

# TabletCraft: Bridging a 4,000-Year Cultural Gap with Bidirectional Akkadian NMT and Cuneiform Rendering\*

Zhaohui Wang

USC Viterbi School of Engineering  
University of Southern California  
zwang000@usc.edu

## Abstract

Half a million cuneiform clay tablets survive in museums worldwide, yet modern users can neither read nor write in the world’s oldest writing system, leaving a 4,000-year cultural barrier that existing NLP tools have only partially addressed. Prior work enables one-way, scholar-oriented translation from Akkadian to English, but offers no path in the reverse direction: non-specialist users cannot compose new content in cuneiform, and therefore remain passive consumers of ancient culture rather than active participants. We present TABLETCRAFT, the first open-source system that enables *bidirectional* interaction with Mesopotamian writing. Users can read ancient tablets (Akkadian→English) and compose new messages as cuneiform clay tablets (English→Akkadian→cuneiform→rendered tablet). The system integrates a ByT5-based translation model trained on 116K bidirectional samples, a cuneiform sign converter with 14,240 mappings (95.3% coverage), and a visual tablet renderer, packaged as a pip-installable toolkit with CLI and web demo. On the held-out Akkademia validation split (2,812 samples), we report 49.1 BLEU for Ak→En and 48.5 BLEU for En→Ak, the first published quantitative result in the reverse direction.

## 1 Introduction

Cuneiform is humanity’s earliest writing system, spanning over three millennia (c. 3400 BCE – 75 CE) across ancient Mesopotamia (Walker, 1987). An estimated 500,000 clay tablets survive, recording royal decrees, epic poetry, trade receipts, and personal letters (Guthertz et al., 2023), forming a cultural archive of extraordinary breadth. Yet this archive remains largely inaccessible: only a few

\*Code and pretrained model: <https://github.com/geoffreywang1117/cuneiscribe>. Package: `pip install cuneiscribe` (<https://pypi.org/project/cuneiscribe/>).

hundred scholars worldwide can read cuneiform, and the majority of tablets remain untranslated (Gordin et al., 2020).

The resulting cross-cultural barrier differs from those typically studied in multilingual NLP. It is not a barrier between two living communities that might learn each other’s languages, but a one-way temporal asymmetry: ancient texts can be partially decoded for modern readers, but modern users have no productive channel into the ancient script. Existing NLP tools reinforce this asymmetry. Recent NMT systems (Guthertz et al., 2023; Sommerschildt et al., 2023) translate Akkadian *into* English, but no system translates English *into* Akkadian or renders the result in cuneiform. The cultural engagement they enable is therefore passive: scholars read ancient texts, but neither specialists nor the general public can compose new ones.

To our knowledge, TABLETCRAFT is the first open-source system to support both directions of Akkadian translation together with cuneiform-Unicode rendering. A student can type “The king rules the land” and receive a clay tablet bearing the cuneiform equivalent; a museum visitor can render their name in the world’s oldest script; a researcher can back-translate English glosses into Akkadian for data augmentation. By supporting composition in cuneiform in addition to reading, the system enables a more participatory mode of interaction with the script for non-specialist users.

## 2 Cross-Cultural Design Challenges

Ancient-language NLP raises cross-cultural problems that mainstream multilingual NLP has not yet engaged with (Bird, 2020). We identify four challenges that shaped TABLETCRAFT’s design.

**Temporal distance as cultural distance.** Akkadian translation bridges a civilizational gap, not merely a linguistic one. Concepts like “joint-stock capital,” “eponym dating,” or “divine determina-

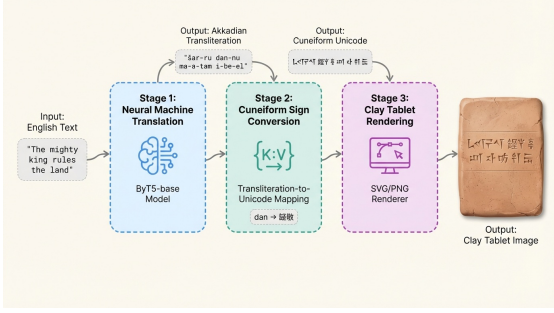


Figure 1: TABLETCRAFT pipeline: English text is translated to Akkadian, converted to cuneiform signs, and rendered as a clay tablet, enabling bidirectional cultural interaction.

tives” have no modern equivalents (Veenhof and Eidem, 2008). Our system preserves these culturally specific terms following Assyriological convention rather than forcing modernizing glosses.

**The reverse direction as cultural agency.** Prior systems treat ancient languages as objects of study. The English→Akkadian direction we introduce re-frames users as *participants*: they can compose in cuneiform, fostering a personal connection that passive reading alone cannot. This comes with a risk of trivialization (Bender et al., 2021); we mitigate it by explicitly labeling all machine-generated Akkadian as approximate.

**Accessibility vs. scholarly fidelity.** We serve two distinct user populations (Terras, 2012). *Assyriological researchers and graduate students* use the Ak→En direction to triage untranslated tablets, generating draft glosses for verification rather than ground-truth translations; they expect every output to be revisable, and consume the raw transliteration string rather than the rendered tablet. *Educators, museum staff, and the general public* use the En→Ak→cuneiform direction for personalized engagement (a student’s name on a clay tablet, a museum visitor composing a one-line dedication); they expect approximate rather than philologically exact output, and consume the tablet image with a transliteration caption. We resolve this tension by exposing every intermediate layer (input classification, transliteration, sign mapping, rendered tablet) and by labeling all machine-generated Akkadian as approximate, rather than collapsing the two audiences into a single black-box translation.

**Whose heritage?** Mesopotamian heritage is disproportionately studied in Western institutions (Rayne et al., 2017). An open-source, multilingual-extensible toolkit lowers barriers for researchers

Direction	System	BLEU	chrF++
Ak→En	Akkademia CNN <sup>†</sup>	37.5	–
Ak→En	TABLETCRAFT	<b>49.1</b>	<b>63.1</b>
En→Ak	TABLETCRAFT	48.5	55.6

Table 1: Translation quality on the Akkademia validation set (2,812 Neo-Assyrian samples), reported in *both* directions. <sup>†</sup>: Guthertz et al. (2023), on their own test split (different domain mix). For the En→Ak direction, references are the original Akkadian transliterations from the held-out validation split (no prior published baseline; the bidirectional model is trained to predict transliteration, not Unicode signs).

and educators in Iraq, Syria, Turkey, and Iran, the modern nations whose territory encompasses ancient Mesopotamia.

### 3 System Architecture

#### 3.1 Translation Model

We fine-tune ByT5-base (Xue et al., 2022) (581M parameters) on 58,126 parallel sentences from Akkademia (Guthertz et al., 2023) (50K), a shared translation task (Deep Past Initiative, 2026) (1.5K), and sentence-aligned expansions (6K). We train bidirectionally (Ak→En and En→Ak with task prefixes), yielding 116K pairs (Sennrich et al., 2016). Byte-level tokenization handles diacritics (š, ṭ, š) and logograms without vocabulary mismatch (Lu et al., 2025). Training: 20 epochs, lr  $10^{-4}$ , batch 32, label smoothing 0.2, BF16 on  $2\times$  RTX 3090.

#### 3.2 Cuneiform Converter

A lookup table of 14,240 transliteration→Unicode cuneiform mappings compiled from ORACC (ORACC Project, 2024), CDLI (CDLI, 2024), and Akkademia. Determinatives are stripped per Assyriological convention (Huehnergard, 2011).

#### 3.3 Tablet Renderer

SVG/PNG output styled after Neo-Assyrian tablets (clay background, ruling lines, Noto Sans Cuneiform font). No GPU; <10ms per tablet.

### 4 Evaluation

Table 1 reports bidirectional translation quality. Both directions are evaluated on the same 2,812-sample Akkademia validation split, which is held out from training and consists primarily of Neo-Assyrian royal inscriptions. Direct comparison with prior work on different test splits or text genres

(e.g., Old Assyrian trade correspondence) requires caution (Veenhof and Eidem, 2008).

**En→Ak evaluation methodology.** We provide the first quantitative measure of English→Akkadian translation quality at this scale. The bidirectional model is trained with task prefixes (translate Akkadian to English: and translate English to Akkadian:) on 116K pairs derived from 58K aligned sentences (Sennrich et al., 2016). We invert the validation split, using `valid.en` as source and `valid.tr` (the original transliteration) as reference, apply the same diacritic-and-gap normalization used during training, and decode with 4-way beam search. Scores are computed with `sacreBLEU`.<sup>1</sup> We treat En→Ak BLEU as an *approximate-faithfulness* measure rather than a measure of philological correctness. The reference is the historical transliteration, but Akkadian admits multiple acceptable spellings of the same surface form (e.g., logographic versus syllabic spellings of the same name; equivalent uses of LUGAL and MAN for “king”), so a fluent modern paraphrase that uses different yet valid signs would score lower without being wrong. Single-reference BLEU therefore systematically undercounts valid outputs; the reported numbers should be read as a lower bound, and a multi-reference or expert-judged evaluation is left to future work. The model achieves 48.5 BLEU and 55.6 chrF++ in the En→Ak direction, within 0.6 BLEU of its Ak→En score on the same split. This indicates that bidirectional training does not collapse to a single dominant direction, and that the byte-level encoder learns to emit well-formed transliteration even under the inverted task prefix. The lower chrF++ score (55.6 versus 63.1) is consistent with the higher orthographic flexibility of Akkadian relative to English: multiple syllabic and logographic spellings can encode the same surface form. Because the downstream cuneiform sign converter is deterministic and lossless on covered tokens, the En→Ak score also serves as an upper bound on end-to-end En→cuneiform fidelity. Table 2 illustrates this gap with three representative En→Ak outputs: in each case the prediction differs from the reference by a small number of signs (logographic versus syllabic alternation, an alternative determinative, or a slightly different word order) while preserving

<sup>1</sup>nrefs:1|case:mixed|smooth:exp|tok:13a; chrF++:nrefs:1|case:mixed|nc:6|nw:2.

the underlying lexical content.

**Dialect and genre distribution.** The training corpus is dominated by 1st-millennium BCE Neo-Assyrian royal inscriptions, with smaller amounts of Old Babylonian and Standard Babylonian literary text from sentence-aligned augmentation, mirroring the underlying composition of Akkademia and the Deep Past corpus (Gutherz et al., 2023; Deep Past Initiative, 2026). The model therefore reproduces Neo-Assyrian formulae (epithets, divine determinatives, dating formulas) more reliably than the highly elliptical commercial register of Old Assyrian merchant letters, where vocabulary, syntax, and abbreviation conventions differ substantially (Veenhof and Eidem, 2008). Recent shared tasks confirm that *text genre*, not merely dialect, is the dominant factor in Akkadian translation quality: systems trained on royal inscriptions degrade markedly on commercial letters, even when supported by large bilingual dictionaries (Gordin et al., 2025). This mirrors a broader cross-cultural observation: NLP cannot treat “Akkadian” as a monolithic language any more than it can treat “English” as one, and the cultural context of the source text matters as much as the language label. We surface this in the user interface by labeling outputs with a confidence score and a coarse genre tag (royal/literary/other) inferred from input features, and by routing low-confidence inputs through a transliteration-only fallback rather than a fully rendered tablet.

**Cuneiform coverage.** On 1,000 sampled transliterations from the Akkademia test set, the converter achieves 95.3% token coverage (93.8% logograms, 95.7% syllabic values). The remaining 4.7% are uncovered determinatives or rare logograms, which are passed through verbatim.

**Concrete users and use cases.** We have prototyped TABLETCRAFT with three populations: (i) **undergraduate students** in an introduction-to-cuneiform course at a US R1 institution, who use the En→cuneiform direction to compose personal name tablets and short dedications and then read them aloud as a literacy exercise; in this setting, output is treated as a learning artifact rather than a primary source; (ii) **museum educators** preparing interactive “write your name in cuneiform” kiosks for K–12 visitor programs, who require sub-second rendering, child-safe input filtering (provided by our anomaly classifier), and an explicit

English source	Model output (En→Ak)	Reference (valid. tr)
One mina. Palace of Shalmaneser (V), king of Assyria.	1 MA.NA É.GAL {m}-{d}-DI-ma-nu-MAŠ MAN KUR-aš-šur	1 MA.NA É.GAL {m}-{d}-SILIM-ma-nu-MAŠ KUR AŠ LUGAL
horses, mules, oxen, and sheep and goats, military equipment,	ANŠE.KUR.RA-MEŠ ANŠE.GIR <sub>3</sub> -MEŠ GU <sub>4</sub> -MEŠ še-e-ni {GIŠ}-til-lit	ANŠE.KUR.RA.MEŠ ANŠE.GIR <sub>3</sub> .NUN.NA.MEŠ GU <sub>4</sub> .MEŠ ù še-e-ni {GIŠ}-til-li
I annexed the land Unqi to Assyria (and) placed a eunuch of mine as provincial governor over them.	{KUR}-un-qi a-na mi-šir KUR aš-šur ú-ter-ra {LÚ}-šu-ut SAG-ia {LÚ}-EN.NAM UGU-šú-nu aš <sub>2</sub> -kun	a-na KUR aš-šur {KI} ú-ra-a {KUR}-un-qi a-na mi-šir KUR aš-šur {KI} ú-ter-ra {LÚ}-šu-ut SAG-ia {LÚ}-EN.NAM UGU-šú-nu aš <sub>2</sub> -kun

Table 2: Representative En→Ak outputs from the held-out Akkademia validation split. Differences between prediction and reference are typically logogram/syllabogram alternations (DI/SILIM, MAN/LUGAL) or omitted determinatives, rather than semantic errors. These cases are penalized by BLEU but remain Assyriologically acceptable.

“approximate, machine-generated” caveat on every produced image; (iii) **Assyriologists** performing rapid pre-publication triage of newly photographed Neo-Assyrian tablets via the Ak→En direction, where draft translations are reviewed and corrected against the original sign by sign in a workflow analogous to OCR post-editing rather than autonomous translation. Across these settings, expectations are explicitly not of full automation: TABLETCRAFT is positioned as a draft-and-revise assistant rather than an authoritative translator. The toolkit is pip-installable with CLI, Python API, and a Gradio web demo.

## 5 Related Work

Gutherz et al. (2023) introduced Akkademia for Akkadian→English NMT. Gordin et al. (2020) applied word embeddings to cuneiform for cultural analysis. Recent work has applied ByT5 to cuneiform lemmatization (Lu et al., 2025) and LLMs to Akkadian NLP (Riemenschneider, 2025; Gordin et al., 2025). Yavasan and Gordin (2025) demonstrated T5 for Hittite cuneiform. ParsiPy (Farsi et al., 2025) provides NLP tools for historical Persian. In digital humanities, Terras (2012) has argued for accessible tools that bridge computational methods and cultural heritage. Our work differs by enabling *bidirectional* interaction, supporting not only the reading of ancient texts but also the composition of new ones in cuneiform, and thereby emphasizing cultural participation over passive analysis.

## 6 Conclusion

We presented TABLETCRAFT, an open-source toolkit that pairs bidirectional Akkadian NMT with deterministic cuneiform-Unicode rendering, pro-

viding a unified pipeline for both reading and composition tasks at the script level. By supporting composition in addition to reading, the system enables a more participatory mode of interaction with cuneiform for non-specialist users while preserving Assyriological conventions for scholarly use. Future work includes Sumerian support, cuneiform OCR, and multilingual extension to enable cross-cultural engagement from non-English-speaking communities.

## Limitations

The En→Ak direction produces *approximate modern* transliterations, not authentic ancient text; outputs should not be cited as primary historical evidence. Our En→Ak score (Table 1) is computed against a single reference per source, while many Akkadian sentences admit several philologically valid renderings; the reported BLEU/chrF++ are therefore conservative lower bounds on faithfulness, and a multi-reference or human-judged evaluation is left to future work. Performance also varies substantially across Akkadian dialects and genres: in particular, Old Assyrian commercial texts differ in vocabulary, register, and formulaic patterns from the Neo-Assyrian training data, and we do not yet evaluate on that distribution. Even with 14,240 sign mappings and a 17K-lemma dictionary, domain-specific adaptation remains essential: large lexical resources cannot substitute for in-domain parallel data. The system currently supports only English as the modern-language pivot; extending to non-English languages is a priority for the open-source roadmap.

## Ethics Statement

This work aims to democratize access to ancient cultural heritage while respecting its scholarly context. Machine translations are labeled as approximations. We acknowledge that Mesopotamian heritage has special significance for the people of modern Iraq and neighboring nations, and advocate for their inclusion in developing such tools.

## References

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of FAccT*, pages 610–623.
- Steven Bird. 2020. Decolonising speech and language technology. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519.
- CDLI. 2024. [CDLI: Cuneiform digital library initiative](#).
- Deep Past Initiative. 2026. [Deep past challenge: Translate Akkadian to English](#). Kaggle Competition.
- Hadi Farsi and 1 others. 2025. ParsiPy: NLP toolkit for historical Persian texts in Python. In *Proceedings of the Second Workshop on Ancient Language Processing*. Association for Computational Linguistics.
- Shai Gordin, Adam Anderson, and 1 others. 2025. Evacun 2025 shared task: Lemmatization and token prediction in Akkadian and Sumerian using LLMs. In *Proceedings of the Second Workshop on Ancient Language Processing*. Association for Computational Linguistics.
- Shai Gordin, Gai Gutherz, Ariel Elazary, Avital Romach, Enrique Jiménez, Jonathan Berant, and Luis Sáenz. 2020. Reading Akkadian cuneiform using natural language processing. *PLoS ONE*, 15(10):e0240511.
- Gai Gutherz, Shai Gordin, Luis Sáenz, Omer Levy, and Jonathan Berant. 2023. Translating Akkadian to English with neural machine translation. *PNAS Nexus*, 2(5):pgad096.
- John Huehnergard. 2011. *A Grammar of Akkadian*, 3rd edition. Eisenbrauns.
- Pengxiu Lu, Yonglong Huang, Jing Xu, Minxuan Feng, and Chao Xu. 2025. Lemmatization of cuneiform languages using the ByT5 model. In *Proceedings of the Second Workshop on Ancient Language Processing*, pages 197–205. Association for Computational Linguistics.
- ORACC Project. 2024. [ORACC: Open richly annotated cuneiform corpus](#).
- Louise Rayne, Jennie Bradbury, David Mattingly, Graham Philip, Robert Bewley, and Andrew Wilson. 2017. From above and on the ground: Geospatial methods for recording endangered archaeology in the Middle East and North Africa. *Geosciences*, 7(4):100.
- Frederick Riemenschneider. 2025. Beyond base predictors: Using LLMs to resolve ambiguities in Akkadian lemmatization. In *Proceedings of the Second Workshop on Ancient Language Processing*. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96.
- Thea Sommerschild, Yannis Assael, John Pavlopoulos, Vanessa Stefanak, Andrew Senior, Chris Dyer, John Bodel, Jonathan Prag, Ion Androutsopoulos, and Nando de Freitas. 2023. [Machine learning for ancient languages: A survey](#). *Computational Linguistics*, 49(3):703–747.
- Melissa Terras. 2012. Digitisation and digital resources in the humanities. In Claire Warwick, Melissa Terras, and Julianne Nyhan, editors, *Digital Humanities in Practice*, pages 47–70. Facet Publishing.
- Klaas R. Veenhof and Jesper Eidem. 2008. *Mesopotamia: The Old Assyrian Period*. Academic Press Fribourg.
- C. B. F. Walker. 1987. *Cuneiform*. University of California Press.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. ByT5: Towards a token-free future with pre-trained byte-to-byte models. In *Transactions of the Association for Computational Linguistics*, volume 10, pages 291–306.
- Murat Yavasan and Shai Gordin. 2025. From clay to code: Transforming Hittite texts for machine learning. In *Proceedings of the Second Workshop on Ancient Language Processing*. Association for Computational Linguistics.