TECHNICAL FINDINGS REPORT

ABAD Adaptive ZD30 v3.5

Final EVA y: Boundary Decomposition and Six-Glyph Ablation

Date	14 August 2026
Decision	MIDLINE SURVIVES
Stream	Currier B + Hands 2/3 + non-Herbal
Validation	Whole-page held-out folds; 200 selection-aware nulls

Executive finding

Outcome: The final-y signal is strongest inside ordinary mid-line text, remains positive after its complexity charge, and survives familywise null correction. It disappears across line boundaries. A simple layout-marker explanation is therefore not supported.

The clean mid-line result was +0.01966 raw and +0.00388 complexity-adjusted bits/token, with familywise p = 0.00498. Across-line prediction was slightly negative and non-significant.

1. Why this test mattered

v3.4 showed that y predicted independently learned next-token surface classes. v3.5 asked whether that result merely came from line endings, paragraph starts, or page-specific writing habits. Entire pages were withheld so no line from a test page could enter training.

Frozen element	Operational rule
Targets	Four surface-form classes learned inside each outer training fold
Models	y alone; six-glyph core; six-glyph core without y
Positions	Line-initial, mid-line, into line-final, across-line, paragraph-start
Null	Shuffle targets within page and position stratum; repeat 200 times
Multiplicity	Maximum standardized null across all 15 model-position cells

2. y-only position results

Position	Tokens	Raw gain	Penalized	Familywise p
First token predicts second	1,997	+0.01421	-0.06960	0.00498
Ordinary mid-line transition	13,200	+0.01966	+0.00388	0.00498
Penultimate predicts final token	2,015	+0.00832	-0.07485	0.00498
Line-final predicts next line	2,068	-0.00120	-0.08249	1.00000
Across a marked paragraph start	27	-0.02458	-2.17696	1.00000

Key contrast: Mid-line was the only stratum that remained positive after complexity correction. Across-line and paragraph-start transitions carried no detectable y signal.

Line-initial and into-line-final transitions had positive raw effects, but their smaller samples could not pay the absolute parameter cost. Paragraph-start evidence was especially underpowered: only 27 held-out transitions.

3. Six-glyph ablation inside mid-line text

Frozen ending model	Tokens	Raw gain	Penalized	Familywise p
y only	13,200	+0.01966	+0.00388	0.00498
l, n, o, r, s, y	13,200	+0.02974	-0.02398	0.00498
l, n, o, r, s	13,200	+0.02902	-0.01711	0.00498

Parsimony result: The full core and the core without y captured more raw variation, but both became negative after paying for their extra contexts. y alone was the only model with a positive adjusted result.

Removing y from the six-glyph core reduced raw gain by only about 0.00073 bits/token. Thus y is not the sole carrier of raw information. Its importance is that it supplies the strongest compact one-glyph summary; the other endings collectively require too many parameters for the available gain.

4. What we learned

Finding	Implication
Mid-line survives	The signal is part of ordinary token sequencing, not only page or line furniture.
Across-line fails	Final y is not acting simply as a line-to-line control marker.
Ablation retains raw signal without y	Other core endings share weaker information.
Only y pays for itself	The parsimonious mechanism is highly concentrated.
Whole-page holdout passes	The result is not explained by training on other lines from the same page.

5. Interpretation

Best-supported conclusion: Final y behaves like a compact within-sequence state marker. That is compatible with a grammatical or morphological ending, but also with a nonlexical token-generation state. The present test distinguishes it from simple layout behavior; it does not translate it.

The result strengthens the claim that strong Currier B contains real local organization. It still does not identify English—or any language—and it does not make care/caring a validated translation anchor.

6. Next logical test

Recommended v3.6: Condition y on the identity and surface family of the current token. Compare qok-, she-, che-, dai-, and all other families under whole-page holdout. Test whether y predicts the next class across many stems or only inside a few repetitive carrier families.

Diffuse y behavior across many stems would support a productive suffix-like role. Concentration inside a few high-frequency families would favor a token-template or repetitive carrier mechanism.

Integrity

Specification hash:
e17ab56b5c14be399a3466a7ca24112fe146fd81fbb992b2c5303e8541f98417

Terminal receipt hash:
7a91b8cdcf23014db4ecd7897ae6e17f914e1280a94fb2eca77c3eff1f8bca5c

Interpretation boundary: structural position and parsimony test only; no language identification or decryption claim.