina, Avner Algom, Khalid Alnajjar, Chiara Alzetta, Erik Andersen, Lene Antonsen, Tatsuya Aoyama, Katya Aplonova, Angelina Aquino, Carolina Aragon, Glyd Aranes, Maria Jesus Aranzabe, Bilge Nas Arican, Hórunn Arnardóttir, Gashaw Arutie, Jessica Naraiswari Arwidarasti, Masayuki Asahara, Katla Ásgeirsdóttir, Deniz Baran Aslan, Cengiz Asmazoğlu, Luma Ateyah, Furkan Atmaca, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Mariana Avelãs, Elena Badmaeva, Keerthana Balasubramani, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Starkaður Barkarson, Rodolfo Basile, Victoria Basmov, Colin Batchelor, John Bauer, Seyyit Talha Bedir, Shabnam Behzad, Kepa Bengoetxea, İbrahim Benli, Yifat Ben Moshe, Gözde Berk, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnė Bielinskienė, Kristín Bjarnadóttir, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Anouck Braggaar, António Branco, Kristina Brokaitė, Aljoscha Burchardt, Marisa Campos, Marie Candito, Bernard Caron, Gauthier Caron, Catarina Carvalho, Rita Carvalho, Lauren Cassidy, Maria Clara Castro, Sérgio Castro, Tatiana Cavalcanti, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Neslihan Cesur, Savas Cetin, Özlem Çetinoğlu, Fabricio Chalub, Liyanage Chamila, Shweta Chauhan, Ethan Chi, Taishi Chika, Yongseok Cho, Jinho Choi, Jayeol Chun, Juyeon Chung, Alessandra T. Cignarella, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Daniela Corbetta, Francisco Costa, Marine Courtin, Mihaela Cristescu, Ingerid Løyning Dale, Philemon Daniel, Elizabeth Davidson, Leonel Figueiredo de Alencar, Mathieu Dehouck, Martina de Laurentiis, Marie-Catherine de Marneffe, Valeria de Paiva, Mehmet Oguz Derin, Elvis de Souza, Arantza Diaz de Ilarraza, Carly Dickerson, Arawinda Dinakaramani, Elisa Di Nuovo, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Adrian Doyle, Timothy Dozat, Kira Drogonova, Puneet Dwivedi, Christian Ebert, Hanne Eckhoff, Masaki Eguchi, Sandra Eiche, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Olga Erina, Tomaž Erjavec, Farah Essaidi, Aline Etienne, Wograinne Evelyn, Sidney Farcundes, Richárd Farkas, Federica Favero, Jannatul Ferdaousi, Marília Fernanda, Hector Fernandez Alcalde, Amal Fethi, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Federica Gamba, Marcos Garcia, Moa Gärdenfors, Fabrício Ferraz Gerardi, Kim Gerdes, Luke Gessler, Filip Ginter, Gustavo Godoy, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Bernadeta Griciūtė, Matias Grioni, Loïc Grobol, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-

Barbance, Tunga Güngör, Nizar Habash, Hinrik Hafsteinsson, Jan Hajič, Jan Hajič jr., Mika Hämäläinen, Linh Hà Mỹ, Na-Rae Han, Muhammad Yudiantira Hanifmuti, Takahiro Harada, Sam Hardwick, Kim Harris, Dag Haug, Johannes Heinecke, Oliver Hellwig, Felix Hennig, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Marivel Huerta Mendez, Jena Hwang, Takumi Ikeda, Anton Karl Ingason, Radu Ion, Elena Irimia, Olájjidé Ishola, Artan Islamaj, Kaoru Ito, Siratun Jannat, Tomáš Jelínek, Apoorva Jha, Katharine Jiang, Anders Johannsen, Hildur Jónsdóttir, Fredrik Jørgensen, Markus Juutinen, Hüner Kaşıkara, Nadezhda Kabaeva, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Neslihan Kara, Ritván Karahóga, Andre Kåsen, Tolga Kayadelen, Sarveswaran Kengatharaiyer, Václava Kettnerová, Jesse Kirchner, Elena Klementieva, Elena Klyachko, Arne Köhn, Abdulatif Köksal, Kamil Kopacewicz, Timo Korikiakangas, Mehmet Köse, Alexey Koshevoy, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Parameswari Krishnamurthy, Sandra Kübler, Adrian Kuçi, Oğuzhan Kuyrukçu, Aslı Kuzgun, Sookyoung Kwak, Kris Kyle, Veronika Laippala, Lorenzo Lambertino, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phuong Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Maria Levina, Lauren Levine, Cheuk Ying Li, Josie Li, Keying Li, Yixuan Li, Yuan Li, KyungTae Lim, Bruna Lima Padovani, Yi-Ju Jessica Lin, Krister Lindén, Yang Janet Liu, Nikola Ljubešić, Olga Loginova, Stefano Lusito, Andry Luthfi, Mikko Luukko, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Menel Mahamdi, Jean Maillard, Ilya Makarchuk, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Büşra Marşan, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Stella Markantonatou, Héctor Martínez Alonso, Lorena Martín Rodríguez, André Martins, Cláudia Martins, Jan Mašek, Hiroshi Matsuda, Yuji Matsumoto, Alessandro Mazzei, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Tatiana Merzhevich, Niko Miekka, Aaron Miller, Karina Mischenkova, Anna Missilä, Cătălin Mititelu, Maria Mitrofan, Yusuke Miyao, AmirHossein Mojiri Foroushani, Judit Molnár, Amirsaeid Moloodi, Simonetta Montemagni, Amir More, Laura Moreno Romero, Giovanni Moretti, Shinsuke Mori, Tomohiko Morioka, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Robert Munro, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Mariam Nakhlé, Juan Ignacio Navarro Horñiacek, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Manuela Nevaci, Luong Nguyễn Thị, Huyền Nguyễn Thị Minh, Yoshihiro Nikaido, Vitaly Nikolaev, Rattima Nitisaroj, Alireza Nourian, Hanna Nurmi, Stina Ojala, Atul Kr. Ojha, Hulda Óladóttir, Adedayo Olúòkun, Mai Omura, Emeka Onwuegbuzia, Noam Ordan, Petya Osenova, Robert Östling, Lilja Óvrelid, Şaziye Betül Özateş, Merve Özçelik, Arzucan Özgür, Balkız Öztürk Başaran, Teresa Paccosi, Alessio Palmero Aprosio, Anastasia Panova, Hyunji Hayley Park, Niko Partanen, Elena Pascual, Marco Passarotti,

Agnieszka Patejuk, Guilherme Paulino-Passos, Giulia Pedonese, Angelika Peljak-Łapińska, Siyao Peng, Siyao Logan Peng, Rita Pereira, Sílvia Pereira, Cene Augusto Perez, Natalia Perkova, Guy Perrier, Slav Petrov, Daria Petrova, Andrea Peverelli, Jason Pheilan, Jussi Piitulainen, Yuval Pinter, Clara Pinto, Tommi A Pirinen, Emily Pitler, Magdalena Plamada, Barbara Plank, Thierry Poibeau, Larisa Ponomareva, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Robert Pugh, Tiina Puolakainen, Sampo Pyysalo, Peng Qi, Andreia Querido, Andriela Rääbis, Alexandre Rademaker, Mizanur Rahoman, Taraka Rama, Loganathan Ramasamy, Joana Ramos, Fam Rashel, Mohammad Sadegh Rasooli, Vinit Ravishankar, Livy Real, Petru Rebeja, Siva Reddy, Mathilde Regnault, Georg Rehm, Arij Riabi, Ivan Riabov, Michael Rießler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Putri Rizqiyah, Luisa Rocha, Eiríkur Rögnvaldsson, Ivan Roksandic, Mykhailo Romanenko, Rudolf Rosa, Valentin Roşca, Davide Rovati, Ben Rozonoyer, Olga Rudina, Jack Rueter, Kristján Rúnarsson, Shoval Sadde, Pegah Safari, Aleksí Sahala, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Ezgi Sanyar, Dage Särg, Marta Sartor, Mitsuya Sasaki, Baiba Saulīte, Yanin Sawanakunanon, Shefali Saxena, Kevin Scannell, Salvatore Scarlata, Nathan Schneider, Sebastian Schuster, Lane Schwartz, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Syeda Shahzadi, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Yana Shishkina, Muh Shohibussirri, Maria Shvedova, Janine Siewert, Einar Freyr Sigurðsson, João Silva, Aline Silveira, Natalia Silveira, Sara Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Haukur Barri Símonarson, Kiril Simov, Dmitri Sitchinava, Ted Sither, Maria Skachedubova, Aaron Smith, Isabela Soares-Bastos, Per Erik Solberg, Barbara Sonnenhauser, Shafi Sourov, Rachele Sprugnoli, Vian Stamou, Steinhór Steingrímsson, Antonio Stella, Abishek Stephen, Milan Straka, Emmett Strickland, Jana Strnadová, Alane Suhr, Yogi Lesmana Sulestio, Umut Sulubacak, Shingo Suzuki, Daniel Swanson, Zsolt Szántó, Chihiro Taguchi, Dima Taji, Fabio Tamburini, Mary Ann C. Tan, Takaaki Tanaka, Dipta Tanaya, Mirko Tavoni, Samson Tella, Isabelle Tellier, Marinella Testori, Guillaume Thomas, Sara Tonelli, Liisi Torga, Marsida Toska, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Utku Türk, Francis Tyers, Sveinbjörn Hórdarson, Vilhjálmur Hörsteinsson, Sumire Uematsu, Roman Untilov, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Andrius Utka, Elena Vagnoni, Sowmya Vajjala, Socrates Vak, Rob van der Goot, Martine Vanhove, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Uliana Vedenina, Giulia Venturi, Veronika Vincze, Natalia Vlasova, Aya Wakasa, Joel C. Wallenberg, Lars Wallin, Abigail Walsh, Jonathan North Washington, Maximilian Wendt, Paul Widmer, Shira Wigderson, Sri Hartati Wijono, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Kayo Yamashita, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka, Marat M.

Yavrumyan, Arife Betül Yenice, Olcay Taner Yıldız, Zhuoran Yu, Arlisa Yuliawati, Zdeněk Žabokrtský, Shorouq Zahra, Amir Zeldes, He Zhou, Hanzhi Zhu, Yilun Zhu, Anna Zhuravleva, and Rayan Ziane. 2023. [Universal dependencies 2.12](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. [A Robustly Optimized BERT Pre-training Approach with Post-training](#). In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 1218–1227, Huhhot, China. Chinese Information Processing Society of China.

Donnchadh Ó Corráin, Hiram Morgan, Beatrix Färber, Benjamin Hazard, Emer Purcell, Caoimhín Ó Dónaill, Hilary Lavelle, Julianne Nyhan, and Emma McCarthy. 1997. [CELT: Corpus of Electronic Texts](#). Retrieved: March 15, 2021.