

# Finetuning LLMs for EvaCun 2025 Token Prediction Shared Task

Josef Jon, Ondřej Bojar

Charles University, Faculty of Mathematics and Physics  
surname@ufal.mff.cuni.cz

## Abstract

In this paper, we present our submission for the token prediction task of EvaCun 2025. Our systems are based on LLMs (Command-R, Mistral, and Aya Expanse) fine-tuned on the task data provided by the organizers. As we only possess a very superficial knowledge of the subject field and the languages of the task, we simply used the training data without any task-specific adjustments, preprocessing, or filtering. We compare 3 different approaches (based on 3 different prompts) of obtaining the predictions, and we evaluate them on a held-out part of the data.

## 1 Introduction

The EvaCun token prediction shared task focuses on missing word restoration in languages originally written in cuneiform – Akkadian and Sumerian. The script is one of the earliest known forms of writing with a history spanning over 3000 years, evolving originally from the proto-cuneiform that was used for accounting and record keeping. One of the most challenging part of interpreting and translating Akkadian and Sumerian text is the polyvalence of cuneiform signs – a single sign can be used as a logogram, i.e. representing a whole word (which is further complicated by the fact that one symbol can represent many different possible words, and that Akkadian texts can contain Sumerian words, even though the languages are not otherwise related), or as a syllable (one sign can represent multiple syllables) or as a determinative that denotes a semantic category of the previous word (diety, person, place, etc.). In the context of the task, our work is greatly simplified by the fact that the task data are already interpreted and transliterated into the Latin alphabet instead of being in the original cuneiform script. As we do not have any knowledge of the languages of the task in our team, we pursued a purely engineering approach of finetuning 3 different LLMs

– Aya Expanse 8B (Dang et al., 2024), Command-R v0.1 34B (Cohere4AI team, 2025) and Mistral Small 3 24B (Mistral team, 2025) – on the task data, with 3 slightly different formulations of the problem. We offer our solution as a baseline to be compared with the more informed and task-specific approaches.

## 2 Related work

A more focused effort in NLP for languages written in cuneiform started only recently. A basis for all future work are databases and datasets like the Electronic Text Corpus of Sumerian Literature (Black et al., 2016), Cuneiform Digital Library Initiative, (CDLI contributors, 2025), CuneiML (Chen et al., 2023) and the Open Richly Annotated Cuneiform Corpus (Oracc Team, 2025).

Simmons et al. (2024) created a new corpus based on these previously released datasets that pairs digital Unicode transcription of cuneiform texts with their transliterations as well as a baseline system trained on this dataset to perform this task. Similarly, Gordin et al. (2020) present a method for automatic translation of Akkadian cuneiform. ? present an MT system for Sumerian with the final goal of an information retrieval pipeline for this language.

## 3 Methods

We fine-tune autoregressive LLMs with 3 different prompts to predict the masked word. A masked language model would be a more natural choice for this task, however causal (autoregressive) language modeling is currently a more popular approach with a larger selection of pretrained models.

### 3.1 Data preprocessing and prompts

The dataset provided by the organizers contains a list of tokens, each token accompanied by its document id, line number in the context of the doc-

Method	Prompt
All	Fill in the missing {language} words, masked by the [MASK] token. Output "WORDS:" and a comma-separated list of the missing words in original {language}: {masked_document}
One by one	Fill in the missing {language} word masked by the [MASK] token: {masked_document_with_unks}
Restore	Complete the missing {language} words masked by the [MASK] tokens and print out the restored document: {masked_document}

Table 1: Prompts for the three token prediction methods we compared.

ument, the word index on that line, language, and extra information, for example, the place where the tablet was found or the type of text. We only make use of the language, word order, and document id, i.e., we do not for example split the inputs into lines or use the additional information. In each document, we mask 15% randomly sampled words with a [MASK] token. For each document, we create at most 15 unique variants with different masks (fewer if the overall possible number of combinations for the given document length is lower). We frame the prediction task in three different ways. The model is shown the masked document in the prompt, and it is asked to:

- Produce a list containing all the original words corresponding to the masked predictions in the correct order (we call this method *All* further in the text).
- All [MASK] tokens except for one are replaced by an [UNK] token, and we ask the model to predict the original word for the single remaining [MASK] token. This is repeated for all [MASK] tokens in the masked document (*One by one*).
- We ask the model to output the full restored text of the masked document. We finetune separate models on each of these prompts (*Restore*).

The specific prompts are shown in Table 1. Our baseline approach, *All*, needs the least effort for data preprocessing and training and inference compute time. However, it could suffer from error propagation due to the autoregressive nature of the inference – the model bases the predictions on previously predicted words as well. *One by one* approach could mitigate this issue, as only one masked word is predicted for each example (others remain masked). *Restore* approach is based on the basic next-word prediction training objective for autoregressive models, but the decoding is complicated by the need for forcing the unmasked parts of

the text and keeping the word lengths the same for whole text for both masked and unmasked versions.

## 4 Experiments

We describe the experimental setup, hyperparameters and results in this section.

### 4.1 Data

The full training data from the organizers contains 913252 tokens in 22777 documents. We set aside 1% (227 documents) for the dev set (we filter out single word documents from the dev set). For our evaluation, we used a subset of this dev set containing 135 documents with 1500 different unique masked examples in total.

### 4.2 LLM finetuning

We finetune the pretrained models using QLoRA (Detmers et al., 2023). We experimented with 3 LLMs: *Command-R V0.1* (4-bit quantized, *CohereForAI/c4ai-command-r-v01-4bit*), *Aya Expansive 8B* and *Mistral Small 3 24B Instruct* (4-bit quantized, *unsloth/Mistral-Small-24B-Instruct-2501-unsloth-bnb-4bit*). We use the transformers (Wolf et al., 2020), peft and trl libraries for the training. We experimented with LoRA rank sizes 8, 16, 32, 64 and 92,  $\alpha = r/2$ . We finetuned the models by AdamW optimizer, with warmup ratio of 0.03 and learning rate  $lr = 2e - 4$ . We used batch sizes 40, 36 and 35 for Aya, Command-R and Mistral models, respectively. We trained on a heterogeneous cluster on a mix of Nvidia L40, A40 and H100 GPUs. We trained for a maximum of epochs, but the checkpoints we actually used for the prediction were from earlier parts of the training, as we describe in the results section.

### 4.3 Results

We sampled from the models with temperature  $t = 0.2$  to obtain the predictions (greedily, without the use of algorithms like beam search) and we measured the accuracy of the predictions on the held-out validation set, i.e the fraction of masked




Figure 1: Relative frequencies of the top-20 generated words for masked positions by best checkpoint for each prompt (*All*, *One by one*, *Restore* on top left, top right and bottom left, respectively) and of the reference masked words (bottom right).

	Aya Expanse	Command-R	Mistral
<b>All</b>	0.202	0.209	0.221
<b>One by one</b>	0.157	0.167	0.205
<b>Restore</b>	0.136	0.139	0.137
<b>Majority voting</b>	0.269 (0.377)		
<b>Most common word</b>	0.04		

Table 2: Accuracies of all combinations of models and prompts on held-out part of the task data. We also report top-3 accuracy for the majority voting in the parentheses. The final row shows the accuracy of predicting the most common word for the given language (based on the training data) for each masked position.

	Aya Expanse	Command-R	Mistral
<b>All</b>	6300 (0.75)	6300 (0.67)	5400 (0.55)
<b>One by one</b>	8100 (0.15)	5400 (0.10)	900 (0.02)
<b>Restore</b>	4500 (0.53)	9000 (0.96)	2700 (0.27)

Table 3: Number of updates (and the corresponding fraction of an epoch in parentheses) that the best-performing models were trained for.

positions where the missing word was predicted correctly out of all masked positions. For the *Restore* method, we generate the restored document in parts, by force decoding the known parts and only generating one complete word (possibly consisting of multiple subwords) for each [MASK] token (i.e. after force decoding the unmasked part of the document, we select the most probable subword that starts with a beginning of word symbol and generate next subwords until we reach another beginning of word subword, we discard this last subword and start with force decoding the known continuation again). We ensemble the results by majority voting, pick 60 best-performing checkpoints and select the most common prediction for each position.

We present the results of the 3 methods on the held-out validation set in Table 2. Overall, finetuning the Mistral models resulted in the best accuracies. However, the differences are not large and with a different choice of hyperparameters in the finetuning, we might see different ranking. From the methods point of view, *All* performed the best. Table 3 shows the number of steps and a corresponding fraction of an epoch that the best-scoring checkpoints were trained on. The final line shows the majority baseline – for each language, we only predict the most common word from the training data for all masked positions. For example, in

Akkadian, the most common word is the proposition *ina*, meaning *in, on, onto, at, to, from* and other possible meanings in compound expressions.

We also show the relative frequencies of the top-20 predicted words by each method and of the reference words in Figure 1. We see that while the list of top 10 words is similar to the reference list for all methods, the LLMs overestimate the probability (frequency) of the most popular words. This is a common issue in text-generating models. As a result, probabilities of less common words are underestimated – the methods generated 1567, 1217 and 3164 unique words for *All*, *One by one* and *Restore* respectively, while the reference contains 2317 unique tokens. We believe that the large number of unique tokens for the *Restore* method is caused by our prediction mechanism that ensures the same word length of both the prediction and the original text but can force the selection of suboptimal predictions as a fallback. Also, note that we did not disallow the generation of [MASK] token in the *Restore* method by mistake. This negatively affects the resulting accuracy of this method.

For the final test set submission, we ran the inference with the best 60 checkpoints on the test dataset and performed the majority voting to obtain top-3 predictions.

## 5 Conclusion

We finetuned various autoregressive LLMs on the token restoration task posed in 3 different ways. We show that the best single model can accurately predict 22.1% of masked tokens on our held-out dev set, while by combining predictions of multiple models by voting, we can reach 26.9% accuracy. However, there might be biases and aspects of the dataset like repetitiveness, which could lead to overestimating the real capabilities of our approach.

In the future, we plan to focus on the much more difficult task of direct translation of cuneiform script into English, either using Unicode transcriptions of the tablets, or a visual LLM to read the tablet photos directly.

## 6 Acknowledgments

This work was supported by Czech Ministry of Education, Youth and Sports (grant MŠMT OP JAK Mezisektorová spolupráce CZ.02.01.01/00/23\_020/0008518) and National Recovery Plan funded project MPO 60273/24/21300/21000 CEDMO 2.0 NPO. The

computational resources were provided by Ministry of Education, Youth and Sports of the Czech Republic Project Nr. LM2023062 LINDAT/CLARIAH-CZ.

## References

- J.A Black, Cunningham, Ebeling, FlückigerHawker, Robson, Taylor, and Zólyomi. 2016. Electronic text corpus of sumerian literature. <https://etcsl.orinst.ox.ac.uk/>. [Accessed 03-03-2025].
- CDLI contributors. 2025. *Cdli*. <https://cdli.mpiwg-berlin.mpg.de/>. [Online; accessed 2025-03-03].
- Danlu Chen, Aditi Agarwal, Taylor Berg-Kirkpatrick, and Jacobo Myerston. 2023. *Cuneiml: A cuneiform dataset for machine learning*. *Journal of Open Humanities Data*.
- Coherer4AI team. 2025. *Command r: Retrieval-augmented generation at production scale*. <https://cohere.com/blog/command-r>. [Online; accessed 2025-03-03].
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, Acyr Locatelli, Hangyu Lin, Dwarak Talupuru, Bharat Venkitesh, David Cairuz, Bowen Yang, Tim Chung, Wei-Yin Ko, Sylvie Shang Shi, Amir Shukayev, Sammie Bae, Aleksandra Piktus, Roman Castagné, Felipe Cruz-Salinas, Eddie Kim, Lucas Crawhall-Stein, Adrien Morisot, Sudip Roy, Phil Blunsom, Ivan Zhang, Aidan Gomez, Nick Frosst, Marzieh Fadaee, Beyza Ermis, Ahmet Üstün, and Sara Hooker. 2024. *Aya expanse: Combining research breakthroughs for a new multilingual frontier*. *Preprint*, arXiv:2412.04261.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. *Qlora: Efficient finetuning of quantized llms*. *Preprint*, arXiv:2305.14314.
- 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):1–16.
- Mistral team. 2025. *Mistral small 3*. <https://mistral.ai/news/mistral-small-3>. [Online; accessed 2025-03-03].
- Oracc Team. 2025. *Oracc*. <http://oracc.org>. [Accessed 03-03-2025].
- Cole Simmons, Richard Diehl Martinez, and Dan Jurafsky. 2024. *SumTablets: A transliteration dataset of Sumerian tablets*. In *Proceedings of the 1st Workshop on Machine Learning for Ancient Languages (ML4AL 2024)*, pages 192–202, Hybrid in Bangkok,

Thailand and online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2020. [Huggingface's transformers: State-of-the-art natural language processing](#). *Preprint*, arXiv:1910.03771.