

EXPLAINING INTERTEXTUALITY THROUGH AI

SO MIYAGAWA *

Abstract *This article argues that AI-assisted intertextuality should be judged by two linked capacities: the capacity to detect textual relations and the capacity to explain why those relations matter. Its originality is methodological rather than infrastructural: it proposes a constrained division of labor in which deterministic tools discover candidate textual relations, retrieval-augmented generation (RAG) explains them, and philologists adjudicate the claim. The test case is Coptic monastic literature, especially the writings of Shenoute and Besa, leaders of the White Monastery federation in late antique Egypt. Their works are difficult for non-specialists because Coptic requires segmentation, biblical quotations are often transformed, and scriptural language functions as social authority rather than ornament. Building on Miyagawa’s TRACER-based dissertation and recent THOTH.AI experiments, the article compares TRACER, passim, and RAG. TRACER and passim remain necessary for reproducible large-scale discovery of quotations and near quotations. RAG-based AI contributes differently: it can retrieve lexical and textual evidence, translate and segment Coptic, explain altered wording, and make allusive hypotheses explicit. The article proposes a conservative hybrid workflow: deterministic tools discover candidates, RAG explains them under system-level constraints, and philologists validate the final claim.*

Keywords: Coptic, intertextuality, retrieval-augmented generation, THOTH.AI, biblical quotation

1. INTRODUCTION

The problem addressed in this article is simple but methodologically important. Intertextuality research requires both detection and explanation. Detection asks whether one passage resembles another: is there an exact quotation, a near quotation, a paraphrase, an allusion, or only a coincidental resemblance? Explanation asks a further question: what does the relation do in the target text? In late antique Christian writing, this second question is often the more important one, because an author may cite Scripture to rebuke, console, authorize, threaten, educate, or construct a public persona.

Coptic monastic literature makes this distinction especially clear. Coptic is the last major phase of the Egyptian language, written mainly in the Greek alphabet with additional signs from earlier Egyptian writing. By the fourth and fifth centuries, it had become a major written language of Egyptian Christianity. Shenoute of Atripe and his successor Besa used it to govern, teach, and admonish the White Monastery federation in Upper Egypt. Their texts are saturated

* Institute of Humanities and Social Sciences; Research Center for West Asian Civilization, University of Tsukuba

with biblical language. Yet that language is rarely a decorative quotation. It is part of a social act: the abbot speaks to the monastery through the Bible.

The argument proceeds in three steps. First, biblical citation in the writings of Shenoute and Besa should be understood as an act of authority within a disciplined monastic society, not merely as literary borrowing. Second, TRACER and passim remain indispensable because they produce reproducible candidate sets for verbal reuse. Third, a RAG-based system such as THOTH.AI can add a different kind of value: it can explain transformed quotations and allusions in a way that is useful to both specialists and non-specialists. The article therefore does not claim that AI replaces philology. It argues that the best workflow is layered: algorithms discover, RAG explains, and philologists decide.

Because the system overview and specifications of THOTH.AI have been introduced elsewhere, this article does not present another technical system description. Its contribution is to define an evidence-centered workflow for AI-assisted intertextuality, especially for low-resource historical corpora. In this workflow, TRACER and passim remain the reproducible discovery layer; THOTH.AI serves as a constrained explanatory layer; and expert review remains the adjudicative layer. This division of labor is intended to be portable beyond Coptic to other linguistic, literary, and historical corpora in which verbal reuse, semantic resemblance, and social context must be evaluated together.

The article uses only cautious claims about AI. Published studies of low-resource ancient-language translation show that large language models can assist experts, but they also show that hallucinations remain a serious risk, especially for Coptic (Wannaz and Miyagawa 2024). For this reason, every claim below distinguishes between confirmed evidence, author-supplied pilot data, and hypotheses requiring validation.

2. COPTIC MONASTIC LITERATURE AND BIBLICAL AUTHORITY

2.1 Coptic for Readers Outside the Field

For readers unfamiliar with Coptic, two points are essential. First, Coptic is Egyptian, not Greek, even though its script is largely Greek and its Christian vocabulary contains many Greek loanwords (Lambdin 1983; Layton 2011). Second, Coptic does not always offer computational tools the same kind of word division that modern English does. A Coptic text often has to be normalized, tokenized, lemmatized, and syntactically analyzed before a machine can compare it with another text (Zeldes and Schroeder 2016; Zeldes and Abrams 2018). In practical terms, this means that pre-processing is not a technical afterthought. It determines what the software can recognize.

A small example illustrates the issue. The Coptic phrase *pefran* means ‘his name.’ A trained reader sees two parts: *pef-*, the possessive article with the masculine singular possessee and third-person masculine singular possessor ‘his,’ and the masculine singular noun *ran*, ‘name.’ A string-matching system that receives only unsegmented or inconsistently segmented text may fail to compare the phrase properly. A RAG-based assistant can be useful here because it can retrieve lexical evidence and explain the segmentation in ordinary language. This is a major

advantage for readers who do not know Coptic but need to evaluate the argument.

Figure 1 gives a schematic example of what a Coptic passage and a THOTH.AI-style explanation can look like to readers outside Coptology. It displays the Coptic string, romanization, translation, retrieved source, shared lexical items, changed wording, and provisional classification. The purpose is visual orientation, not proof of the proposed intertextual relation.

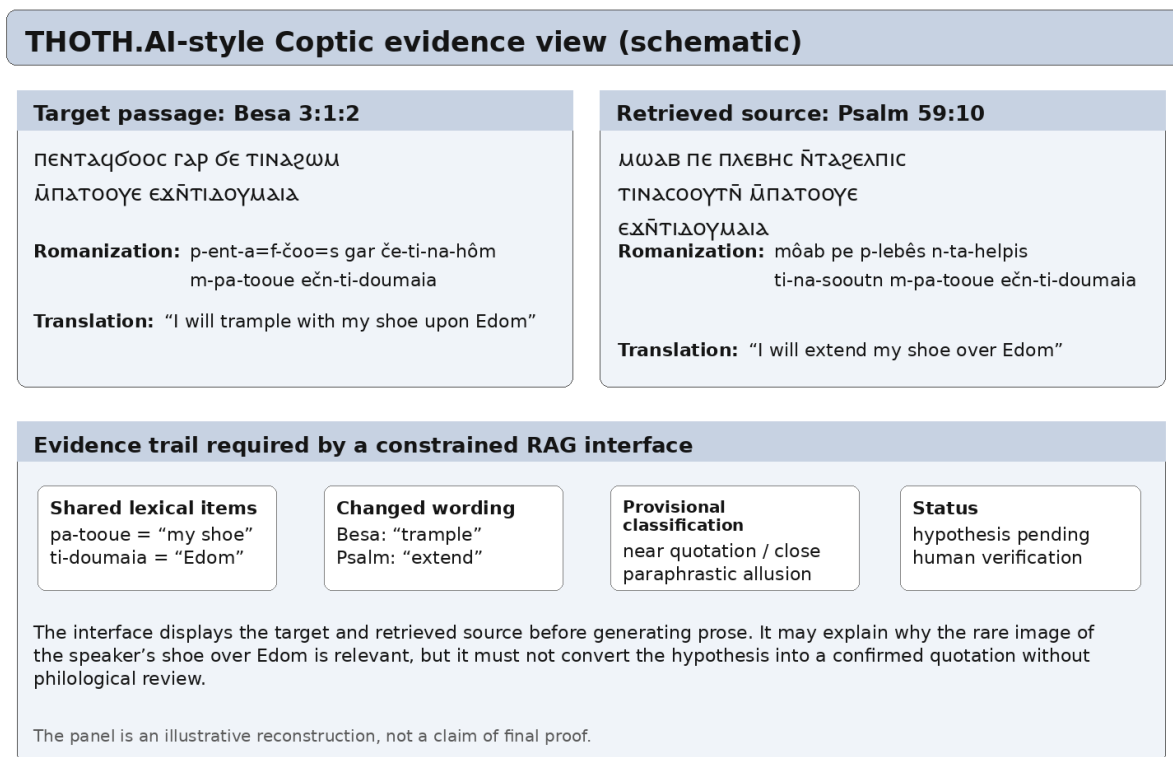


Figure 1. Schematic Coptic evidence view in a THOTH.AI-style explanation. The panel reconstructs the Besa 3:1:2 / Psalm 59:10 example and shows the kind of evidence trail required before a hypothesis can be reviewed philologically.

2.2 Shenoute, Besa, and a Scriptural Monastery

Shenoute and Besa wrote for a monastic federation in which Scripture shaped discipline, memory, and obedience. The White Monastery federation was not simply a private ascetic community. It was a regulated social body of monks and nuns, governed through rules, letters, sermons, assemblies, punishment, and biblical interpretation (Layton 2002, 2007, 2014; Krawiec 2002; Schroeder 2013; Brakke and Crislip 2015; Blanke 2019). In such a setting, a biblical quotation did not function like a modern footnote. It allowed the abbot to make a present disciplinary problem sound like part of the continuing biblical story.

Miyagawa's dissertation summarizes this process through the concept of collective memory. Shenoute and Besa blended biblical phrases and concepts with monastic ideals, often without explicit quotation signals. Their audience could hear the current crisis of the monastery through the remembered language of Israel, the Psalms, the prophets, the disciples, and the apostolic church (Miyagawa 2022, 291-92). This is why intertextuality matters historically. Biblical

language helped define what the community was, who had authority, and what obedience meant.

2.3 Why Citation Conferred Authority

The social force of quotation can be stated as a three-part chain. Scripture was the highest shared authority of the community. Monastic formation trained monks and nuns to hear Scripture as a rule for conduct. Therefore, an abbot who could make Scripture speak to the present situation could present himself not only as a teacher, but also as interpreter, judge, father, sufferer, prophet, and guardian. The words quoted were not merely evidence; they shaped the speaker's role.

This is especially important for Shenoute. Scholars have described his self-presentation as prophetic, Pauline, and that of a suffering servant (Krawiec 2002; Schroeder 2013; López 2013). The point is not that Shenoute merely imitated biblical figures. Rather, his rhetoric made the monastery's conflicts legible through biblical roles. The disobedient could become the wicked; the community could become Israel; the abbot could speak in the tones of a prophet, a psalmist, or an apostolic father. Besa inherited this scriptural rhetoric, even when his style differed from Shenoute's.

The Psalms were especially central. They supplied language for lament, accusation, repentance, persecution, judgment, and trust in God. Miyagawa's corpus study shows that the Psalms occupy a major position in both Shenoute and Besa: they are the most frequently quoted book in the examined parts of Shenoute's Canon 6 and the second most frequently quoted book in Besa's Letters and Sermons (Miyagawa 2022, 215). This explains why the article's examples focus on Psalmic reuse.

This also explains why monastic culture must remain part of the argument. The more a community learned Scripture through recitation, discipline, and repeated hearing, the less an abbot needed to mark every quotation explicitly. Studies of Shenoute's quotation practice and of monastic education show that biblical language could function simultaneously as teaching, memory work, and institutional discipline (Behlmer 2017; Timbie 2013, 2016). For readers outside Coptology, the question is therefore not only whether Shenoute or Besa quoted the Bible, but how quotation made their speech recognizable as authoritative within a biblically trained community

3. EVIDENCE, DETECTION, AND INTERPRETATION

3.1 A Conservative Evidence Ladder

AI-assisted intertextuality becomes unreliable when every suggested connection is treated as equally strong. A rigorous method needs an evidence ladder. Exact quotation has the strongest evidence; structural allusion has weaker but still potentially meaningful evidence; broad thematic resemblance is not enough. Table 1 gives the operational standard used in this article.

Recent computational work on intertextuality usually distinguishes two kinds of signal: verbal reuse, captured by sequence alignment or passage matching, and semantic similarity, captured by distributional resources such as embeddings or wordnets. These signals are

complementary. Verbal methods are reproducible and easy to audit, but they may miss transformed allusions. Semantic methods can broaden candidate discovery, but they may also overgenerate thematic parallels. The workflow proposed here therefore treats both forms of computation as candidate discovery rather than final proof.

Table 1 Evidence Ladder for Coptic Intertextuality

Level	Type	Minimum Evidence	Main Risk
1	Exact quotation	Identical or virtually identical sequence in target and source	Low, if source text is secure
2	Near quotation	Distinctive shared wording with explainable grammatical change	Moderate; manuscript variation must be considered
3	Paraphrase	Same proposition, image, or argument with altered wording	Moderate to high; needs contextual support
4	Allusion	Distinctive motif, structure, or cluster without sustained verbal overlap	High; generic biblical language must be excluded
5	Coincidence	Only common vocabulary or broad religious theme	Should not be counted as intertextual evidence

Note: The ladder is heuristic. It is designed to prevent AI output from being mistaken for evidence. The classification must be based on visible Coptic, a candidate source, and a reasoned explanation.

3.2 What TRACER and Passim Contribute

TRACER is a historical text-reuse tool and software framework developed in Büchler’s doctoral dissertation and later described in the *Journal of Data Mining & Digital Humanities* as a multilevel framework for historical text-reuse detection. It processes textual corpora through configurable stages of pre-processing, feature selection, linking, scoring, and post-processing (Büchler 2013; Büchler et al. 2014, 2018). Conceptually, it asks whether two passages share enough ordered lexical material to be treated as a candidate case of reuse. The user can adjust how the corpus is normalized, what counts as a comparable unit, which distance or similarity measures are used, and which thresholds are retained. In Miyagawa’s dissertation, TRACER was applied to prepared Coptic corpora: the Psalms, Shenoute’s Canon 6, and Besa’s Letters and Sermons. The work required Unicode conversion, normalization, tokenization, lemmatization, and careful manual analysis of the resulting candidates (Miyagawa 2022, 161-62, 216-18). This is crucial: the computational result depended on philological preparation.

For Coptic materials, this preparation is especially important because a formally small decision can change the result. A possessive construction may be written as one orthographic string but analyzed as several grammatical parts; Greek loanwords and Egyptian lexical items may require different lexical treatment; and orthographic variation can obscure an otherwise clear quotation. TRACER was therefore used not as an automatic annotator but as a reproducible screening instrument: it compared the normalized Psalm corpus with normalized monastic target texts and produced alignments that could then be checked against editions, translations, and context.

The results were significant. TRACER identified 23 Psalm quotations in Shenoute’s Canon

6 and 46 in Besa’s Letters and Sermons. Fourteen of the Shenoute quotations and five of the Besa quotations were newly identified (Miyagawa 2022, 289). Figure 2 summarizes these TRACER results visually, distinguishing the total detected Psalm quotations from the newly identified quotations in the two corpora. These results show that deterministic text-reuse tools can find passages that even excellent human annotation may miss. They also show the limit of the method: TRACER produced candidates, not final interpretations. The dissertation explicitly notes that the software cannot by itself decide whether a candidate is a quotation, an idiom, an allusion, or noise (Miyagawa 2022, 216-17).




Figure 2 TRACER results for Psalm quotations in Shenoute and Besa. Source: Aggregate conclusion data from Miyagawa 2022, 289.

Passim has a related but distinct value. It is an open-source tool for finding text reuse in large digital archives by comparing tokenized passages and grouping similar passages into reuse clusters (Romanello and Hengchen 2021). In the present Coptic context, passim can be used after the same philological preparation required for TRACER: the text must be converted to Unicode, segmented consistently, and normalized so that recurring lexical material can be compared. Its strength is scalable passage-level comparison; its weakness is that transformed quotations, paraphrases, and structural allusions may still fall below the verbal threshold. Thus, TRACER and passim are not rivals to philology. They are ways of making the first stage of philology - candidate discovery - explicit and reproducible.

3.3 Why Detection Alone Is Not Enough

The hardest cases are rarely exact quotations. In Coptic monastic writing, an author may change person, number, tense, auxiliary, word order, or addressee in order to move Scripture into a new social situation. These changes are not merely errors or memory slips. They may be the rhetorical point. A Psalm addressed to God may become a rebuke to the monastery. A line about Israel may be redirected to monks and nuns. A biblical curse may become a disciplinary threat. Detection can show the overlap; interpretation must explain the transformation.

For that reason, ‘explanation’ in this article does not mean revealing a model’s hidden computations. It means producing a public, checkable account of evidence: which words

correspond, what source is proposed, what changed, why the change matters, what alternative explanations exist, and what degree of confidence is justified. This is the specific sense in which THOTH.AI is useful for intertextual research.

4. RAG-BASED AI AND THOTH.AI

4.1 RAG as a Philological Architecture

Retrieval-augmented generation combines a generative model with an external retrieval system. In the standard formulation, RAG joins a model's parametric memory with non-parametric memory such as a dense vector index, allowing generated output to be conditioned on retrieved passages rather than on model weights alone (Lewis et al. 2020). For historical philology, the attraction is not that the model becomes omniscient. The attraction is that its answer can be grounded in retrieved lexica, parallel passages, and grammatical evidence. For Coptic, a resource such as the Comprehensive Coptic Lexicon supplies the kind of stable lexical data that can be retrieved instead of improvised by the model (Burns et al. 2020).

THOTH.AI applies this architecture to Ancient Egyptian and Coptic. The 2025 THOTH AI study reports that RAG improves Middle Egyptian-to-English translation by injecting specialized lexica and morphological data into the prompt (Miyagawa 2025). The same principle is relevant to Coptic intertextuality: the model should not be trusted because it sounds fluent. It should be trusted only when it retrieves evidence, displays the relevant text, and explains its reasoning in a way that a human scholar can test.

The unpublished THOTH.AI pilot material also reports project-internal Bohairic-to-English results for the *Vita Sinuthii*: a base-model BLEU of 31.40 and a THOTH.AI 2.0 with RAG score of 40.63, a gain of 9.23 BLEU (Miyagawa 2026). Figure 3 visualizes these lecture-slide data. It is included only as evidence for the methodological point that retrieval can improve domain-specific output. It is not presented as a peer-reviewed intertextuality benchmark.

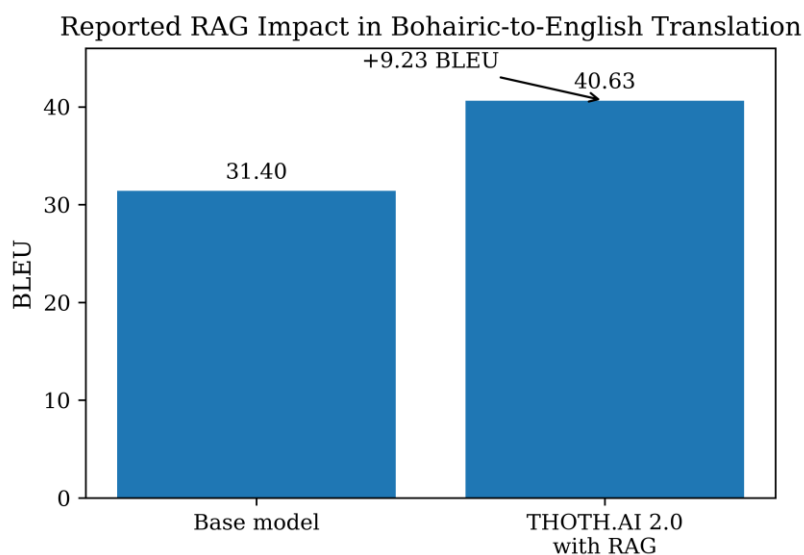


Figure 3 Reported RAG impact in a Bohairic-to-English translation evaluation. Source:

Author-supplied THOTH.AI lecture-slide data; this is not a corpus-wide intertextuality benchmark.

4.2 The Specific Merit of THOTH.AI: Explanation

For intertextual research, the most valuable use of THOTH.AI is not simply to produce more candidates. More candidates can increase confusion if they are not ranked and explained. Its distinctive value is explanatory: it can translate a passage, segment it, identify distinctive words, retrieve candidate biblical sources, and formulate a hypothesis such as ‘near quotation,’ ‘paraphrase,’ or ‘structural allusion.’ This is especially useful when the evidence is not a long exact string but a distinctive image, motif, or rhetorical pattern.

The safe use of such a system requires strict constraints. It should not be asked to invent bibliography, fill gaps from memory, or turn uncertain parallels into certainties. It should be asked to reason from supplied Coptic passages and retrieved sources. Any AI-generated claim must be reducible to the target passage, the proposed source, the alignment or motif relation, the translation, and the stated reason for classification. Without that evidence trail, the claim should not enter the argument. This caution remains necessary even when semantic resources are added: experiments with Coptic Wordnet improve candidate generation, but they do not remove the need for philological classification and verification (Miyagawa et al. 2025).

These constraints should be implemented in the interface itself rather than left only to the user’s discretion. A responsible Coptic intertextuality mode should behave less like an unrestricted chat window and more like a guided annotation environment. The user should select a target passage and permitted source corpora; the system should retrieve only from indexed editions, lexica, and bibliographic records; and the output should be blocked or labeled incomplete when the required evidence fields are missing. Table 2 summarizes the corresponding system-level constraints that would make THOTH.AI’s explanatory output reviewable rather than merely fluent.

Table 2 System-Level Constraints for a Coptic Intertextuality Interface

Control	Concrete implementation	Purpose
Source-gated retrieval	User selects the target passage and allowed corpora or lexica; the model may cite only retrieved records.	Prevents invented sources, editions, and bibliography.
Mandatory evidence fields	The interface requires target text, candidate source, translation, alignment or motif relation, classification, rationale, and uncertainty level.	Makes every claim inspectable by a human reader.
No-evidence behavior	If retrieval fails, the system returns “no verifiable candidate” rather than a fluent explanation.	Prevents gap-filling from model memory.
Evidence-ladder labels	The user must choose exact quotation, near quotation, paraphrase, allusion, or coincidence before exporting the result.	Keeps hypotheses separate from confirmed evidence.

Control	Concrete implementation	Purpose
Provenance logging	The system records retrieved passages, prompt template, model version, date, and later human decision.	Supports reproducibility, teaching review, and audit.
Pedagogical mode	For students, the interface asks them to justify the classification before showing the model’s proposed label.	Turns model instability into guided re-evaluation.

Note: These controls describe system design requirements for the explanatory layer. They do not make AI output authoritative; they make the output reviewable.

5. COPTIC TEST CASES

The Coptic examples below use Kuhn’s edition of Besa’s Letters and Sermons as the textual base, and English translations are adapted from Kuhn’s translation where indicated (Kuhn 1956a, 1956b).

5.1 Besa 4/I/2: A Chain of Psalmic Judgment

The first example shows why deterministic discovery is still necessary. In Besa 4/I/2, a disciplinary passage brings together several biblical images: pursuit with whirlwind and wrath, fire devouring the guilty, an angel removing offenders, dust scattered by wind, and a curse that comes upon those who love it. The unpublished THOTH.AI pilot material highlights Psalm 82:16, Psalm 20:10, Psalm 1:4, Exodus 23:20-23, and Psalm 108:17 in this cluster. TRACER helped reveal Psalmic links that were not always marked in earlier annotation (Miyagawa 2022, 244; Miyagawa 2026).

One short alignment is enough to show the method. Besa’s passage includes the phrase translated ‘and a fire shall devour them.’ The Coptic target phrase is:

Romanized: nte=ou-sate ouom=ou

Translation: and a fire shall devour them

The corresponding part of Psalm 20:10 includes:

Romanized: nte=ou-sate ouom=ou

Translation: and a fire will devour them

Here the evidence is strong because the wording is short but distinctive within a larger cluster of judgment imagery. Yet the interpretation is not finished when the match is found. Besa is not merely listing Psalm quotations. He is composing a disciplinary threat in biblical language. The quotation helps make monastic punishment sound like divine judgment. That is a social and rhetorical claim, not only a string-matching result.

5.2 Besa 3:1:2 and Psalm 59:10: A Transformed Near Quotation

The second example shows why RAG-based explanation can help, but it must be stated cautiously. In Besa 3:1:2, the unpublished THOTH.AI pilot material reports a possible relation that was not reported by Kuhn, TRACER, or passim in that local pilot comparison. The passage

contains the phrase translated by Kuhn as ‘I will trample with my shoe upon Edom.’ The proposed biblical source is Psalm 59:10 in the Septuagint numbering, where the Coptic Psalm says, ‘I will extend my shoe over Edom.’ The shared image of the speaker’s shoe over Edom is distinctive, but the claim remains a hypothesis until the target passage, the Psalter text, and any relevant manuscript or edition variants are checked philologically.

Target passage: Besa 3:1:2

Romanized: p-ent-a=f-čoo=s gar če-ti-na-hôm m-pa-tooue ečn-ti-doumaia

Translation: For he who has said, ‘I will trample with my shoe upon Edom’ (adapted from Kuhn 1956b)

Candidate source: Psalm 59:10

Romanized: môab pe p-lebês n-ta-helpis ti-na-sooutn m-pa-tooue ečn-ti-doumaia

Translation: Moab is the basin of my hope; I will extend my shoe over Edom

The situational context makes the similarity easier to understand. Psalm 59:10 belongs to a victory and dominion idiom in which surrounding peoples such as Moab and Edom are named as subordinated outsiders. The shoe image gives physical form to domination: Edom is placed under or beneath the speaker’s shoe. Besa does not need to retell the Psalmic scene. In a disciplinary or polemical setting, the image can be reactivated as a compact scriptural sign of subjugation and judgment. This context explains why the parallel matters even though the verb has changed.

For a non-Coptic reader, the central evidence is easy to state. Both passages contain ‘my shoe’ (pa-tooue) and ‘Edom’ (ti-doumaia). The verb differs: Besa’s passage uses a verb translated ‘trample,’ while the Psalm uses a verb translated ‘extend.’ A rigid exact-string approach may underweight the parallel because the repeated material is not a long continuous phrase. A RAG-based explanation can articulate the reason for the hypothesis: a rare biblical image of domination over Edom may have been reactivated with a transformed verb. The cautious working classification is therefore near quotation or close paraphrastic allusion, pending full comparison of the Coptic source text and manuscript evidence.

5.3 Psalm 1:1 and Negative Inversion

The third example is weaker as verbal evidence but important as explanatory evidence. Psalm 1:1 begins by blessing the one who does not enter progressively into evil: he does not walk in the counsel of the wicked, stand in the way of sinners, or sit in the seat of pestilent persons. The rhetorical movement is from motion to permanence: walking, standing, sitting. The blessed person refuses to settle in evil.

Source: Psalm 1:1

Romanized: naiat=f m-p-rôme ete-mpe=f-bôk hm-p-šočne n-n-asebês mpe=f-ahe-rat=f hi-te-hiê n-n-ref-r-nobe mpe=f-hmoos hi-t-kathêdra n-n-loimos

Translation: Blessed is the man who did not go in the counsel of the wicked, did not stand in the way of sinners, and did not sit on the seat of pestilent persons

Besa’s *To an Erring Monk* describes the opposite moral condition. Those who become dark are fed upon by foxes; those who transgress the commandments become the place where the

devil and his demons feed. The relevant Coptic phrase is:

Romanized: net-ne-r-kake n-bašor na-moone hi-čô=ou ... net-na-paraba n-n-entolê p-diabolos
auô nef-daimôn na-moone n-hêt=ou

Translation: On those who become dark, the foxes shall feed ... in those who transgress the
commandments will the devil and his demons feed (adapted from Kuhn 1956b, 22)

The contrast is therefore not a logical contradiction between two propositions but an inversion of moral space. Psalm 1:1 praises the person who refuses three increasingly settled positions in evil: going, standing, and sitting. Besa describes people who have crossed that boundary. By transgressing the commandments, they become the place in which foxes, the devil, and demons act. The shared logic is spatial and moral: Psalm 1 imagines blessedness as refusing to inhabit the sphere of the wicked, while Besa imagines disobedience as becoming inhabited by evil. In this pilot test, THOTH.AI’s useful contribution is to name that structure rather than merely count shared words. This should not be presented as a secure quotation. It is better called a possible structural allusion or ‘negative inversion.’ The label is valuable because it makes the claim testable: the scholar can now ask whether the pattern is distinctive enough, whether Besa uses the Psalm elsewhere, and whether alternative biblical sources explain the passage better.

A related pilot example concerns Besa 3:2:3 and Psalm 59:6. The unpublished THOTH.AI pilot material reports that Kuhn marked the relation with ‘cf.’ that TRACER and passim did not identify it in the pilot run, and that THOTH.AI explained the relation through the motif cluster of a sign and flight from the bow (Miyagawa 2026). This is exactly the kind of claim that should be treated as a hypothesis: promising, explainable, and still dependent on human verification

5.4 Pilot Comparison and Its Limits

The pilot report summarizes a selected Besa comparison in which Kuhn identified 9 relevant cases, TRACER 11, passim 12, and THOTH.AI 16. Figure 4 visualizes the data. The figure must be read cautiously. It is not corpus-wide precision, recall, or F1. It does not prove that THOTH.AI is generally more accurate than TRACER or passim. It shows only that, in one selected test case, a RAG-based system produced additional explainable candidates, especially where transformed wording and allusion mattered.

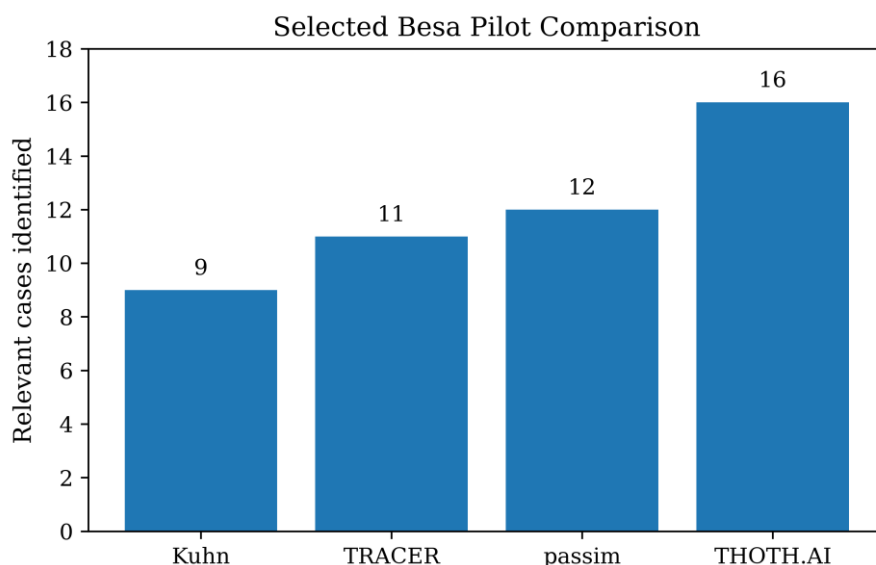


Figure 4 Pilot comparison in a selected Besa test case. Source: Author-supplied THOTH.AI pilot data; the figure reports local case-study counts, not corpus-wide precision or recall.

6. A HYBRID WORKFLOW FOR RESPONSIBLE AI INTERTEXTUALITY

The argument can now be converted into a workflow. First, build a clean corpus and record every pre-processing decision: manuscript base, edition, normalization, tokenization, and lemmatization. Second, run TRACER or passim to obtain reproducible candidates. Third, inspect alignments and remove obvious noise, especially idiomatic phrases and stop-word matches. Fourth, use THOTH.AI or another RAG-based system to explain difficult candidates. Fifth, verify each proposed relation against the target passage, candidate source, manuscripts, editions, translations, lexica, and prior scholarship. Sixth, publish both the interpretation and the evidence trail. Table 3 summarizes this division of labor by setting out the roles of TRACER/passim, THOTH.AI with RAG, and the corresponding validation controls at each stage of the workflow.

Table 3 Tool Roles and Validation Controls

Stage	TRACER / passim	THOTH.AI with RAG	Validation Control
Corpus preparation	Requires clean input	Uses retrieved data only if indexed	Record normalization and segmentation
Candidate discovery	Strong for exact and near quotation	Useful for focused prompts, not bulk scanning	Keep deterministic candidate lists
Difficult cases	May miss paraphrase and structural allusion	Can explain motif and contextual relations	Require visible source and target

Stage	TRACER / <i>passim</i>	THOTH.AI with RAG	Validation Control
Risk	False positives and idioms	Hallucinated sources or fluent overclaiming	Reject unverifiable claims
Best use	Broad, reproducible screening	Deep explanation and hypothesis formation	Final human adjudication

Note: The table presents a workflow, not a ranking. No single tool is sufficient for responsible intertextual interpretation.

The same controls make educational use safer. For beginning students of Coptic, THOTH.AI should not present itself as an answer machine. It should display segmentation, lexical retrieval, translations, and alternative labels while requiring the student to justify a classification with the evidence ladder before accepting the model’s proposed label. In this form, the instability of generative AI becomes a prompt for careful re-evaluation rather than a source of unreviewed authority.

This workflow also defines what counts as hallucination in philology. It is not only a false answer. It may be a fabricated source, an invented quotation, a plausible hagiographic detail that is absent from the Coptic, or a translation that imports familiar Christian language not found in the target passage. The THOTH.AI pilot material itself gives examples of weaker systems generating unsupported content. RAG reduces this risk only when retrieval is authoritative and when the user requires evidence, citation, segmentation, and verification.

The central rule is therefore strict: an AI-generated intertextual claim is not evidence until its target passage, candidate source, and reasoning are visible. A model may assist discovery and explanation. It does not adjudicate the claim. This rule protects philology while allowing AI to make the evidence more accessible.

7. CONCLUSION

AI-assisted intertextuality is most persuasive when it is logically modest and philologically ambitious. It should not promise automatic interpretation. It should make interpretation clearer. In the Coptic monastic material examined here, TRACER and *passim* remain indispensable for discovering verbal reuse across prepared corpora. THOTH.AI contributes something different: it can explain why a proposed relation may matter, especially when the evidence is transformed, thematic, grammatical, or structural rather than exactly verbal. That contribution becomes responsible only when the interface enforces source-gated retrieval, mandatory evidence fields, uncertainty labels, and final human review.

For Shenoute and Besa, this explanatory level is essential. Biblical citation was one of the ways an abbot spoke with authority. Scripture authorized rebuke, shaped communal memory, and allowed monastic conflicts to be narrated through the language of the Psalms, prophets, apostles, and Israel. A successful AI system for this material must therefore do more than count overlap. It must help readers see how a Coptic phrase becomes an act of scriptural performance. The best formula remains simple: algorithms discover, RAG explains, and philologists decide.

APPENDIX A. PROMPT TEMPLATE FOR RAG-ASSISTED EXPLANATION

The following template gives the prompt pattern used for the RAG-assisted explanatory step described in Section 4.2. It is included to make the workflow reproducible. The template is deliberately restrictive: the model is asked to explain a candidate relation, not to decide that the relation is true.

Prompt template:

You are assisting a Coptic philologist. Explain, but do not prove, a possible intertextual relation between the supplied Coptic target passage and the retrieved sources. Use only the supplied target passage, retrieved Coptic sources, lexicon entries, and bibliographic records. Do not invent sources, editions, manuscript data, bibliography, translations, or historical context. If the evidence is insufficient, say so explicitly.

Required output fields:

1. Target passage: quote the supplied Coptic, romanization, and translation.
2. Candidate source: quote the retrieved Coptic source, reference, and translation.
3. Alignment or motif relation: list shared words, changed words, and grammatical or rhetorical transformations.
4. Evidence level: choose one label from the evidence ladder: exact quotation, near quotation, paraphrase, allusion, or coincidence.
5. Rationale: explain why this label is justified and what evidence would weaken it.
6. Alternatives: list other possible sources or reasons for coincidence, if any are visible in the retrieved material.
7. Human checks required: editions, manuscripts, lexica, translations, and prior scholarship to be verified.
8. Final status: mark the claim as confirmed, rejected, or hypothesis pending philological review. Use “hypothesis pending review” unless a human expert has verified the relation.

A second verification prompt can then ask the system to list only what must be checked by a human reviewer, without generating any new bibliography or new candidate source. This separates explanation from adjudication and keeps the final scholarly decision outside the generative model.

ORCID

So Miyagawa  <https://orcid.org/0000-0002-2950-7193>

REFERENCES

- Behlmer, Heike. 2017. 'The Use of the Psalms in Shenoute's Tractate He Who Sits Upon His Throne.' In *From Gnostics to Monastics: Studies in Coptic and Early Christianity in Honor of Bentley Layton*, edited by David Brakke, Stephen J. Davis, and Stephen Emmel, 315-30. *Orientalia Lovaniensia Analecta* 263. Leuven: Peeters.
- Blanke, Louise. 2019. *An Archaeology of Egyptian Monasticism: Settlement, Economy and Daily Life at the White Monastery Federation*. Yale Egyptological Publications 2. New Haven, CT: Yale Egyptology.
- Brakke, David, and Andrew Crislip. 2015. *Selected Discourses of Shenoute the Great: Community, Theology, and Social Conflict in Late Antique Egypt*. Cambridge: Cambridge University Press.
- Büchler, Marco. 2013. *Informationstechnische Aspekte des Historical Text Re-use*. PhD diss., Universität Leipzig.
- Büchler, Marco, Philip R. Burns, Martin Müller, Emily Franzini, and Greta Franzini. 2014. 'Towards a Historical Text Re-use Detection.' In *Text Mining, Theory and Applications of Natural Language Processing*, edited by Chris Biemann and Alexander Mehler, 221-38. Cham: Springer. https://doi.org/10.1007/978-3-319-12655-5_11.
- Büchler, Marco, Greta Franzini, Emily Franzini, Maria Moritz, and Kirill Bulert. 2018. 'TRACER: A Multilevel Framework for Historical Text Reuse Detection.' *Journal of Data Mining & Digital Humanities*, special issue 'Computer-Aided Processing of Intertextuality in Ancient Languages.'
- Burns, Dylan Michael, Frank Feder, Katrin John, Maxim Kupreyev, et al. 2020. *Comprehensive Coptic Lexicon: Including Loanwords from Ancient Greek v1.2*. Version 1.2. Berlin: Freie Universität Berlin. <https://doi.org/10.17169/refubium-27566>.
- Krawiec, Rebecca. 2002. *Shenoute and the Women of the White Monastery: Egyptian Monasticism in Late Antiquity*. Oxford: Oxford University Press.
- Kuhn, Karl-Heinz, ed. 1956a. *Letters and Sermons of Besa*. *Corpus Scriptorum Christianorum Orientalium* 157, *Scriptores Coptici* 21. Louvain: Imprimerie Orientaliste L. Durbecq.
- Kuhn, Karl-Heinz, trans. 1956b. *Letters and Sermons of Besa*. *Corpus Scriptorum Christianorum Orientalium* 158, *Scriptores Coptici* 22. Louvain: Imprimerie Orientaliste L. Durbecq.
- Lambdin, Thomas O. 1983. *Introduction to Sahidic Coptic*. Macon, GA: Mercer University Press.
- Layton, Bentley. 2002. 'Social Structure and Food Consumption in an Early Christian Monastery: The Evidence of Shenoute's Canons and the White Monastery Federation A.D. 385-465.' *Le Muséon* 115 (1-2): 25-55. <https://doi.org/10.2143/MUS.115.1.321>.
- Layton, Bentley. 2007. 'Rules, Patterns, and the Exercise of Power in Shenoute's Monastery: The Problem of World Replacement and Identity Maintenance.' *Journal of Early Christian Studies* 15 (1): 45-73. <https://doi.org/10.1353/earl.2007.0015>.
- Layton, Bentley. 2011. *A Coptic Grammar: With Chrestomathy and Glossary: Sahidic Dialect*. 3rd ed. *Porta Linguarum Orientalium Neue Serie* 20. Wiesbaden: Harrassowitz Verlag.
- Layton, Bentley. 2014. *The Canons of Our Fathers: Monastic Rules of Shenoute*. Oxford: Oxford University Press.
- Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. 'Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.' In *Advances in Neural Information Processing Systems* 33, 9459-74.

- López, Ariel G. 2013. *Shenoute of Atripe and the Uses of Poverty: Rural Patronage, Religious Conflict, and Monasticism in Late Antique Egypt*. Berkeley: University of California Press.
- Miyagawa, So. 2022. *Shenoute, Besa and the Bible: Digital Text Reuse Analysis of Selected Monastic Writings from Egypt*. Doctoral thesis, Georg-August-Universität Göttingen. <https://doi.org/10.53846/goediss-9082>.
- Miyagawa, So. 2025. 'RAG-Enhanced Neural Machine Translation of Ancient Egyptian Text: A Case Study of THOTH AI.' In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, 33-40. Albuquerque, USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.nlp4dh-1.4>.
- Miyagawa, So. 2026. 'THOTH.AI: A Large Language Model for Ancient Egyptian and Coptic.' Lecture Series 'Digital Coptic Textual Studies: A Decade of Digital Research on the Written Heritage of Christian Egypt in Göttingen and Beyond,' Georg-August-Universität Göttingen, 14 January 2026. Unpublished presentation slides.
- Miyagawa, So, Luis Morgado da Costa, Laura Slaughter, and Heike Behlmer. 2025. 'Automatic Detection of Coptic Text Reuse: Applying Coptic Wordnet to Intertextuality Studies in Selected Coptic Monastic Writings.' In *Proceedings of the 13th Global Wordnet Conference*, 179-84. Pavia, Italy: Global Wordnet Association. <https://doi.org/10.18653/v1/2025.gwc-1.21>.
- Romanello, Matteo, and Simon Hengchen. 2021. 'Detecting Text Reuse with Passim.' *Programming Historian* 10. <https://doi.org/10.46430/phen0092>.
- Schroeder, Caroline T. 2013. *Monastic Bodies: Discipline and Salvation in Shenoute of Atripe*. Philadelphia: University of Pennsylvania Press.
- Timbie, Janet A. 2013. 'Writing Rules and Quoting Scripture in Early Coptic Monastic Texts.' In *Ascetic Culture: Essays in Honor of Philip Rousseau*, edited by Blake Leyerle and Robin Darling Young, 29-48. Notre Dame, IN: University of Notre Dame Press.
- Timbie, Janet A. 2016. 'The Education of Shenoute and Other Cenobitic Leaders: Inside and Outside of the Monastery.' In *Education and Religion in Late Antiquity: Reflections, Social Contexts and Genres*, edited by Peter Gemeinhardt, Lieve van Hoof, and Peter Van Nuffelen, 34-46. London: Routledge.
- Wannaz, Audric-Charles, 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*, 463-71. Miami, USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.nlp4dh-1.44>.
- Zeldes, Amir, and Mitchell Abrams. 2018. 'The Coptic Universal Dependency Treebank.' In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, 192-201. Brussels, Belgium: Association for Computational Linguistics. <https://doi.org/10.18653/v1/W18-6022>.
- Zeldes, Amir, and Caroline T. Schroeder. 2016. 'An NLP Pipeline for Coptic.' In *Proceedings of the 10th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, 146-55. Berlin, Germany: Association for Computational Linguistics. <https://doi.org/10.18653/v1/W16-2119>.