TECHNICAL FINDINGS REPORT

ABAD Adaptive ZD30 v3.4

Representation-Matched Surface Latent-Class Test

Date	14 August 2026
Decision	PASS — narrow structural-survival gate
Primary question	Does the ending signal survive without ABAD or POS target classes?
Interpretation boundary	Surface structure only; no language or decryption claim

Executive finding

Outcome: Strong Currier B retained a small but statistically clean ending-to-next-surface-class signal after ABAD targets and POS labels were removed. Currier A remained flat. This shows that the B effect is not solely manufactured by ABAD class definitions, but it does not establish that the mechanism is linguistic.

Final EVA y was selected in all five strong-B folds; r appeared in one fold. The surviving complexity-adjusted gain was only +0.00111 bits/token, versus +0.07142 in v3.3 under ABAD classes. Because the targets differ, that contrast is descriptive rather than a formal decomposition.

1. What changed from v3.3

Component	v3.3	v3.4
Next-token classes	ABAD or POS/onset classes	Four independently learned surface classes
Class inputs	Externally defined labels	Token length, initial/final character, unigrams, bigrams
Ending search	Training-only, up to six	Training-only, up to six
Null	Repeat ending selection	Repeat ending selection after within-segment class shuffle

Leakage control: Class learning used only the token’s own surface form. It did not use the preceding token, ABAD labels, POS tags, dictionaries, semantics, or language identity.

2. Frozen test design

Each outer fold learned four clusters from training-corpus token forms. Held-out tokens were assigned to those frozen centroids. The current token's ending then predicted the following token's learned class. Nested folds selected zero to six ending types, and every added context paid the same BIC-style complexity cost used in v3.3. Two hundred selection-aware nulls were evaluated per corpus.

3. Primary results

Corpus	Tokens	Raw gain	Penalized	p	Selected endings / folds
Voynich strong B	17,194	+0.01512	+0.00111	0.00498	r:1, y:5
Voynich Currier A	9,816	+0.00000	-0.01132	1.00000	none
English Brown	25,000	+0.00329	-0.00501	1.00000	a:2, e:1, f:1
Spanish CESS	25,000	+0.04013	+0.01983	0.00498	a:1, e:4, l:5, o:4, s:4
Portuguese Floresta	25,000	+0.03516	+0.01917	0.00498	e:2, m:2, o:5, s:4
English monoalphabetic	25,000	+0.00909	-0.00263	0.00498	q:1, t:2, y:5
English homophonic	25,000	+0.00000	-0.00484	1.00000	none
English shuffled	25,000	-0.00000	-0.00485	1.00000	none
Surface template	25,000	+0.14970	+0.12685	0.00498	l:5, n:4, o:3, r:5, s:4

Strong-B gate

PASS: Strong B produced +0.01512 raw and +0.00111 complexity-adjusted bits/token; empirical p = 0.00498. Currier A selected no endings, produced a negative adjusted result, and had p = 1.000.

Natural-language comparison

Spanish and Portuguese passed with adjusted gains near +0.020 bits/token—roughly eighteen times the surviving Voynich-B magnitude. English Brown showed a small raw tendency but did not overcome its complexity charge. Therefore v3.4 does not support a specific English match.

The shuffled-English and homophonic controls selected no stable endings and failed. The hard surface-template generator was far stronger than every natural or Voynich stream, at +0.12685 adjusted bits/token.

4. What was learned

Finding	Meaning
B survives independent	The strong-B ending effect is not entirely an ABAD-label artifact.

class construction	
A remains flat	The established Currier strength difference survives representation matching.
y repeats in all B folds	The signal remains highly concentrated in one final glyph.
Romance controls are stronger	Surface-class predictability is compatible with natural morphology, but Voynich is unusually weak.
English fails after penalty	This experiment gives no positive support to a specifically English base.
Template is much stronger	Voynich does not behave like the tested rigid four-template generator.

5. Interpretation

Best-supported conclusion: Strong B contains a real but weak local dependency between current-token endings and the surface shape of the next token. The dependency survives removal of the ABAD target layer, is concentrated in y, and is absent in Currier A under the same engine.

The result narrows the mechanism but does not choose between morphology, scribal/token-generation rules, line-position conventions, or another structured encoding. The mixed language controls show that the test detects some natural-language systems, but not uniformly; corpus morphology and the learned surface representation both matter.

6. Important cautions

Caution	Effect on interpretation
Different target spaces	v3.3 and v3.4 gain magnitudes cannot be treated as directly interchangeable.
Intrinsic-form clustering	The classes describe spelling/glyph shape, not grammatical categories.
Corpus variation	English, Spanish, and Portuguese differ in morphology, tagging source, and document structure.
Small UDHR controls	German and Italian samples are too small for strong negative conclusions.
No historical abbreviation corpus	The most period-appropriate comparator remains untested.

7. Next decisive test

Recommended v3.5: Freeze y and decompose its effect by line position and boundary status: mid-line, line-final, line-initial transition, paragraph start, and across-line transition. Run y alone, the six-glyph core, and the six-glyph core without y under identical selection-aware nulls.

If y remains predictive strictly within mid-line sequences, a morphology-like mechanism remains plausible. If the gain concentrates at line or paragraph boundaries, layout or control-marker behavior becomes the stronger explanation.

Integrity

Specification hash:
941d0dfafe09d19364811443ff995515b238c3a86d36f64d95ede4a5016802d0

Terminal receipt hash:
b4adf071d1052365b402ab35757df0de3aa73cd3db56b13536c6997bc0b9b454

Decision scope: PASS means only that strong B retained positive, significant held-out surface-class prediction under the frozen representation-matched protocol.