

Every model here earned its status.

Model claims are easy to inflate and hard to audit — so we keep a registry that shows production models next to the entries that never earned the label, and a changelog that records the corrections, not just the wins.

SERIES

Model Card MC-001

VERSION

1.0

UPDATED

9 October 2026

STATUS

Published

AUTHORS

Criterica Intelligence Research

Corresponding author: Jack Cosentino

Cleared individually

PRODUCTION STATUS

Each model passes temporal-holdout and calibration review on its own

Every attempt recorded

THE FULL REGISTRY

Models that failed stay on the record, labeled as failed

Judge-level

JUDICIAL ASSIGNMENT

A first-class model feature, not a footnote

Federal and state

COVERAGE

US states and federal circuits, plus Canada, the UK and Australia

One model = one jurisdiction × one case type × one prediction target. The entries that stay experimental, stubbed awaiting sufficient data, or recorded as failed are not overhead — they are the gate working. Status changes only through the same verification protocol that governs every published claim.

METHODOLOGY CHANGELOG

Prediction infrastructure that never corrects itself is not being measured. These are the material governance events on the fleet — including the ones where we found our own numbers wanting and withdrew them.

JUL 2026

GOVERNANCE

Below-gate models demoted from production

A fleet-wide temporal-holdout review identified a small set of production models performing below chance on the most recent evaluation window. All were moved back to experimental status pending label and evaluation audit, and the production count was reduced accordingly.

JUL 2026

CORRECTION

Fleet-wide calibration metric withdrawn for re-derivation

An internal audit found a previously computed fleet-level calibration summary had been evaluated on an unrepresentative sample. The figure was withdrawn from every surface rather than corrected in place, and per-model calibration review continues under the standard gates.

JUL 2026

DATA INTEGRITY

Court-record corpus deduplication verified

The primary corpus was deduplicated on a composite source key after cross-source identifier collisions were found, and the unique count was verified by a gated row-count check before the production data was swapped. The pre-deduplication file was archived, not deleted.

JUL 2026

GOVERNANCE

Models with unverifiable training provenance quarantined

A provenance review found a small model family whose training data could not be traced to verifiable source records. The family was reverted to experimental status and its previously published metrics were withdrawn.

APR 2026

TRAINING

Universal retrain on real data completed

Every production model was retrained exclusively on real filed records, completing the removal of all synthetic training data from the production fleet.

VERIFY IT YOURSELF

Registry inspection

Under NDA: walk the registry itself — model entries, training windows, holdout results, version history, status transitions. What this page says reconciles to that file.

Methodology walk-through

A working session with the team on pipeline, gates, and evaluation protocol. Bring your hardest statistician.

Audit on your book

The strongest test: we score your historical matters with known outcomes and you measure the models against what actually happened. Your data, your baseline.

ARRANGE A VERIFICATION SESSION →

CITE THIS DOCUMENT

Criteria Intelligence Research (2026). The Model Registry and Public Changelog. Model Card MC-001, version 1.0. Criteria. <https://critericaintelligence.com/registry>

```
@misc{criterica_mc001_2026,  
  author = {{Criteria Intelligence Research} and Cosentino, Jack},  
  title = {The Model Registry and Public Changelog},  
  publisher = {Criteria},  
  howpublished = {Model Card},  
  number = {MC-001},  
  year = {2026},  
  month = oct,  
  note = {Version 1.0},  
  url = {https://critericaintelligence.com/registry}  
}
```

DOI: pending