

Filling the Gaps in Ancient Akkadian Texts: A Masked Language Modelling Approach

Koren Lazar[◇] Benny Saret[♣] Asaf Yehudai[◇]

Wayne Horowitz[♣] Nathan Wasserman[♣] Gabriel Stanovsky[◇]

[◇]School of Computer Science and Engineering [♣]The Institute of Archaeology
The Hebrew University of Jerusalem, Jerusalem, Israel
{koren.lazar, gabriel.stanovsky}@mail.huji.ac.il

Abstract

We present models which complete missing text given transliterations of ancient Mesopotamian documents, originally written on cuneiform clay tablets (2500 BCE - 100 CE). Due to the tablets' deterioration, scholars often rely on contextual cues to manually fill in missing parts in the text in a subjective and time-consuming process. We identify that this challenge can be formulated as a masked language modelling task, used mostly as a pre-training objective for contextualized language models. Following, we develop several architectures focusing on the Akkadian language, the lingua franca of the time. We find that despite data scarcity (1M tokens) we can achieve state of the art performance on missing tokens prediction (89% hit@5) using a greedy decoding scheme and pretraining on data from other languages and different time periods. Finally, we conduct human evaluations showing the applicability of our models in assisting experts to transcribe texts in extinct languages.

1 Introduction

The Akkadian language was the lingua franca of the Middle East and Egypt in the Late Bronze and Early Iron Ages, spoken or in use from 2500 BCE until its gradual extinction around 100 CE (Oppenheim, 2013). It was written in cuneiform signs — wedge-shaped imprints on clay tablets, as depicted in Figure 1 (Walker, 1987). These tablets are the main record from the Mesopotamian cultures, including religious texts, bureaucratic records, royal decrees, and more. Therefore they are a target of extensive transcription and transliteration efforts. One such transcription is exemplified by the Latinized text to the right of the tablet in Figure 1.

The Open Richly Annotated Cuneiform Corpus (Oracc)¹ is one of the major Akkadian transcription collections, culminating in approximately

¹<http://oracc.org>



*a-na LUGAL EN-ia
ARAD-ka mDUG—IM—aš-šur
lu DI-mu a-na LUGAL EN-ia
ša ina É.GAL aš-pur-an-ni
nu-uk ID-MEŠ lu-x¹+x¹-ru^{*1}
li-ki-ri-ku ina x¹ ma²-dak²¹-te^{*}
né-bu-ru le-pu¹-[šiu]
ú-ma-a m²arba-il-[a.a]
LÚ^v.03.U₅ it-[tal-ka]
ma-a ID ina^{*} [me-lí]-šú²¹
i^{*1}-ba²-ši a-[x x x]-i¹
a-na-ku ina bat-i¹ti¹ [am-mi]-i¹ti^{*1}
la e-ti-ik [la al-lik]
ina UGU-*hi* šú-u ina i¹É^{*}.GAL^{*1}
a-sa-par ú-ma-a
am-ma-ka-ma [né-bu]-ru^{*1}
qa²-ni^{*} ma²-dak²¹-[ti]
le-pu-šú¹*

Figure 1: A clay tablet from Oracc (left) with its corresponding Latin transliteration (right). Words are delimited by spaces, while signs are delimited by hyphens or dots. A sign which is missing due to deterioration is denoted by ‘x’ and highlighted in red in the figure. We develop models which automatically complete these missing signs based on the surrounding context.

2.3M transcribed signs from 10K tablets. As further evidenced in Figure 1, many of the signs in the tablets were eroded over time and some parts were broken or lost, forcing editors to “fill in the gaps” where possible, based on the context of the surrounding words.

In this paper, we identify that the task of masked language modeling, used ubiquitously in recent years for pretraining other downstream tasks (Peters et al., 2018; Howard and Ruder, 2018; Liu et al., 2019) lends itself directly to missing sign prediction in the transliterated texts. We experiment with various adaptations of BERT-based models (Devlin et al., 2019) trained and tested on Oracc, combined with a greedy decoding scheme to extend the prediction from single tokens to multiple words. We specifically focus on the effect multilingual pretraining has on downstream performance, which was recently shown beneficial for low-resource settings (Chau et al., 2020).

In an automatic evaluation, we find that a combination of large-scale multilingual pretraining with Akkadian finetuning achieves state-of-the-art results, with a top 5 accuracy of 89.5%, vastly improving over other models and baselines. Interestingly, we find that the multilingual pretraining signal seems to be more important than the signal of the target small-scale Akkadian data, as the *zero-shot* performance of a multilingual language model surpasses that of a monolingual Akkadian model by about 10%.

Finally, we show the model’s potential applicability in assisting transcription by filling in missing parts. To account for the challenges in human assessment of an extinct language, we created a controlled setup where domain experts are asked to identify plausible predictions out of a combination of model predictions, the original masked sequences, and noise. We find that in a majority of cases, the annotators found at least one of the model’s top 3 predictions useful, while the performance degrades on longer sequences. Future work can improve the model by designing more elaborate decoding schemes and exploring the specific effect of related languages (e.g., Arabic and Hebrew) on downstream performance. Our code and trained models are made publicly available at www.github.com/SLAB-NLP/Akk.

Our main contributions are:

- We identify that the longstanding challenge of filling in gaps in Akkadian texts directly corresponds to advances in masked language modeling.
- We train the first Akkadian language model, which can serve as a pretrained starting point for other downstream tasks such as Akkadian morphological analysis.
- We develop state-of-the-art models for completing missing signs by combining large-scale multilingual pretraining with Akkadian language finetuning.
- We devise a controlled user study, showing the potential applicability of our model in assisting scholars fill in gaps in real-world Akkadian texts.

2 Background

In this section, we will introduce the Akkadian language and the Open Richly Annotated Cuneiform

Corpus (Oracc). While it is one of the largest sources of the Akkadian language, it is of orders of magnitude smaller compared to resources for other languages, such as English or German. Then, we will introduce masked language modeling, which will serve as the basis for our sign prediction model.

2.1 The Akkadian Language and the Oracc Dataset

Akkadian is a Semitic language, related to several languages spoken today, such as Hebrew, Aramaic, Amharic, Maltese, and Arabic. It has been documented from the 3rd millennium B.C.E. until the first century of the common era, in modern Iraq, between the Euphrates and the Tigris rivers, as well as in modern Syria, east Turkey, and the Northern Levant (Huehnergard, 2011). In this work, we will use the Open Richly Annotated Cuneiform Corpus (Oracc), one of the largest international cooperative projects gathering cuneiform texts from many archaeological sites.

Most relevant to this work, Oracc contains Latinized transliterations of the cuneiform texts, as can be seen in Figure 1, depicting a clay tablet and its transliteration in Oracc. It also contains English translations for parts of the texts. In total, as can be seen in Table 1, Oracc consists of about 10K texts (each a transliteration of a single tablet), containing 1M words and 2.3M signs, as well as 9K translated texts in English containing 1.2M English words. Importantly, the editors can often visually estimate the number of missing signs in a deteriorated or missing part and denote each with ‘x’ in the transliteration (marked in red in Figure 1). Therefore, in the following sections, we will assume that the number of missing signs is given as input to our models.

	# Texts	# Words	# Signs
Akkadian Train	8K	950K	1.8M
Akkadian Test	2K	250K	500K
English Train	7K	950K	–
English Test	2K	250K	–

Table 1: Number of texts, words, and signs in our preprocessed version of Oracc, English texts are corresponding translations of the Akkadian texts.

2.2 Multilingual Masked Language Modeling

In masked language modeling (MLM), a model is asked to predict masked parts in a text given their surrounding context. Recent years have seen large gains for almost all NLP tasks by using the token representations learned during MLM as a starting point for downstream applications. In particular, recent work has noticed that joint training on various languages greatly helps downstream applications, especially where labeled data is sparse (Pires et al., 2019; Chau et al., 2020; Conneau et al., 2020).

In this work we identify that the MLM objective directly corresponds to the task of filling in gaps in Akkadian texts and train several MLM variants on it. In the following sections, we will especially examine the effect of multilingual pretraining on our task.

3 Task Definition

Intuitively, our task, as demonstrated in Figure 2, is to predict missing tokens or signs given their context in transliterated Akkadian documents. Human experts achieve this when compiling Oracc by considering not only the surrounding context in the tablet, but also its wider, external context, such as its corpus, or the time and location where the text was originally written or found. In many cases, researchers can estimate the number of missing signs even after their physical deterioration, and mark them as sequences of ‘x’s. E.g., note the sequence of 2 ‘x’s marked in red in Figure 2. We will use this signal as input to our model, which specifies the number of signs to be predicted.²

Formally, let $T = (s_1, \dots, s_n) \in \Sigma^n$ be a transliterated Akkadian document comprised of a concatenation of n signs, where Σ is the set of all Akkadian signs. Let $I \subseteq [n]$ such that $\forall i \in I : s_i = x$, where x denotes a missing sign. The number of missing signs is assumed to be known a priori, based on the editor’s examination of the tablets. Therefore, the model should output $(p_1, \dots, p_{|I|}) \in \Sigma^{|I|}$ predictions for the missing signs in T .

4 Model

In this section, we will introduce BERT-based models aiming to solve the task of predicting miss-

²We filter cases where the editors can not estimate the number of missing signs.

ing signs in Akkadian texts. We chose these models since their pretraining task is also our downstream task. The high-level diagram of the model is presented in Figure 2 and is elaborated below. First, in Section 4.1, we outline the preprocessing of Oracc, aiming to remove annotations that are external to the original text. Then in Section 4.2, we propose two models for predicting missing signs. Lastly, in Section 4.3, we present an algorithm to extend BERT sub-word level prediction to multiple signs and words. In the following two sections we will test these models in both automatic and human evaluation setups.

4.1 Preprocessing

Oracc is a collaborative effort to transliterate Mesopotamian tablets, mainly in Akkadian. Figure 1 exemplifies different characteristics of the corpus. We removed signs added by editors in the transliteration process as they were not part of the original text. For example, we removed signs which indicate how certain the editors are in their reading of the tablet. As an example, note that in Figure 2 the first sign in the transliterated text is marked as uncertain with the ‘□’ characters before preprocessing. In addition, we also remove superscripts and subscripts, which indicate different readings of the Akkadian cuneiform text, e.g., an ‘^m’ superscript is preceding the last word in the transliterated text.

During training, similarly to Devlin et al. (2019), we train the model to predict known tokens by masking them at random. During inference, we mask each missing sign, indicated by ‘x’ in Oracc,

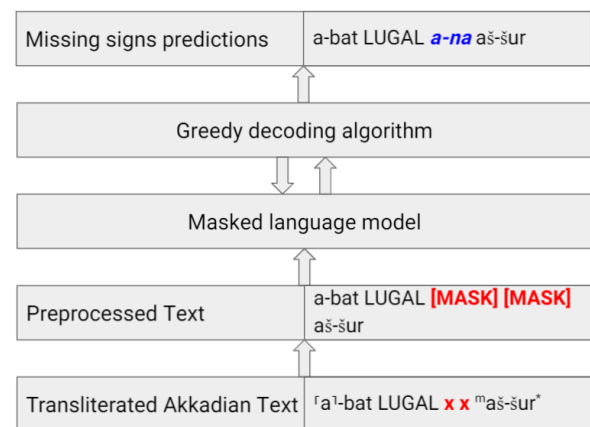


Figure 2: High-level diagram of our model, producing a sequence of signs (marked in blue) given input from Oracc with missing signs (red ‘x’s). We experiment with different language models and pretraining data.

and iteratively predict each of the tokens composing it.

4.2 Masked Language Models

We experimented with monolingual and multilingual versions of BERT.

First, we pretrained from scratch a monolingual BERT model with a reduced number of parameters (750K) following conclusions from Kaplan et al. (2020). Second, following recent research suggesting that pretraining on similar languages is beneficial for many NLP tasks, including in low-resource settings (Pires et al., 2019; Wu and Dredze, 2019; Chau et al., 2020; Conneau et al., 2020), we finetuned a pretrained multilingual BERT (M-BERT) model (Devlin et al., 2019).³ M-BERT was trained on the 104 most common languages of Wikipedia, including Hebrew and Arabic - Semitic languages that are typologically similar to Akkadian.

To adapt M-BERT to Akkadian, we assign its 99 available free tokens, optimizing for maximum likelihood by the WordPiece tokenization algorithm (Schuster and Nakajima, 2012; Wu et al., 2016).

4.3 Decoding: From Tokens to Signs

While the MLM task is designed to predict single tokens, in our setting, multiple signs and words may be omitted due to deterioration. To bridge this gap, we greedily extend the token level prediction by adapting the k-beams algorithm such that it outputs possible predictions given an Akkadian text with a sequence of missing signs. See the example at the top of Figure 2, where the two ‘x’ signs in the input are predicted as *a-na*. To achieve this, we count the number of sign delimiters (space, dot, hyphens) predicted at each time step, and choose the best *k* candidates according to the following conditional probability:

$$p(X_1, \dots, X_n, C) = \prod_{i=1}^n p(X_i | X_1, \dots, X_{i-1}, C) \quad (1)$$

Where X_i denotes the i_{th} masked token, and C denotes the observed context. For example, in Figure 2, *a-na* is composed of three sub-sign tokens: ‘a’, ‘-’, ‘na’, while $C = ('a-bat LUGAL', 'aš-šur')$, and the sequence probability is $p(na|- , a, C) \cdot p(-|a, C) \cdot p(a|C)$.

³<https://huggingface.co/bert-base-multilingual-cased>

5 Automatic Evaluation

We present an automatic evaluation of our models’ predictions for missing signs in ancient Akkadian texts, testing several masked language modeling variants for single token prediction, as well as our greedy extension to multiple tokens and signs. In all evaluations, we mask known tokens and evaluate the model’s ability to predict the original masked tokens. This setup allows us to test against large amounts of texts in Oracc from different periods of time, locations or genres.

5.1 Models and Datasets

We use two strong baselines: (1) the LSTM model that was proposed by Fetaya et al. (2020), and was retrained on our dataset using their default configuration;^{4,5} and (2) the cased BERT-base multilingual model, without finetuning over Oracc.⁶

We compare these two baselines against our models, as presented in 4.2, trained in three configurations: (1) BERT+AKK(mono) refers to the reduced size BERT model, trained from scratch on the Akkadian texts from Oracc; (2) MBERT+Akk is a finetuned version of M-BERT on the Akkadian texts, using the model’s additional free tokens to encode sub-word tokens from Oracc; and (3) MBERT+Akk+Eng further finetunes on the English translations available in Oracc to introduce additional domain-specific signal. We test all models against 5 different genres of Akkadian texts tagged in Oracc, masking 15% of the tokens. The genres can be largely divided into two groups. First, the Royal Inscription, Monumental, and Astrological Reports are the most common genres in the dataset and consist of longer coherent texts, mostly of essays and correspondence. Second, we test on two other genres: Lexical which consists mostly of tabular information (lists of synonyms and translations), and Decree that contains concatenated non-contextualized short sentences.

5.2 Experimental Setup

For all our experiments, we used a random 80% - 20% split for train and test (see Table 1). For the monolingual model, we trained our reduced-parameters BERT model from scratch for 300 epochs with 4 NVIDIA Tesla M60 GPUs for 2 hours. For the multilingual experiments, we fine-

⁴<https://github.com/DigitalPasts/Atrahasis>

⁵<https://github.com/DigitalPasts/Akkademia>

⁶<https://huggingface.co/bert-base-multilingual-cased>

Genre	Metric	LSTM	MBERT-base	BERT+AKK(mono)	MBERT+Akk	MBERT+Akk+Eng
Royal	MRR	.52	.57	.57	.83	.83
Inscription	Hit@5	.60	.65	.56	.90	.90
Royal or	MRR	.51	.61	.61	.84	.83
Monumental	Hit@5	.61	.69	.69	.90	.90
Astrological	MRR	.53	.55	.55	.81	.80
Report	Hit@5	.60	.64	.64	.88	.88
Lexical	MRR	.10	.61	.69	.69	.66
	Hit@5	.10	.76	.76	.85	.85
Decree	MRR	.49	.67	.39	.71	.74
	Hit@5	.60	.73	.51	.76	.76
Overall	MRR	.52	.60	.50	.83	.83
	Hit@5	.59	.67	.60	.89	.89

Table 2: MRR and Hit@5 precision by genre. The first two models from the left are our baselines: LSTM refers to the model from (Fetaya et al., 2020) retrained on our data, MBERT-base refers to the *zero-shot* M-BERT model without training on Oracc. The following three models are introduced in Section 4.2: BERT+AKK(mono) is trained mono-lingually from scratch on Oracc Akkadian texts; MBERT+Akk finetunes on Oracc Akkadian texts; and MBERT+Akk+Eng is also finetuned on their English translations. The three genres at the top of the Table (Royal Inscription, Monumental, Astrological) are the most common in our test dataset and contain longer, more coherent texts. The two genres at the bottom (Lexical and Decree) contain tabular texts and non-contextualized, short sentences.

tuned M-BERT for 20 epochs similarly to (Chau et al., 2020), with 8 NVIDIA Tesla M60 GPUs for 2-3 hours. We used the original architecture of M-BERT, adding a masked language modeling head for prediction. For the LSTM model of Fetaya et al. (2020), we train for 200 epochs, with 1 NVIDIA Tesla M60 GPU for 68 hours.

5.3 Metrics

We report performance according to the Hit@k and mean reciprocal rank (MRR) metrics, as defined below:

$$MRR = \frac{1}{N} \sum_{i=1}^N \frac{1}{rank_i} \quad (2)$$

$$Hit@k = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{[rank_i \leq k]} \quad (3)$$

Where N is the number of masked instances, $rank_i$ is the rank of the original masked token in the model’s predictions, and $\mathbb{1}$ is the indicator function.

The Hit@k metric directly measures applicability in our target application, i.e., how likely is the correct prediction to appear if we present the user with our model’s top k predictions. MRR complements Hit@k by providing a finer-grained evalua-

tion, as the model receives partial credit in correlation with every ranking.

5.4 Results

Table 2 compares token level evaluation across our different models and genres, while Figure 3 presents an evaluation of the prediction of multiple signs and words. We note several interesting observations based on these results.

Multilingual pretraining + Akkadian finetuning achieves state-of-the-art performance. On average, the two M-BERT models, which were finetuned over Oracc texts, outperform all other models by at least 20% on both metrics. This is particularly pronounced in the more natural first set of genres, where the multilingual models often surpass 85% in both MRR and Hit@5.

Zero-shot multilingual pretraining outperforms monolingual training. Surprisingly, in most tested settings, the zero-shot version of M-BERT outperforms both BERT+AKK(mono) and the LSTM models, despite never training on Akkadian. This suggests that the signal from pretraining is stronger than that of the Akkadian texts, likely due to the relatively small amounts of data. Moreover, as M-BERT was trained over the MLM task in other languages during its pretraining, this evaluation can be seen as a

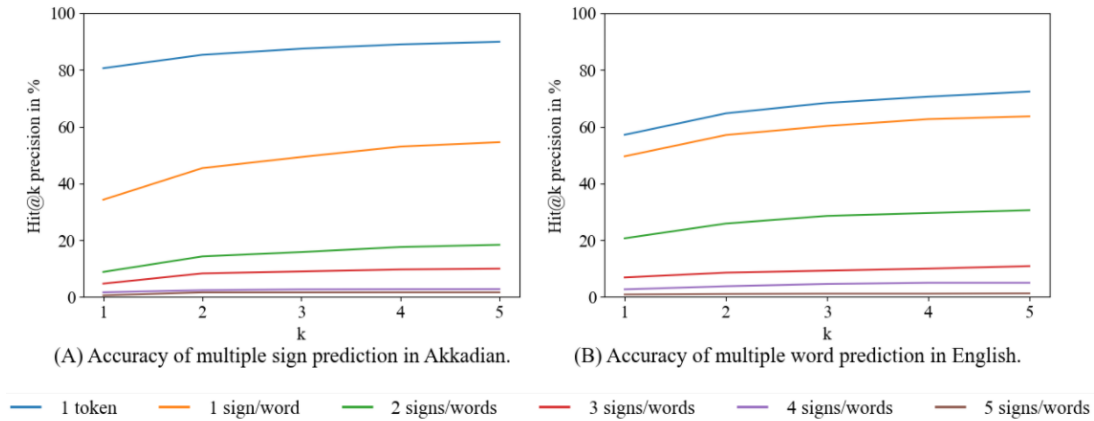


Figure 3: Hit@k precision for sequences of varying lengths in Akkadian (A) and English (B). We find that both languages do well on 1 token and 1 sign, where the correct answer is expected to be in the models’ top 5 predictions for half of the instances. Performance drops sharply for longer sequences, possibly due to the large search space. We directly measure the model’s applicability in user studies in Section 6.

zero-shot cross-lingual transfer learning, on which M-BERT was found to be competitive in many NLP tasks (Pires et al., 2019; Wu and Dredze, 2019; Conneau et al., 2020).

Performance degrades on the Lexical genre.

The gains of the multilingual models are reduced in the Lexical genre. Specifically, they are on par with BERT+AKK(mono) in this genre. This may indicate that this genre’s idiosyncratic syntax does not benefit much from multilingual pretraining.

Context matters after finetuning M-BERT.

The performance of the finetuned M-BERT is the lowest in the Decree genre and is very close to that of the MBERT-base. This is perhaps not surprising as the Decree texts are concatenations of unrelated short sentences, while one of BERT’s main advantages is its learned contextualized representations of different domains.

Finetuning on English Oracc translations does not improve performance.

Finetuning M-BERT only on Akkadian (MBERT+Akk) leads to results on par with additional finetuning on English (MBERT+Akk+Eng), possibly indicating that the amount of Akkadian texts and English translations is not enough to make M-BERT align between the two languages in Oracc’s unique domains.

Performance degrades on longer masked sequences for both English and Akkadian.

Figure 3 compares our best-performing model in predicting a varying number of signs against M-BERT on English texts, where both use our greedy

decoding strategy to extend their predictions to multiple signs and words. We note similar patterns for both languages. The performance for a single sign and word is high, and it deteriorates when more elements are predicted. In the following section, we extend this evaluation by conducting a human evaluation that aims to test the model’s applicability in a real-world setting.

6 Human Evaluation and User Studies

We note that the automatic evaluation presented in the previous section offers only an upper bound of the model’s ability to suggest reasonable completions, since the original text is often only one out of many other equiprobable completions of the masked text. Consider, for example, the masked English text at the top of Figure 4. While the original text was “*of the former*”, the model’s top predictions (“*of the previous*”, “*of the first*”) may also be acceptable to scholars. This may also explain the degradation in performance in Figure 3, as the number of plausible completions rises in correlation with the length of the predicted span.

To address this, we conduct a direct manual evaluation of the top performing model’s predictions (M-BERT finetuned over Oracc) in a controlled environment, on both the original Akkadian, as well as its corresponding English translation. We begin by describing the experiment setup, which aims to cope with the inherent noise of human analysis in the MLM task, especially in an extinct language. Then, we discuss our findings, which show that the model provides sensible




Figure 4: Human evaluation interface for English (top) and transliterated Akkadian (bottom). Given the textual context from the tablet and a missing span of text (marked by red X's), the annotator decides whether each presented option is plausible. The options consist of the top three model predictions (marked in blue) and two controls: the original masked span (marked in yellow) and a randomly sampled span of text functioning as a distractor (marked in red).




Figure 5: Human evaluation results. The X-axis represents the number of signs (in Akkadian) or words (in English) in a predicted sequence, and the Y-axis represents the average number of model predictions that our human experts approved for the given predicted sequence. The upper error bars represent false negatives, where the gold sequence was labeled not plausible. The lower error bars represent false positives, where the distractor was labeled as plausible. We find that annotators tend to introduce false negatives, while they are less prone to falsely label distractors as plausible.

suggestions in most instances, while the comparison with English reveals that there is room for improvement, especially on longer sequences.

6.1 Experiment Setup: Coping with Noisy Human Evaluation

Our human evaluation of missing sign prediction in Akkadian was done by two of the authors, who are professional Assyriologists. They can read Akkadian at an academic level, and represent the users who work on cuneiform transliteration and may benefit from our model's predictions. Despite their unique expertise, they do not speak the language fluently like native speakers did, and the lan-

guage's natural variations over thousands of years makes the reading even more difficult.

To address this, we created an annotation scheme⁷ which evaluates the model's predictions and estimates the noise introduced in the annotation process. As exemplified in Figure 4, for each annotation instance, we generated 5 suggestions: 3 model predictions, the original masked term, and a distractor sequence that was randomly sampled from the Akkadian texts.⁸ The annotators observe the 5 suggestions in a randomized order, oblivious to which ones are model predictions. They are then required to mark each suggestion as either plausible or implausible, given the document's surrounding context.

Inserting the original masked sequence and the distractor enabled us to quantitatively estimate two sources of noise. First, the percentage of gold samples which were marked as incorrect reflects an *underestimation* of the model's ability as these are samples which in fact occurred in the original ancient texts, yet were ruled out by our experts. Similarly, the percentage of distractors marked as plausible reflects an *overestimation* of the model's performance.

By combining the estimated model accuracy (the percentage of the predictions marked as plausible) with both sources of noise, we can estimate a range in which the actual performance of the model may lie. Finally, for comparison with a high-resource language, we asked two fluent English speakers to annotate instances from the English translations of Oracc when predictions were generated by English BERT-base uncased model in the same experimental setup, as demonstrated at the top of Figure 4.

We conclude this part with an example human annotation and its corresponding analysis.

Annotation example. Consider the English annotation instance presented in Figure 4, and assume the annotator marked as plausible the following four items: the artificially introduced noise (“of Enlil’s”); two of the model predictions: “of the first”, “of the previous”; and the gold instance (“of the former”), while the remaining model prediction (“, your father”) is considered wrong by the human annotator. In which case, we compute

⁷Created with docanno (Nakayama et al., 2018).

⁸In case the model predicted the gold sequence, we added an additional model prediction, to ensure we always present 5 options.

the annotator’s quality assessment for this instance as $\frac{2}{3}$, while we record that they tend to *overestimate* the model performance, as they marked the artificial noise as plausible. Both of these metrics (accuracy and error estimation) are aggregated and averaged over the entire annotation.

6.2 Results

Each of our two annotators marked the top 5 model predictions for 70 different missing sequences, resulting in 700 binary annotations overall. 150 of these annotations were doubly annotated to compute agreement, overall finding good levels of agreement (.81κ for English and .79κ for Akkadian). These were drawn from royal inscriptions, as tagged in Oracc. This genre contains straight-forward, yet elaborate syntax and is well known by our annotators. We can make several observations based on Figure 5 which depicts the results of the human evaluation, based on the number of missing signs and the tested language (Akkadian versus English).

Our model’s Akkadian predictions are applicably useful... Per sequence of one or two signs, the annotators tended to accept on average at least one suggestion as plausible, while for three signs, they accepted on average about one suggestion per two sequences. From an applicative point of view, this functionality readily lends itself to aid transliteration of missing signs for sequences of such lengths, which constitute the majority (57%) of missing spans in Oracc.⁹

... yet performance degrades with the number of missing tokens. In Figure 5, we observe that the performance of the Akkadian model (in orange) degrades faster than the English model (in blue) the longer the predicted sequence gets. This indicates that the greedy decoding from a single span to multiple spans works better for English than for Akkadian. Designing a better decoding scheme is left as an interesting avenue for future work.

Humans tend to underestimate the model performance. By examining the assessments for the artificially introduced gold and distractor sequences we can estimate that the actual model performance may be higher than our experts estimated. We see that for both languages and in

⁹E.g., imagine a virtual keyboard auto-complete feature that suggests plausible completions in half of the cases.

most tested scenarios, our annotators were able to rule out the distractor, while they tended to also wrongly discarded the gold sequence (shown by the upper error bar), indicating that they may have also ruled out other plausible predictions made by the model.

7 Related Work

Most related to our work, Fetaya et al. (2020) designed an LSTM model which similarly aims to complete fragmentary sequences in Babylonian texts. They differ from us in two major aspects. First, they focus on small-scale highly-structured texts, for example, lists (parataxis), such as receipts or census documents (Jursa, 2004). Second, their LSTM model does not use multilingual pretraining, instead, it is trained on monolingual Akkadian data and its parameters are randomly initialized. In Section 5, we retrain their model on our data, showing that it underperforms on all genres compared to models which were pretrained using multilingual data, even in a zero-shot setting, further attesting to the valuable signal of multilingual pretraining in low-resource settings.

Other works have used Oracc and other Akkadian resources and may benefit from our language model for Akkadian. Jauhiainen et al. (2019) used Oracc for a shared task around language and dialect identification. Luukko et al. (2020) recently introduced a syntactic treebank for Akkadian over texts from Oracc, while Sahala et al. (2020) built a morphological analyzer using annotations from Oracc. Finally, Gordin et al. (2020) automatically transliterated Unicode cuneiform glyphs into the Latinized transliterated form.

Several recent works also noticed the cross-lingual transfer capabilities of M-BERT. Wu and Dredze (2019) and Conneau et al. (2020) found that M-BERT can successfully learn various NLP tasks in a zero-shot setting using cross-lingual transfer, pointing at the shared parameters across languages as the most important factor. Pires et al. (2019) showed that M-BERT is capable of zero-shot transfer learning even between languages with different writing systems.

8 Conclusions and Future Work

We presented a state-of-the-art model for missing sign completion in Akkadian texts, using multilingual pretraining and finetuning on Akkadian texts. Interestingly, we discovered that in such a low-

resource setting, the signal from pretraining may be more important than the finetuning objective. Evidently, a zero-shot model outperforms monolingual Akkadian models. Finally, we conducted a controlled user study showing the model’s potential applicability in aiding human editors.

Our work sets the ground for various avenues of future work. First, A more elaborate decoding scheme can be designed to mitigate the degradation of performance for longer masked sequences, for example by employing SpanBERT (Joshi et al., 2020) to represent the missing sequences during training and inference. Second, our findings suggest that an exploration of the specific utility of similar languages, e.g., Arabic or Hebrew, may yield improvements in missing sign prediction.

Acknowledgements

We thank Ethan Fetaya and Shai Gordin for insightful discussions and suggestions and the anonymous reviewers for their helpful comments and feedback. This work was supported in part by a research gift from the Allen Institute for AI.

References

- Ethan C. Chau, Lucy H. Lin, and Noah A. Smith. 2020. [Parsing with multilingual BERT, a small corpus, and a small treebank](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1324–1334, Online. Association for Computational Linguistics.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, 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.
- Ethan Fetaya, Yonatan Lifshitz, Elad Aaron, and Shai Gordin. 2020. [Restoration of fragmentary babylonian texts using recurrent neural networks](#). *Proceedings of the National Academy of Sciences*, 117(37):22743–22751.
- Shai Gordin, Gai Gutherz, Ariel Elazary, Avital Romach, Enrique Jiménez, Jonathan Berant, and Yoram Cohen. 2020. Reading akkadian cuneiform using natural language processing. *PloS one*, 15(10):e0240511.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- John Huehnergard. 2011. Introduction. *A Grammar of Akkadian*, pages xxiii–xlii.
- Tommi Jauhiainen, Heidi Jauhiainen, Tero Alstola, and Krister Lindén. 2019. [Language and dialect identification of cuneiform texts](#). In *Proceedings of the Sixth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 89–98, Ann Arbor, Michigan. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Michael Jursa. 2004. Accounting in neo-babylonian institutional archives: structure, usage, implications. *Creating Economic Order: Record-keeping, Standardization, and the Development of Accounting in the Ancient Near East, Bethesda*, pages 145–198.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mikko Luukko, Aleksi Sahala, Sam Hardwick, and Krister Lindén. 2020. [Akkadian treebank for early neo-assyrian royal inscriptions](#). In *Proceedings of the 19th International Workshop on Treebanks and Linguistic Theories*, pages 124–134, Düsseldorf, Germany. Association for Computational Linguistics.
- Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang. 2018. [doccano: Text annotation tool for human](#). Software available from <https://github.com/doccano/doccano>.
- A Leo Oppenheim. 2013. *Ancient Mesopotamia: portrait of a dead civilization*. University of Chicago Press.

- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Aleksi Sahala, Miikka Silfverberg, Antti Arppe, and Krister Lindén. 2020. [BabyFST - towards a finite-state based computational model of ancient babylonian](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3886–3894, Marseille, France. European Language Resources Association.
- Mike Schuster and Kaisuke Nakajima. 2012. [Japanese and korean voice search](#). In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152.
- Christopher Bromhead Fleming Walker. 1987. *Cuneiform*, volume 3. Univ of California Press.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). 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 833–844, Hong Kong, China. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, ukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.