

Candidate-Loss Risk in AI Workflows

Public pre-NDA overview | Low-disclosure bounded reproduction test | Implementation details withheld

Do not believe the claim. Test whether one agreed workflow reproduces the signal.

WHY THIS MATTERS

- The issue is not average AI performance alone.
- Valuable, rare, weak-signal, or emerging candidates can be discarded too early.
- Once discarded, later review may only see the surviving path.

WHAT THIS TESTS

- One agreed workflow using representative or non-confidential records.
- Where candidates become approval, execution, rejection, escalation, or closure.
- Whether governed HOLD and re-evaluation reduce high-value misses without unacceptable review burden.

WHAT THIS DOES NOT REQUIRE

- No production deployment.
- No customer data or production logs.
- No model-weight access.
- No implementation disclosure in public materials.

Candidate state

>

HOLD / classify

>

Re-check / return

>

Governed decision

PUBLIC DISCLOSURE BOUNDARY

- This public overview does not request confidential information, internal architecture, customer data, production logs, source code, model weights, or implementation details.
- Exact thresholds, guard conditions, seed-level implementation, precise metric ratios, and patent-sensitive mappings remain outside public disclosure.

POSSIBLE NEXT STEP

- Run a bounded reproduction test on one agreed workflow.
- Use representative or non-confidential records.
- Decide whether NDA-level technical due diligence is justified.