

RAG-Enhanced Neural Machine Translation of Ancient Egyptian Text: A Case Study of THOTH AI

So Miyagawa

Research Center for West Asian Civilization

University of Tsukuba

1-1-1 Tennodai, Bldg. B506

Tsukuba, Ibaraki 305-0006, Japan

miyagawa.so.kb@u.tsukuba.ac.jp

Abstract

This paper demonstrates how Retrieval-Augmented Generation (RAG) significantly improves translation accuracy for Middle Egyptian, a historically rich but low-resource language. We integrate a vectorized Coptic-Egyptian lexicon and morphological database into a specialized tool called THOTH AI. By supplying domain-specific linguistic knowledge to Large Language Models (LLMs) like Claude 3.5 Sonnet, our system yields translations that are more contextually grounded and semantically precise. We compare THOTH AI against various mainstream models, including Gemini 2.0, DeepSeek R1, and GPT variants, evaluating performance with BLEU, SacreBLEU, METEOR, ROUGE, and chrF. Experimental results on the coronation decree of Thutmose I (18th Dynasty) show that THOTH AI's RAG approach provides the most accurate translations, highlighting the critical value of domain knowledge in natural language processing for ancient, specialized corpora. Furthermore, we discuss how our method benefits e-learning, digital humanities, and language revitalization efforts, bridging the gap between purely data-driven approaches and expert-driven resources in historical linguistics.

1 Introduction

Ancient Egyptian is an Afro-Asiatic language dating back five millennia, encompassing multiple historical phases—Old Egyptian, Middle Egyptian, Late Egyptian, Demotic, and Coptic—as well as a complex set of writing systems (hieroglyphic, hieratic, demotic, and the Greek-based Coptic script). Despite its linguistic and cultural importance, it remains a low-resource language for natural language processing (NLP) tasks, primarily due to limited digitized parallel corpora and the intricate

orthographic and grammatical features of its scripts. While large-scale neural networks and Large Language Models (LLMs) have revolutionized machine translation in high-resource languages, these models tend to underperform when domain-specific data are scarce.

Retrieval-Augmented Generation (RAG) is an emerging strategy that mitigates data scarcity by pairing LLMs with external knowledge repositories. Rather than relying solely on the implicit knowledge encoded in a model's parameters, RAG injects relevant external information—such as specialized lexicons, dictionaries, and grammatical annotations—directly into the model's prompt. In this way, the model's generative process is “grounded” in domain knowledge it might otherwise lack. Our research aims to show how RAG-based methods can achieve substantial improvements in translating Middle Egyptian texts, focusing on the coronation decree of the 18th Dynasty pharaoh Thutmose I.

We developed THOTH AI, an interactive translation system that unifies Claude 3.5 Sonnet with a vectorized lexicon curated from the Comprehensive Coptic Lexicon (Burns et al., 2020), produced by the Thesaurus Linguae Aegyptiae project, the Coptic SCRIPTORIUM project (Schroeder and Zeldes, 2016), and the KELLIA project. The system leverages the Dify platform's default vectorizer to embed specialized lexical data, enabling instant retrieval of morphological, semantic, and etymological details across Ancient Egyptian's historical stages. THOTH AI then passes these retrieved entries to Claude 3.5 Sonnet for a RAG-enhanced translation. To measure effectiveness, we compare THOTH AI's performance against state-of-the-art models such as Gemini 2.0, DeepSeek R1, and GPT variants (GPT 4o, GPT o1 Pro, GPT o3-mini-high). Consis-

tent improvements in BLEU, SacreBLEU, METEOR, ROUGE, and chrF scores underscore the benefits of domain-specific retrieval.

Finally, we discuss how this RAG-based approach, when integrated with OCR tools for handling hieratic or hieroglyphic script images, can transform e-learning by automatically offering morphological insights and dictionary lookups to students. Moreover, we highlight its value for digital humanities and Coptic-language revival, showcasing a method that is equally useful for academically trained Egyptologists and broader communities seeking to engage with Egypt’s ancient legacy.

2 Background and Related Work

2.1 Middle Egyptian as a Low-Resource Language

Although Middle Egyptian (see Table 1) emerged around the 21st–17th centuries BCE and continued in use (often in administrative or religious texts) for more than a millennium, modern NLP research on it remains limited. Most mainstream NLP resources focus on well-documented languages with extensive digital corpora. In contrast, Middle Egyptian scholarship frequently relies on manual philological analysis, with only partial or inconsistent digitization of core texts (Miyagawa and Kawai, 2024).

Standard neural machine translation (NMT) systems often require large volumes of parallel data to train robust models. Middle Egyptian’s severe data shortage means few large-scale NMT solutions exist. Also, morphological and orthographic complexity compounds the data challenge: for instance, many texts feature ideograms, phonograms, and determinatives in hieroglyphic writing. A typical approach is to transliterate these forms into a Latin-based system, which enables computational handling but can obscure nuances if the transliteration conventions differ or if morphological boundaries are not clearly marked.

2.2 Challenges in Ancient and Coptic Studies

Coptic, the final stage of Ancient Egyptian, has experienced a minor revival among certain communities and in scholarly domains (Miyagawa, 2024a,b; Saeed et al., 2024). Although

it is written in a modified Greek alphabet, its usage data is still sparse, mostly liturgical documents or specialized dictionaries (Feder et al., 2018). Tools like Coptic SCRIPTORIUM have emerged, providing annotated corpora, but advanced tasks such as machine translation, morphological tagging, or dictionary linking still pose significant challenges.

From a linguistic standpoint, bridging Middle Egyptian and Coptic data demands detailed knowledge of phonological and morphological evolution. Many forms of Coptic can be traced back to earlier Egyptian stages, but the correspondences are not always transparent. For instance, certain hieroglyphic forms converge into a single Coptic lemma, while other forms diverge or disappear altogether.

2.3 Emergence of RAG for Low-Resource Languages

Recent years have seen increased interest in Retrieval-Augmented Generation (RAG) (Gao et al., 2024), which addresses data scarcity by pairing LLMs with an external vector database. Instead of hoping an LLM has memorized a wide range of rare or archaic lexemes, RAG retrieves relevant dictionary entries or parallel texts to supply domain knowledge explicitly. This approach has shown promise in various specialized domains, from law to biomedical text. In the context of Ancient Egyptian, RAG can fetch morphological notes, definitions, and exemplars that strongly inform the generative process, significantly improving accuracy and reducing hallucinations (Enis and Hopkins, 2024).

3 Methodology

3.1 Source Text Selection and Experimental Scope

To highlight RAG’s impact on translation quality, we selected a classical text: the coronation decree of Thutmose I, an 18th Dynasty pharaoh (ca. 1504–1492 BCE). The text, found in Sethe (1927), contains phrases that blend religious, administrative, and formulaic elements common in official inscriptions. This text is challenging enough to expose the limitations of general-purpose LLMs but still comprehensible enough to have available reference translations (e.g., de Buck 1948; Nederhof 2023).

Table 1: Historical Stages of the Egyptian Language (Miyagawa and Kawai, 2024, 70), based on (Kammerzell, 2000, 97)

Stage	Period	Script
Pre-Old Egyptian	32nd–27th c. BCE	Early Hieroglyphs
Old Egyptian	27th–21st c. BCE, 8th c. BCE (archaic)	Hieroglyphs, Hieratic
Middle Egyptian	23rd c. BCE–4th c. CE	Hieroglyphs, Hieratic
Late Egyptian	14th–7th c. BCE	Hieroglyphs, Hieratic
Demotic	8th c. BCE–5th c. CE	Demotic
Coptic	3rd c. CE–21st c.	Coptic script

All tested systems received the same portion of transliterated Middle Egyptian text, ensuring a fair comparison. We used [Nederhof’s](#) translation as the gold-standard reference for quantitative scoring. Our experiments focused on how each model handles archaic vocabulary, honorific epithets, morphological markers, and elliptical constructions typical of Middle Egyptian.

3.2 Models Evaluated

We compared eight different large language models or variants to see how well they translated Middle Egyptian. The first model was THOTH AI (RAG-Enhanced), our custom system built upon Claude 3.5 Sonnet but further enhanced by retrieving specialized lexicon entries. In contrast, Claude 3.5 Sonnet (baseline) was tested in its raw state, without domain-specific retrieval. We also included Gemini 2.0 in two different modes (Pro and Flash Thinking), both designed to offer advanced context reasoning. Another competitor was DeepSeek R1, a smaller model trained with a focus on low-resource languages, although not explicitly engineered for Ancient Egyptian. Finally, we examined three GPT-based variants (GPT 4o, GPT o1 Pro, and GPT o3-mini-high), each providing different parameter scales and pretraining data coverage. Altogether, this diversity of models allowed us to evaluate the effect of specialized retrieval in contrast to a variety of LLM architectures and capabilities.

3.3 Vectorizing the Comprehensive Coptic Lexicon

A crucial aspect of our RAG setup is the Comprehensive Coptic Lexicon, including etymological information, which gathers lexical, morphological, and historical data spanning ev-

ery stage of Ancient Egyptian up to Coptic. This repository contains a wide variety of information, including etymological relationships across Old Egyptian, Middle Egyptian, Demotic, and Coptic, as well as morphological fields such as nominal forms, suffix conjugations, and determinatives that highlight how certain words evolve over time. Some entries also document subtle semantic distinctions that emerged or disappeared across different chronological layers of the language.

To integrate this information into our translation pipeline, we employed Dify’s default vectorizer to produce vector embeddings of these lexicon entries. Each entry’s lemma forms, glosses, and morphological details were embedded into a shared semantic space. When a user inputs a snippet of transliterated Middle Egyptian, THOTH AI creates vector embeddings for those tokens and calculates their similarity to the embedded lexicon data. Once the system identifies the top matching entries, it appends them to the prompt before passing it to the LLM. This way, archaic terms or unusual morphological constructs are tied to their relevant definitions and historical annotations, effectively giving the LLM a “cheat sheet” of domain-specific knowledge.

3.4 THOTH AI Pipeline

We built THOTH AI on the open-source Dify platform,¹ which supports custom vector search modules and integrates seamlessly with LLM APIs. In the final system, there are several core steps in the workflow.

First, if the user has only scanned imagery of hieroglyphic or hieratic texts, we have an optional OCR module that can produce translit-

¹<https://github.com/langgenius/dify> (accessed on March 24, 2025)

erated text. Although OCR for ancient scripts is still experimental, it can provide a useful first pass for typed text. Next, THOTH AI takes this transliterated text and divides it into smaller segments if it is particularly lengthy. For each segment, the system calculates vector embeddings and conducts a similarity search in the Comprehensive Coptic Lexicon database. It retrieves the top few entries that match or partially match the user’s input, focusing on lexical, morphological, or historical data relevant to those segments.

These retrieved items are then concatenated directly with the user’s input so that Claude 3.5 Sonnet is not generating a translation in isolation. Instead, the model receives explicit domain knowledge that clarifies archaic or polysemous words. Finally, the model’s translation and short morphological notes are displayed on the THOTH AI interface (see Figure 3). This interface clearly labels the user’s input, the final translation, and a list of dictionary or lexicon entries that inform the translation process.

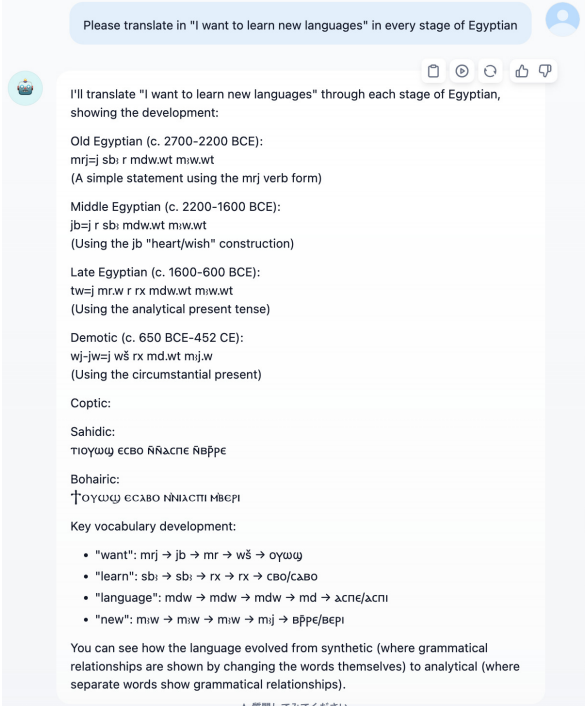


Figure 1: User interface of THOTH AI

4 Evaluation and Results

4.1 Quantitative Metrics

We measured translation quality in multiple ways to capture both exact *n*-gram overlaps and broader semantic fidelity. We used the BLEU (Papineni et al., 2002) and SacreBLEU (Post, 2018) metrics to compare *n*-gram overlaps against a reference translation, thus assessing how well each model captured the literal content of the Middle Egyptian text. Beyond *n*-gram precision, we relied on METEOR (Banerjee and Lavie, 2005), which considers synonyms and stems, and ROUGE (Lin, 2004), often utilized in text summarization but valuable here for evaluating recall of key phrases. Finally, we employed chrF (Popović, 2017, 2015), a character-level F-score metric especially suitable for languages that exhibit frequent and subtle morphological changes. All these metrics used Nederhof (2023)’s English rendition as the reference.

4.2 Test Data and Reference

For consistency, every system translated the same portion of Thutmose I’s coronation decree, specifically No. 30 in Sethe (1927), taking its Latin transliteration from de Buck (1948). The text is dense with formal epithets, references to gods, and references to the pharaoh’s lineage. Because these expressions can be formulaic, they serve as an ideal stress test for LLM-based translators. If an LLM has never seen specific epithets or morphological forms, it may guess incorrectly or omit them, thereby reducing its overall accuracy.

4.3 Overall Scores

Table 2 summarizes our main findings. As shown, THOTH AI, the RAG-based system, yields the highest BLEU score (0.354) among all tested models. It also leads in metrics like ROUGE and METEOR, underlining that RAG fosters not only literal fidelity (as reflected in BLEU) but also coverage of keywords and morphological consistency (as indicated by ROUGE and chrF). Claude 3.5 Sonnet, used as a standalone baseline, achieves the second-best BLEU (0.325), which is close but still notably behind THOTH AI. Gemini 2.0 Pro ranks third (0.288 BLEU), showing some promise but lacking the specialized retrieval

that helped THOTH AI excel. Other systems such as DeepSeek R1 and the GPT-based variants (GPT 4o, GPT o1 Pro, GPT o3-mini-high) produce more modest scores, presumably because they do not incorporate domain-specific references during generation.

4.4 Qualitative Observations

The quantitative results in Table 2 align with our qualitative observations during a manual review of the translations shown in Table 3. First, we found that RAG helps particularly with rare or archaic lexemes: for example, certain Middle Egyptian expressions that refer to the king’s divine roles or that mention obscure place names. Without retrieval, some models simply substitute placeholders or produce incomplete translations. THOTH AI consistently retrieved the correct glosses.

We also observed a higher morphological accuracy when the RAG-based approach provided relevant dictionary entries detailing suffix pronouns or determinatives. In non-RAG models, these morphological items often caused confusion, with the result that entire clauses might be mistranslated. Finally, using the vectorized lexicon to unify synonyms under a single lemma also yielded more consistent renderings of key epithets across the text. In contrast, non-RAG models sometimes varied their translations of the same term from line to line.

5 Discussion

5.1 Practical Benefits of RAG in Ancient Egyptian Studies

RAG-based methods like THOTH AI can be transformative for students and researchers in several different ways. They simplify dictionary lookups, reducing the manual labor typically required for analyzing each new inscription or papyrus fragment. By presenting morphological notes automatically, they also help novices see how suffix pronouns attach to verbs or how determinatives function, providing instant educational feedback. Moreover, large-scale text processing for Digital Humanities becomes more approachable: scholars can upload entire corpora, obtain rough translations, and refine them using specialized knowledge. Over time, these incremental enhancements further enrich the vector database, improving overall coverage

and accuracy.

5.2 Coptic Revitalization Efforts

Because the Comprehensive Coptic Lexicon includes references extending beyond Middle Egyptian into Coptic, THOTH AI can help illustrate the linguistic connections between earlier Egyptian forms and their Coptic descendants (Wannaz and Miyagawa, 2024). Its base, namely the Claude 3.5 Sonnet, is state-of-the-art in the translation ability of Coptic into English (Wannaz and Miyagawa, 2024; Saeed et al., 2024). Scholars and community members interested in promoting Coptic usage can use THOTH AI to generate bilingual texts, bridging older and contemporary liturgical contexts. This is vital for language revitalization, which typically benefits from robust digital tools and accessible learning platforms.

As it is a RAG on comprehensive Coptic–Ancient Egyptian etymological data, THOTH AI is good with answering etymological questions such as Coptic personal names such as *Papnoute* (Figure 2).

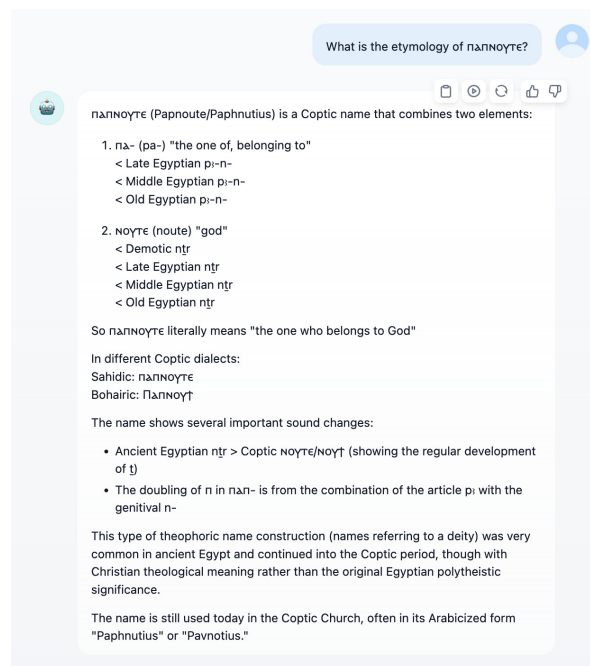


Figure 2: Using THOTH.AI for the etymology of *Papnoute* a Coptic personal name

THOTH AI is also useful for composing Coptic texts. For example, there are good cases of using THOTH AI to craft lyrics for a new Coptic song with music composition by SUNO AI, a song/music generation application. With its

Table 2: Middle Egyptian (Latin Transliteration) → English MT: Evaluation Scores on the Thutmose I Decree

Model	BLEU	SacreBLEU	ROUGE-1	ROUGE-L	chrF	METEOR
Nederhof (Ref)	1.000	100.000	1.000	1.000	100.000	1.000
THOTH AI (RAG)	0.354	35.431	0.730	0.680	61.052	0.650
Claude 3.5 Sonnet	0.325	32.457	0.717	0.652	58.064	0.640
Gemini 2.0 Pro	0.288	28.772	0.714	0.681	59.937	0.581
Gemini 2.0 Flash	0.256	25.590	0.697	0.664	59.298	0.571
DeepSeek R1	0.222	22.163	0.675	0.587	53.735	0.510
GPT o1 Pro	0.216	21.598	0.625	0.534	53.167	0.509
GPT 4o	0.196	19.615	0.581	0.484	51.109	0.417
GPT o3-mini-high	0.034	3.413	0.375	0.233	35.148	0.249

intuitive user interface and no fee for general users, this application can be a good tool for learning and revitalizing the Coptic language.

5.3 Limitations

Despite the strong performance of THOTH AI, there are notable limitations. One limitation is lexicon coverage. Even with a substantial resource like the Comprehensive Coptic Lexicon, the base of Coptic Dictionary Online (Feder et al., 2018), certain specialized religious texts or local dialect variants may remain undocumented, causing RAG to miss crucial definitions. Another challenge is OCR accuracy for hieratic or cursive hieroglyphics, as these scripts are visually complex and often damaged. Although our OCR module helps with initial transliteration, human oversight is still essential. Finally, contextualized cultural knowledge cannot be fully encoded in dictionary entries. Some references to minor deities or ephemeral socio-historical events require domain experts to interpret. Hence, while RAG grounds translations in lexical facts, it does not replace the deeper cultural or historical perspective provided by Egyptologists.

6 Conclusion

Adding a retrieval layer to mainstream Large Language Models has significantly boosted the quality of Middle Egyptian translations. Our system, THOTH AI, pairs Claude 3.5 Sonnet with a vectorized Coptic-Egyptian lexicon and morphological database. On a classical text from the reign of Thutmose I, THOTH AI’s translations consistently outperformed standard LLM approaches, as demonstrated by

higher BLEU, METEOR, ROUGE, and chrF scores. These improvements highlight how crucial domain-specific knowledge is for deciphering the archaic vocabulary and grammatical forms inherent to Ancient Egyptian.

The broader implications of this research extend to digital humanities, where scholars face extensive corpora of under-digitized ancient texts, as well as to e-learning, where students can benefit from automatic dictionary lookups and morphological guidance. In addition, the approach offers insights into Coptic revitalization, illustrating how bridging different historical stages of the same language can inform modern usage. Overall, our study suggests that the synergy between RAG and domain-specific lexical databases can substantially advance the machine translation of ancient, low-resource languages, transforming them from challenging realms of manual scholarship to more accessible domains of AI-assisted study.

Acknowledgments

This work was supported by the open-source Dify platform, which enabled straightforward implementation of vector retrieval and LLM integration. We thank the KELLIA project, Coptic SCRIPTORIUM, and the Thesaurus Linguae Aegyptiae for providing and curating the Comprehensive Coptic Lexicon, enabling a robust basis for RAG. Thanks also to colleagues and early testers who offered feedback on intermediate prototypes of THOTH AI, helping refine its user interface and retrieval settings.

Translit.	THOTH AI	Claude 3.5 Sonnet	Gemini 2.0 Pro	GPT o1 pro mode
nsw-bjtj 3- hpr-kj-r ^c s ⁱ -r ^c dhwj-msjw dj n ^h dt	The King of Upper and Lower Egypt, Aakheperkare, Son of Re, Thutmose, given life forever.	The King of Upper and Lower Egypt, Aaheperkare, Son of Ra, Thutmose, given life forever	The King of Upper and Lower Egypt, Aakheperkare, the Son of Ra, Thutmose, given life forever.	The King of Upper and Lower Egypt, Aa-kheper-ka-Ra, the son of Re, Thutmose, given life forever.
dw ⁱ -ntr sp 4 j ^c h-msjw n ^h .tj	Praise god 4 times, Iahmose, may he live!	Give praise four times to Iahmose, may he live	Adoring the god, four times, Iahmes, may you live!	Praising the god four times, (by) Iahmes, (who is) alive.
nfrt-jtj n ^h .tj wdb.tj snb.tj	Neferteri, may she live, prosper and be healthy!	Neferiri, may she live, prosper, and be healthy	Nefertiiry, may you live, prosper, and be healthy!	The one who does good, (who is) living, whole, and healthy.
wd-nsw n s ⁱ - nsw jmj-r ⁱ h ⁱ swt rswt twrtj	A royal decree to the king's son, overseer of the southern foreign lands, Tori:	Royal decree to the King's son, Overseer of the Southern Foreign Lands, Turi	Royal decree to the king's son, overseer of the southern foreign lands, Turi.	Royal decree to the King's Son, the overseer of the southern foreign lands, Ture.
mk jn.tw n=k wd pn n nsw r rdjt rh=k ntt hm=j ^c .w.s. j ^c (w) m nsw- bjtj hr st-hr nt n ^h w	Behold, this royal decree is brought to you to let you know that My Majesty (l.p.h.) has appeared as King of Upper and Lower Egypt upon the throne of Horus of the living ones.	Behold, this royal decree is brought to you to let you know that My Majesty (l.p.h.) has appeared as King of Upper and Lower Egypt upon the Throne of Horus of the Living	Behold, this decree of the king is brought to you to let you know that my majesty (l.p.h.) has appeared as King of Upper and Lower Egypt upon the Horus-throne of the living.	Behold, one brings to you this decree of the King to let you know that My Majesty has appeared as King of Upper and Lower Egypt upon the Living Horus-Throne.
nn whm.tj=j dt jr nhbt=j m	It shall never be repeated again. As for my titulary:	Which shall not be repeated forever. As for my titulary:	without its being repeated forever, my being established in	It shall never be undone forever, for my Nekhbet is with me.
hr k ⁱ -nht mrtj- m ^s t nbtj h ^c j- m-nst ⁱ 3-ph ⁱ tj hr-nbw nfr- rnpwt s ^c n ^h - j ^b w nsw-bjtj 3- hpr-kj-r ^c s ⁱ -r ^c dhwj-msjw	Horus: Mighty Bull, Beloved of Maat. Two Ladies: Appearing in the Divine Flame, Great of Strength. Golden Horus: Perfect of Years, Who Makes Hearts Live. King of Upper and Lower Egypt: Aakheperkare. Son of Re: Thutmose.	Horus: Mighty Bull, Beloved of Maat; Two Ladies: Appearing in the Uraeus, Great of Strength; Golden Horus: Perfect of Years, Who Vivifies Hearts; King of Upper and Lower Egypt: Aaheperkare; Son of Ra: Thutmose	Horus, the Strong Bull, Beloved of Maat; the Two Ladies, Who Appears in the Uraeus, Great of Strength; the Golden Horus, Perfect of Years, Who Makes Hearts Live; the King of Upper and Lower Egypt, Aakheperkare, the Son of Ra, Thutmose,	The Horus "Strong Bull, Beloved of Ma' at," the Two Ladies "He who appears in fire, Great of Power," the Golden Horus "Perfect of Years, who gives life to hearts," (he is) the King of Upper and Lower Egypt, Aa-kheper-ka-Ra, the son of Re, Thutmose.
n ^h (w) dt r nhh	Living forever and ever.	Living forever and ever	may he live forever and eternally.	Alive for all eternity, forever.
jl d ^d =k dj.tw m ^s c htp-ntr n ntrw tpj sm ^c w j ^b w m jrt h ^s swt hr-tp c.w.s. nsw-bjtj	Now you shall cause that divine offerings be given correctly to the gods of the first (nome) of Upper Egypt at Abu (Elephantine), making praises on behalf of (l.p.h.) the King of Upper and Lower Egypt,	Then you shall cause the divine offerings to be presented to the gods who are upon Upper Egypt and Elephantine in making praises on behalf of (l.p.h.) the King of Upper and Lower Egypt	Now you shall cause to make true the divine offerings of the gods who are in Upper Egypt (and) Elephantine, through the performance of the favors of the foremost (l.p.h.), the King of Upper and Lower Egypt,	Then you shall say: "A true offering and a divine offering is given to the gods, the Primeval Ones of Upper Egypt and Elephantine, (taken) from the produce of foreign lands, upon His Majesty's throne, the King of Upper and Lower Egypt."
3-hpr-kj-r ^c dj n ^h jl d ^d =k dj.tw mn n ^h m rn n hm=j ^c c.w.s. ms.n mwt-nsw sn=j- snb	Aakheperkare, given life. Now you shall cause that life remains established in the name of My Majesty (l.p.h.), born of the King's Mother Senseneb.	Aaheperkare, given life. Then you shall cause life to remain established in the name of My Majesty (l.p.h.), born of the King's Mother Seniseneb	Aakheperkare, given life. Now you shall cause 'life' to remain in the name of my majesty (l.p.h), born of the King's Mother, Seniseneb,	Aa-kheper-ka-Ra, given life. Then you shall say: "Lasting life is given in the name of My Majesty, (who was) born of the King's Mother Sen(j)-seneb."
snb.tj h ⁱ b pw r rdjt rh=k st ntt pr-nsw d(w) wdb(w)	May you be healthy! This is a message to let you know that the Royal House is safe and prosperous.	May you be healthy. This is a dispatch to let you know that the Royal House is safe and prosperous	may she be healthy. It is a communication to let you know about them, because the palace is whole and prosperous.	May you live. This festival is to let you know that the King's House is flourishing and intact.
sphr sj m rnp ^t - sp 1 3bd 3 prt sw 21 hrw n hb n h ^c w	Written in regnal year 1, third month of Peret, day 21, the day of the Festival of Appearance.	Copied in Year 1, Month 3 of Winter, Day 21, the day of the Festival of Appearance	Copied in Year 1, third month of Peret (Emergence/Winter), day 21, the day of the festival of the coronation.	It came to pass in Year 1, the third month of Peret (Winter), day 21, (on) the day of the Festival of Appearance (h'w).

Table 3: The comparison of translations by THOTH AI, Claude 3.5 Sonnet, Gemini 2.0 Pro, and GPT o1 pro mode

References

- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72.
- Dylan Michael Burns, Frank Feder, Katrin John, and Maxim Kupreyev. 2020. Comprehensive coptic lexicon: Including loanwords from ancient greek. <http://dx.doi.org/10.17169/refubium-27566> (accessed February 23, 2025).
- Adriaan de Buck. 1948. *Egyptian Readingbook*. Ares Publishers.
- Mohammed Enis and Mark Hopkins. 2024. From LLM to NMT: Advancing low-resource machine translation with claude. *arXiv:2404.13813 [cs.CL]*.
- Frank Feder, Maxim Kupreyev, Elizabeth Manning, Caroline T. Schroeder, and Amir Zeldes. 2018. A linked coptic dictionary online. In *Proceedings of the Second Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 12–21. Association for Computational Linguistics.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangwei Jia, Jinzhu Pan, Yan Bi, Yixin Dai, Jian Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-augmented generation for large language models: A survey. *arXiv:2312.10997v5 [cs.CL]*.
- Frank Kammerzell. 2000. Egyptian possessive constructions: A diachronic typological perspective. *Sprachtypologie und Universalienforschung: STUF*, 53:97–108.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Proceedings of the ACL Workshop: Text Summarization Branches Out*, pages 74–81.
- So Miyagawa. 2024a. Gengo fukkō, bunka hasshin no tame no kikai hon’yaku: Koputo-go bohaira hōgen no nichijō-teki shiyō no kasseika ni mukete [JPN: Machine translation for language revitalization and cultural dissemination: Toward revitalizing the everyday use of the bohairic dialect of coptic]. *Digital Archive Gakkai-shi*, pages 124–128.
- So Miyagawa. 2024b. Koputo-go kyōiku to gengo fukkō undō o shien suru tame no kikai hon’yaku dēta setto no kōchiku [JPN: Building a machine translation dataset to support coptic education and the language revitalization movement]. In *Jinkō Chinō Gakkai Zenkoku Taikai Ronbunshū 2024 [1N4-OS-18]*, pages 1–4.
- So Miyagawa and Nozomi Kawai. 2024. *Yomenai moji ni idonda hitobito: Hieroglyph kaidoku 1600-nenshi [JPN: People who Challenged Undeciphered Script: 1,600-Year History of Decipherment of Egyptian Hieroglyphs]*. Yamakawa Shuppansha.
- Mark Nederhof. 2023. Digital scholarly editions of middle egyptian texts. <http://nederhof.github.io/EgyptianTexts>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. pages 311–318.
- Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Muhammed Saeed, Asim Mohamed, Mukhtar Mohamed, Shady Shehata, and Muhammad Abdul-Mageed. 2024. From Nile sands to digital hands: Machine translation of Coptic texts. In *Proceedings of The Second Arabic Natural Language Processing Conference*, pages 298–308, Bangkok, Thailand. Association for Computational Linguistics.
- Caroline T. Schroeder and Amir Zeldes. 2016. Raiders of the Lost Corpus. *Digital Humanities Quarterly*, 10(2).
- Kurt Sethe. 1927. *Urkunden der 18. Dynastie, Volume I*. Hinrichs.
- Audric-Charles Wannaz and So Miyagawa. 2024. Assessing large language models in translating Coptic and Ancient Greek ostraca. In *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*, pages 463–471, Miami, USA. Association for Computational Linguistics.