RESEARCH REPORT | 11 SEPTEMBER 2026

Four-Gate Identity Transfer

Cross-folio tournament over all 24 gate permutations

Juan Gabriel Molina
Independent Researcher

Framework by Christine Preedy | Independent testing, reconstruction, and report by Juan Gabriel Molina.

Abstract

This report documents a retrospective reconstruction of the four-gate identity transfer test. The same nonoverlapping source/holdout split used in the broad-role transfer is retained, but the four gate glyphs are represented separately as GATE:o, GATE:d, GATE:n and GATE:y. An order-2 categorical count model trained on 13,461 noninitial source targets is evaluated on 3,262 holdout positions across f22r-f26v. The unchanged four-gate identity ranks first among all 24 permutations, with no tied rival, held-out loss 1.325963 nats per target and 51.5021% exact next-role accuracy. The next-best permutation swaps d and y and has loss 1.743266. A consistent train+test relabeling control makes all 24 mappings tie exactly. The result therefore supports stable cross-folio differentiation of the four gate symbols under the trained representation, but it does not by itself establish that the literal names are Sunrise, Noon, Sunset and Midnight.

Research question and scope

Does the four-way gate distinction learned on the source folios transfer unchanged to the held-out f22r-f26v scoring rows better than every alternative assignment of the four gate identities? The experiment isolates identity consistency of the gate symbols within the predictive code. Literal time-of-day semantics require evidence that is not invariant to categorical renaming.

Keywords: Voynich Manuscript; gate cycle; cross-folio transfer; permutation tournament; held-out prediction; reproducibility

1. Data and gate-expanded representation

Training and holdout data are the same source tables used in the companion codebook-transfer report. The difference is representational: broad GATE is expanded into four distinct symbols according to the raw gate glyph while all other broad roles remain unchanged.

Property	Frozen specification
Training folios / targets	39 / 13,461
Holdout folios / rows	10 / 3,262
Gate identities	o, d, n, y
Expanded labels	GATE:o, GATE:d, GATE:n, GATE:y
Tournament size	4! = 24
Smoothing	alpha = 0.2

1.1 What is and is not being tested

The model can test whether the four distinct gate states retain the same identity relationship between training and holdout. It cannot infer the English meanings from categorical prediction alone. For example, renaming GATE:d to a different word everywhere would not change any numerical probability.

2. Model, baselines and rival mappings

The predictive model is the same order-2 count model used in the broad-role tournament: next-state probabilities are estimated from two preceding states with additive smoothing and first-order/unigram backoff for unseen contexts. No folio ID, line position or astronomical data are features in this model.

Loss = mean[-log P(next state | two previous states)]. Lower is better.

2.1 Exhaustive four-way tournament

All 24 permutations of the four gate identities are applied to held-out contexts and targets under the fixed training model. The unchanged assignment is the identity candidate. The exhaustive enumeration eliminates sampling error in the rival set: every possible four-way reassignment is explicitly scored.

2.2 Symmetry control

The same 24 permutations are then applied consistently to training and testing, followed by refitting. All 24 variants return exactly the same loss; maximum numerical difference is 0.0. That outcome is expected for a purely categorical model and is essential to correct interpretation.

The fixed-training tournament tests transfer of identity; the consistent-relabel control shows why this test is not an independent proof of the time-of-day names.

3. Primary results

Metric	Result
Identity rank	1 / 24
Tied rivals	0
Held-out loss	1.3259632757 nats/target
Exact next-role accuracy	51.5021%
First-order baseline loss	1.3883725540
Unigram baseline loss	2.3276781831
Next-best rival	d <-> y swap
Next-best loss	1.7432661974
Consistent relabel max difference	0.0

The identity candidate improves on the first-order baseline by 0.062409 nats per target and on the unigram baseline by 1.001715. The nearest gate permutation is separated by 0.417303 nats per target. These are predictive-score differences under the frozen model, not probabilities that a gate label is historically correct.

3.1 Five best gate permutations

Rank	Mapping	Loss	Accuracy
1	identity	1.325963	51.50%
2	GATE:o->GATE:o; GATE:d->GATE:y; GATE:n->GATE:n; GATE:y->GATE:d	1.743266	43.26%
3	GATE:o->GATE:y; GATE:d->GATE:d; GATE:n->GATE:n; GATE:y->GATE:o	1.746174	38.81%
4	GATE:o->GATE:o; GATE:d->GATE:n; GATE:n->GATE:d; GATE:y->GATE:y	1.827840	47.12%
5	GATE:o->GATE:y; GATE:d->GATE:o; GATE:n->GATE:n; GATE:y->GATE:d	2.042836	31.18%

Table 1. All 24 candidates are included in the package; the table shows the five lowest-loss mappings.

4. Folio-level behavior

Folio	Rows	Order-2 loss	Unigram loss
f22r	411	1.136926	2.315243
f22v	310	1.062910	2.335222
f23r	365	1.172663	2.266300
f23v	325	1.251784	2.274283
f24r	352	1.312314	2.283388
f24v	326	1.332212	2.330252
f25r	211	1.231650	2.302464
f25v	234	1.154692	2.244394
f26r	353	1.744054	2.451961
f26v	375	1.737863	2.429589

Table 2. Order-2 loss remains below the unigram loss for each holdout folio. f26r and f26v are visibly more difficult than the preceding holdout folios, but they do not reverse the aggregate baseline comparison.

4.1 Supported statement

Within this reconstructed model, the four raw gate symbols are not interchangeable across the source/holdout boundary. The trained identity assignment produces the best predictive alignment on the held-out rows among all 24 possibilities.

4.2 Semantic boundary

The literal labels Sunrise, Noon, Sunset and Midnight cannot be established by this tournament because global categorical renaming is exactly symmetric. External anchors, astronomical constraints, parser behavior or other tests that refer to real-world time are required to distinguish those names. This report deliberately does not import such evidence into the transfer score.

5. Limitations and relation to other gate evidence

5.1 Retrospective status

The original historical runner is unavailable. The current test is a transparent September 11 reconstruction using an explicit model and previously exported scoring rows. It should be cited as a reconstructed transfer test, not as a newly blinded challenge.

5.2 Fixed representation

The four identities are defined from the existing gate encoding. The experiment does not ask whether a different four-symbol family elsewhere in the manuscript could provide an equally useful partition, and it does not compare against all alternative Voynich morphological systems.

5.3 Relationship to astronomical gate claims

Astronomical or historical anchors are logically separate. If such evidence independently identifies one or more gate meanings, this transfer result becomes useful because it shows that the distinct gate identities propagate across held-out folios. But the present 1/24 rank alone must not be used as the anchor.

5.4 Rank is not a p-value

The 24 permutations are an exhaustive finite rival set. Reporting rank 1/24 is exact for that set, but it is not an empirical p-value over a stochastic null population and should not be described as one.

5.5 Reproducibility versus independent replication

Executing the packaged runner again reproduces this reconstructed gate-identity tournament. It is a reproducibility check on the same frozen representation and data, not a statistically or historically independent replication of the unavailable original experiment.

6. Reproducibility, provenance and references

The package contains source and holdout data, a standalone runner, all 24 candidate mappings, folio-level losses, result JSON, protocol, dependency file and SHA-256 manifest.

```
python -m pip install -r requirements.txt
python code/run_gate_identity_transfer.py
python verify_manifest.py
```

Package data fingerprints:
source_occurrences.csv: e7f2dc706365d360edfa9bbf97ca0702db5377eb44ace718515bf326f7470bbd
holdout_scoring_data.csv: 17783c60ee3898aaf0c33b0b24547078009e1a21fea22a3be667100f4b4e96d8

References and source records

[1] Preedy, C. and Molina, J. G. (2026). Four-Gate Identity Transfer Handoff, version 2026-09-11. README.md and PROTOCOL.md.

[2] Same package. data/source_occurrences.csv; data/holdout_scoring_data.csv; code/run_gate_identity_transfer.py.

[3] Same package. results/result.json; results/all_candidates.csv; results/per_folio.csv.

Contributions: Christine Preedy originated the framework and proposed interpretations. Juan Gabriel Molina contributed the testing architecture, computational analysis, reproducibility packaging and report.