

How the numbers are derived.

Every figure we share, on request, traces to a defined source and a documented gate. This page explains, in honest terms, what the production model count and court-record corpus mean, how a model earns the production label, and which figures come from third parties.

SERIES

Technical Note TN-001

VERSION

1.0

UPDATED

9 October 2026

STATUS

Published

AUTHORS

Criteria Intelligence Research

Corresponding author: Jack Cosentino

01

What counts as a production model

The production model count refers to entries in the Criteria model registry that have cleared the production promotion gate. The registry is the single source of truth: every entry carries a status (production, stub, experimental, suspended, or failed), a training-data vintage, a target variable, and a jurisdiction-case-type scope.

A "model" in this count is one jurisdiction-and-case-type-and-target combination — for example, plaintiff-win probability for commercial cases in a single federal district. Criteria does not build one generalist classifier and report it as many; the count reflects narrow, jurisdiction-specific models built one at a time. Stub and gated models are tracked separately and excluded from the production figure.

How a model earns its way into production

A model is only counted as production after it proves itself against real, recorded outcomes it was never shown during training, and after it has enough real cases for the specific jurisdiction and case type to be trustworthy. Models that fall short stay in stub status and are not counted as production or offered as deployable.

A model that looks near-perfect is treated as suspect, not as success. It almost always means the answer leaked into the inputs, so the model is held back no matter how strong it looks. Any accuracy figures shown elsewhere describe performance against withheld real outcomes, and illustrate the standard, not a guarantee on any single prediction.

The court-record corpus

The corpus is the deduplicated count of real, filed records in the master index. Coverage spans United States federal and state courts, international tribunals, and regulatory enforcement bodies. Specific sources and construction are proprietary. The corpus is normalized, deduplicated across overlapping records, outcome-labeled, and time-partitioned before any model trains on it.

No synthetic, AI-generated, or imputed records appear anywhere in the production training corpus. Because the index is appended and re-deduplicated on an ongoing basis, Criterica does not publish a running total — the current count is available to institutional counterparties under diligence.

Retraining cadence

Production models are retrained against the current master record index on a periodic basis and whenever a material data append changes the underlying distribution for a jurisdiction. Models are versioned independently; a retrain either re-promotes a model under the same standard or flags it for review. Buyers always receive the current production build through the API — there is no stale-model version management on the client side.

Market-size & duration figures

Market-size figures cited on this site (for example, the size of the US litigation-finance market, corporate legal spend, or claims-reserve exposure) are third-party estimates, not Criterica measurements. These figures are drawn from published third-party industry research; full sourcing is available to institutional counterparties on request.

Operational figures such as average case duration and any reduction in duration exposure are illustrative of the modeled effect and depend on case mix, jurisdiction, and intake process. These figures are illustrative of the modeled effect, derived from docket-level historical data; the underlying basis is available to institutional counterparties on request. They are not predictions or guarantees of any individual outcome.

Statistics shown reflect historical or illustrative model outputs derived from real case data. They are not predictions or guarantees of any individual outcome. Litigation results depend on facts, jurisdiction, judge, and counsel, and vary case by case. Model accuracy is subject to selection effects and changing legal dynamics.

The fastest way to verify the numbers is to run them against your own data.

BOOK A DEMO

Criteria Intelligence Research (2026). Methodology: How the Numbers Are Derived. Technical Note TN-001, version 1.0. Criteria. <https://critericaintelligence.com/methodology>

```
@techreport{criterica_tn001_2026,  
  author = {{Criteria Intelligence Research} and Cosentino, Jack},  
  title = {Methodology: How the Numbers Are Derived},  
  institution = {Criteria},  
  type = {Technical Note},  
  number = {TN-001},  
  year = {2026},  
  month = oct,  
  note = {Version 1.0},  
  url = {https://critericaintelligence.com/methodology}  
}
```

DOI: pending