

VOYNICH ABAD / CPUTA

MASTER RECOVERY & REPRODUCIBILITY REPORT

Consolidated research state through 22 August 2026

Structural findings, negative results, CPUTA-1, replication rules, recovery checklist, and next tests

STATUS: STRUCTURAL / MORPHOLOGICAL SIGNAL DETECTED; PLAINTEXT NOT RECOVERED. This report intentionally separates validated or file-documented results from later chat-derived CPUTA results that still require preservation of executable code/data for full independent reproduction.

1. Executive recovery summary

The research program has progressively moved away from free-form semantic interpretation and toward frozen, falsifiable structural tests. The durable result is that Voynichese contains constrained edge morphology, token-boundary dependencies, recurring structural states, and productive stem-preserving transformations. Historical-language mappings tested so far have not passed the full validation gates.

The most important methodological correction was the transcription audit. EVA is not assumed to be the manuscript's true alphabet. CPUTA-1 (Consensus Physical Unit Alphabet v1) instead uses cross-transcription STA-family evidence and supported composites, then checks results independently in ZL3b, GC/v101 and RF1b.

The current working architecture is: physical glyph/composite units -> token stem plus edge morphology -> functional token state -> local sequence grammar -> unknown semantic layer. Translation remains unresolved because no deterministic token-to-sound or token-to-meaning mapping has yet predicted coherent withheld plaintext.

2. Evidence-status key

Status	Meaning	Use
DOCUMENTED PASS	Result preserved in prior report/ledger with explicit test statistics.	May be treated as established within this research program, subject to outside replication.
CHAT-REPORTED REPLICATION	Result reported during the current CPUTA sequence, but executable analysis package is not preserved in this report.	Treat as high-priority provisional until code/data regenerate the numbers exactly.
MIXED / SENSITIVE	Some folds/transcriptions positive; not universal.	Do not promote to grammar/translation.
FAIL / CLOSED	Predeclared gate failed or later falsification closed the branch.	Do not recycle without a genuinely new model.
DEFERRED	Question may be legitimate, but prerequisites are not frozen.	Resume only after stated prerequisite.

3. Corpus and representation stack

Primary manuscript: Beinecke MS 408 / Voynich Manuscript.

Earlier corpus layer: EVA / IVTFF transcription, approximately 38,800 tokens and approximately 8,100 types in the independent evaluation.

Cross-transcription layer: ZL3b, GC2a/v101 and RF1b, normalized through STA1/reduced physical families.

CPUTA-1 principle: preserve supported composite units (including minim composites such as in/iin where represented synthetically), common bench/ligature families, major gallows distinctions, and explicit uncertainty around ambiguous spaces. CPUTA-1 assigns no phonetic or semantic values.

4. Chronological master ledger - foundations through v3.1

Test	Decision	What it established
v1.0	SEALED	Dictionaries, ontology, expansion grammar, weights and selection-aware null rules frozen before testing.
v1.1	FAIL	Executable adaptive mapping: no candidate language passed.
v1.2	FAIL	Targeted care/caring morphology did not uniquely beat matched Voynich suffix controls.
v1.3	FAIL	Anchor-reverse mapping did not survive frozen reverse testing.
v1.7	FAIL	Structure-gated candidate language mapping: no candidate passed.
v1.8	FAIL	Hand-specific language selection: no hand-specific candidate passed.
v1.9	FAIL	Abbreviation-tolerant English failed coverage/composite gates.
v1.10	FAIL	Latin + Middle English did not rescue held-out performance.
v1.11	FAIL	Old English + Middle English failed.
v1.12	FAIL	Pure Middle English with 229,691-form lexicon failed 40% held-out coverage/composite gates.
v1.13	MIXED/FAIL	Controlled scribal variation failed overall; Hand-2 biological verso pocket was descriptively elevated but did not validate language.
v1.14	FAIL/CLOSED	K=3 homophonic Middle English: nominal p=.0348 alone, but coverage 26.86%, uncertainty 73.14%, worst-page null 86.23%, composite below best null; branch closed.
v2.x structural	MIXED	Multiple structural/entropy/hand/section tests narrowed signal to edge/suffix behavior; several broad or second-order models failed.
v3.0	PASS	Compact shared core: l,n,o,r,s,y in 5/5 folds; penalized gain +0.03897; p=.00995.
v3.1	FAIL AS AUTONOMOUS GRAMMAR	Six-glyph+OTHER sequential grammar failed complexity-adjusted within-stream/second-order tests; core interpreted as boundary-control markers.

5. Important negative results that must remain negative

- Readable English was not recovered under the validated evaluation. The stable qoky/qokain -> care/caring pair was a property of a frozen English-optimized substitution, not proof of English plaintext.
- Numeric/harmonic branches did not isolate plaintext: lag 7 and lag 49 were not uniquely privileged; Madhava-Leibniz, ZD30 shape-matched, combined reverse, and cross-ledger tests failed their stronger controls.
- Rigid question-answer alternation failed; interior modes persisted rather than alternating.
- Fixed global block grammar failed.
- Formula-vs-ending directional lead failed (approximately p=.45 in the later reassessment).
- EDY/IIN reciprocal mirror was closed under CUPA matched controls.
- Exact prefix x suffix shell interaction worsened held-out log loss in ZL, GC and RF.
- Literal ending-level AABA cadence did not reproduce cross-transcriptionally.

6. Transcription robustness checkpoint

The transcription audit asked whether the strongest edge/boundary findings were artifacts of EVA segmentation. The preserved CPUA-1 report promoted the following as transcription-robust:

- Local ending persistence.
- Cross-token boundary dependency.
- Strong positional specialization of glyph families.
- Productive y / a+in paired-ending family after composite-unit reassessment.

The same report kept physical-ending ABAB transcription-sensitive and closed EDY/IIN reciprocal mirroring, exact shell grammar, literal ending-AABA, rigid Q/A alternation, fixed block order, and directional formula/ending lead.

7. CPUA-1 frozen representation

Rule	Frozen treatment
Source layer	STA1 / reduced-family representation; ZL3b, GC2a/v101 and RF1b analyzed independently.
Minim composites	Supported in/iin-like composites may be single physical units rather than forced EVA stroke sequences.
Bench/ligature forms	Common ch/sh/ckh/cth/cph/cfh-style synthetic forms treated as units when represented that way in STA.
Gallows	Major p/f/k/t physical classes remain distinct; no sound is assigned.
Rare variants	Collapse to family-level wildcard unless future image audit proves functional distinction.
Spaces	Certain spaces preserved; uncertain spaces remain ambiguity flags rather than silently accepted lexical boundaries.
Semantics	None. CPUA-1 is physical/structural only.

8. Preserved CPUA near-miss reassessment

Feature	Decision	Result
-y / -ain	PROMOTED STRUCTURALLY	ZL rank 2/101 (~.020); GC 1/101 (~.010); RF 5/101 (~.050) among nearest suffix-pair controls. Supports productive paired endings, not English -ing.
-edy <-> -iin	CLOSED	Matched AABA reciprocal test null in all three representations.
Prefix x suffix shell	CLOSED	Interaction worsened held-out log loss in all three transcriptions.
Physical-ending ABAB	MIXED	ZL p~.012 and RF p~.008 positive; GC null. Keep transcription-sensitive.
Physical-ending AABA	CLOSED AS LITERAL CADENCE	No cross-transcription support.

9. Current-chat CPUA functional results - preservation ledger

IMPORTANT REPRODUCIBILITY NOTE: The following results were reported during the current chat after the saved CPUA-1 PDF. They are preserved here so the hypotheses, reported numbers and pass/fail logic are not lost. They should be regenerated from frozen source files and executable code before publication or an external priority claim.

Result	Status	Preserved claim
CPUA ABAD recurrence	CHAT-REPORTED STRONG	Best role in all three transcriptions reported as small -> medium -> small -> large. Reported z approximately 5.09 (ZL), 5.09 (GC), 5.16 (RF), selection-aware p<=.001. Reported 14/15 held-out folio directions positive, one

		essentially flat.
A/B/D functional profiles	CHAT-REPORTED	A ~14%, longest (~5.41 units), line/paragraph-opening biased; B ~37%, ~4.96 units, persistent and Currier-B/Hands-2/3 enriched; D ~49%, shortest (~2.81), default/position-neutral.
Shared internal stems	CHAT-REPORTED	~59-67% of tokens reportedly belong to internal-body families crossing states; A<->B dominant shared-stem relation (reported 287 ZL, 344 GC, 343 RF families).
A->B edge rules	CHAT-REPORTED	Reported robust opening changes include y->o, ch->q, d->k, ch->sh; endings include o->y, l->r, r->l, y->s, r/l->iin. Multiple rules reportedly recur across all held-out folds.
Opening vs ending directionality	CHAT-REPORTED	y->o and d->k reported as strongly left-facing; o->y reported to alter both sides, especially next-state distribution. Interpretation: prefix-like system relates to left context; suffix-like system more strongly constrains right context.
-o / -y continuation	CHAT-REPORTED	For same stems, -o reportedly favors shorter/more repetitive next tokens and D-state continuation; -y favors longer/more diverse next tokens and B-state continuation; reported 15/15 split direction for next-token length.
Short-token grammar class	CHAT-REPORTED	Short/common tokens after -o reported as grammar/glue-like rather than content-like: repetitive, broad attachment, predictive of following state.
Particle subclasses	CHAT-REPORTED / PROVISIONAL	ol/al initially same-class connector candidates; or/s/r cross-class relational candidates. After individual split: ol strongest same-class bridge; al bridge-like but position-sensitive; or strongest cross-class relation; r right-oriented; s unresolved. Latest grouped same-vs-different-class test was promising but not clean across every withheld split.

10. Current working grammar model

The strongest nonsemantic model currently under test is:

- LEFT CONTEXT -> opening marker + internal stem + ending marker -> RIGHT CONTEXT
- Opening markers may encode a left-facing positional/relational function.
- Ending markers may encode continuation/class-selection constraints.
- Short high-frequency tokens may contain several specialized grammatical-particle classes rather than one undifferentiated glue vocabulary.
- A/B/D are neutral functional states, not yet noun/verb/adjective labels.
- ABAD, if regenerated at the reported CPUA strength, would describe a higher-order recurrence among functional states, not plaintext.

11. Reproducibility protocol - minimum publishable standard

1. Acquire and checksum the exact ZL3b, GC2a/v101 and RF1b source transcription files plus STA mapping tables.
2. Freeze CPUA-1 conversion code and emit a machine-readable mapping table from source glyph sequences to physical-unit families.

3. Preserve folio, line, paragraph, Currier, Hand and section metadata. Never shuffle across boundaries unless the null explicitly requires it.
4. Predeclare each hypothesis, statistic, folds, exclusions, null model, number of permutations, multiple-testing correction and pass/fail gate before looking at held-out results.
5. Use folio-grouped held-out folds. No token from a held-out folio may influence state clustering, mapping choice, hyperparameters or semantic labels.
6. For transcription-robust claims, run ZL, GC and RF independently. Do not pool first and then claim replication.
7. For state-label tests, make nulls selection-aware: if the observed analysis is allowed to choose the best of six A/B/D label assignments, every null replicate must receive the same choice.
8. For suffix/edge-pair tests, match controls on frequency and, where relevant, stem opportunity, token length, line position and section.
9. For contextual morphology, compare the same internal stem across forms whenever possible; otherwise lexical composition can mimic grammar.
10. Report all folds, including reversals and null folds. Never report only corpus-wide p-values.
11. For language/translation tests, freeze token-to-sound/meaning mapping on training folios; decode withheld folios without viewing illustrations; reveal images only after plaintext is fixed.
12. Release negative results and alternative mappings. A coherent narrative is not evidence unless it survives the same frozen pipeline.
13. Independent reproduction should start from raw public transcription plus code, not from this prose report.

12. Exact regeneration recipes

Ending persistence

For each line, reduce token endings to CUPA/STA families. Statistic: number or rate of long same-ending-family runs. Null: shuffle endings within the same line, preserving line length and ending frequencies. Repeat ≥ 1000 . Run independently in ZL/GC/RF.

Boundary mutual information

Measure MI(current token final family; next token initial family). Null: shuffle next-token initial families within appropriate structural blocks/lines while preserving marginals. Compare observed bits to null distribution.

-y/-ain paired family

Define endings before outcome inspection. Strip ending to define stem. Count shared stems/weighted shared occurrences. Compare against frequency-nearest suffix-pair controls; repeat independently across transcriptions.

ABAD state recurrence

Fit the frozen CUPA state classifier/clustering on training folios only. Evaluate all six label assignments if labels are arbitrary. Count A-B-A-D windows on held-out folios. Null: first-order Markov sequences preserving training transition probabilities and test length; selection-aware best-label choice in every replicate.

Shared-stem A/B morphology

Remove frozen opening and closing physical units to obtain internal body. Require identical body across A and B forms. Tabulate opening-only, ending-only and joint changes. Validate transformation ranks on held-out folios.

Directional context

Within same-stem A/B pairs, compare previous-state and next-state distributions separately. Permute A/B form labels within stem family. Report left and right divergence and fold direction.

Particle attachment

For each short token form separately (ol, al, or, r, s candidates), measure left/right state classes, content-like stem diversity, line position and bridge symmetry. Hold out folios. Only assign a functional nickname if the attachment signature reproduces.

13. Failure-resurrection rule

Do not reopen old failures merely because a new representation exists. Reopen only if the old test was near threshold, directionally stable, subgroup-specific, or directly dependent on a segmentation assumption that CPUA materially changes. Hard semantic failures remain closed unless a genuinely new mapping mechanism changes their inputs.

Recommended reopen triggers: nominal $p < .10$; $\geq 3/5$ directional folds; meaningful effect that missed correction; Currier/Hand-specific signal; or direct dependence on EVA character/token boundaries.

14. Translation gate - what must happen before 'decrypted' is defensible

- A deterministic glyph/unit $\rightarrow$ output rule exists and is fully published.
- Rules are frozen before withheld folios are decoded.
- Output on withheld folios is traceable backward to exact source tokens.
- Coherent syntax and semantics appear without illustration-driven word selection.
- Competing languages/encodings lose under the same objective scoring procedure.
- Independent researchers reproduce the same output from raw transcription and code.
- Negative controls and alternative mappings cannot achieve comparable coherence by refitting.
- At least one risky prospective prediction succeeds: the frozen system predicts a previously unexamined passage or externally checkable content.

15. What is presently justified

Defensible current claim: The program has strong evidence for a structured Voynich token system with edge morphology, boundary-control behavior and context-dependent functional states. CPUA-1 reduces dependence on EVA segmentation. The semantic identity, language, authorship and plaintext remain unresolved.

Not justified yet: 'English translation recovered', 'Joachim of Fiore authored it', 'ABAD is the author's intended name for the states', or assigning specific English meanings such as and/of/the to individual short forms.

16. Priority next tests

1. Regenerate every chat-reported CPUA functional result from preserved executable code and write machine-readable CSV/JSON result tables.
2. Complete individual particle test for ol, al, or, r, s with same-class/different-class attachment and held-out folio replication.
3. Re-run six-family boundary model in CPUA units and compare directly with v3.0 EVA-era six-glyph core.
4. Build functional-state classifier that is trained only on physical/distributional features and frozen before section/illustration labels are revealed.
5. Only after functional classes stabilize, test noun-like / verb-like / modifier-like analogies using distributional diagnostics, not translations.
6. Then reopen semantic/language candidates under a new preregistered protocol.

17. Recovery package checklist

- This master report (PDF + editable DOCX).

- ABAD_Voynich_Master_Test_Ledger_v1_to_v3_1.pdf.
- ABAD_Voynich_Master_Evaluation_Report (latest independent evaluation).
- Voynich_CPUA1_Physical_Unit_Alphabet_and_Near_Miss_Reassessment_2026-08-21.pdf.
- Voynich_Transcription_Robustness_and_Reassessment_Audit_2026-08-21.pdf.
- Raw ZL3b, GC2a/v101, RF1b and STA mapping files.
- All scripts/notebooks used for CPUA conversion, state clustering, null simulations and held-out tests.
- Machine-readable result ledger: test ID, date, hypothesis, source hashes, parameters, statistic, folds, p-value, decision.
- README with one-command regeneration instructions and software versions.
- Cryptographic SHA-256 manifest for every file.

18. Source provenance inside the user's research archive

- ABAD_Voynich_Master_Test_Ledger_v1_to_v3_1.pdf - master chronological test ledger through v3.1.
- ABAD_Voynich_Master_Evaluation_Report 8.docx - independent reconstruction/evaluation dated 12 Aug 2026.
- REPORT(4).md - v1.14 homophonic Middle English closeout.
- Voynich_CPUA1_Physical_Unit_Alphabet_and_Near_Miss_Reassessment_2026-08-21.pdf - CPUA-1 rules and first cross-transcription reassessment.
- Voynich_Transcription_Robustness_and_Reassessment_Audit_2026-08-21.pdf - transcription robustness checkpoint.
- Current conversation - later CPUA ABAD/function/morphology/particle sequence; preserved here as chat-reported pending executable regeneration.

19. Numerical snapshot

Measure	Preserved value
Earlier corpus	~38,800 EVA tokens; ~8,100 types
Old ABAD candidate	best z~2.83 vs first-order Markov control
v3.0 six-glyph core	l,n,o,r,s,y; 5/5 folds; +0.03897 penalized gain; p=.00995
v1.14 coverage	26.86% raw; fail vs >=40%
v1.14 uncertainty	73.14%; fail vs <=60%
Old -y/-ain	192 shared stems; 1,208 weighted paired occurrences; matched shared-stem p~.096
CPUA -y/-ain	reported preserved PDF ranks: ZL 2/101, GC 1/101, RF 5/101
CPUA shell interaction	worse held-out log loss in ZL/GC/RF; closed
CPUA ending ABAB	ZL and RF positive; GC null; mixed
Later CPUA ABAD	chat-reported z~5.1 across all three; must be executable-regenerated

20. Final scientific position

The research is no longer well described as a collection of proposed translations. Its strongest component is a progressively constrained structural-linguistic program. The major success is not a word meaning; it is the accumulation of independent evidence that token edges, internal bodies, boundary families and local state transitions are systematically organized. The major unresolved problem is semantic anchoring.

If the later CPUA functional results reproduce from frozen code and raw public transcriptions, the project will have a substantially stronger grammar-level result than the earlier EVA-only phase. That still would not itself constitute a translation. The decisive future transition is from distributional function to a unique, predictive sound/meaning mapping on untouched folios.