Methods note / Computer Science

Evaluating AI Agents for Support Ticket Routing

Vladimir Chernega

BestAgent AI Research Lab • November 25, 2026

Abstract

Support ticket routing provides a bounded task for evaluating AI agents because each ticket has an observable destination and an independently checkable outcome. This methods note uses a fixed fixture of 16 synthetic English tickets to illustrate paired scoring. An always-other baseline correctly routed 4 tickets, while a keyword baseline correctly routed 14. The fixture and rules were constructed for the same demonstration, so these counts describe behavior on the fixture. A proposed protocol for operational evaluation separates execution success from verified routing, records outcome-review coverage, and measures arrival time and human reassignment. The protocol specifies how to compare an agent with the current routing system using held-out tickets and the same eligibility rules.

Keywords: AI agents, support ticket routing, benchmark reproducibility, outcome verification, evaluation design.

1. Define the task and scoring rule

Each eligible ticket enters the system once and receives one predicted destination: billing, account access, technical support, or other. A correct prediction matches the reference label. Before evaluation, specify whether ambiguous tickets go to a designated escalation queue or are excluded under a documented rule. Apply the same rule to both systems and count exclusions. The primary outcome is arrival in the correct queue. A separate execution indicator records whether the workflow successfully completed the routing action.

2. Reproduce the reference comparison

The demonstration fixture contains four hand-written tickets in each of the four categories. The first baseline always predicts other. The second lowercases the text, extracts alphabetic tokens, and checks three category-specific keyword lists. It predicts other when no category matches or when several categories match.

On the fixed fixture, the always-other baseline was correct on 4 of 16 tickets (25%). The keyword baseline was correct on 14 of 16 tickets (87.5%). It routed all billing and other examples correctly, and three of four examples in each of account access and technical support. The account example describing difficulty signing in and the technical example mentioning that an API “times out” both went to other. These two errors illustrate why reporting the cases alongside totals matters. All 16 tickets were scored for both baselines.

The ticket texts and keyword lists were written for the same synthetic exercise. The fixture has a balanced task mixture and no customer data, agent runs, or timing observations. Its counts verify the scoring procedure on the released inputs. They provide no estimate of accuracy in a production ticket population.

3. Design an operational evaluation

Freeze the candidate version, model configuration, routing rules, eligibility criteria, measurement period, and primary outcome before collecting results. Use a held-out set of tickets that was unavailable during rule and prompt development. Sample from the intended workflow. If rare categories are deliberately oversampled, report both their counts and the sampling scheme before calculating an overall score. Two reviewers can label difficult cases independently and adjudicate disagreements using a written queue policy. Reviewers should assign expected destinations without seeing the candidate's prediction.

Run the candidate and the existing routing process on the same tickets. Preserve the paired records so a comparison can show which tickets each system handles correctly, which they both miss, and where they disagree. Report counts in a confusion matrix by expected and predicted queue. Choose the size of the pilot for the minimum useful change and the number of observations required in categories where routing errors carry higher costs.

4. Measure outcomes and operational cost

Report execution completion across eligible runs, review coverage across eligible runs, and verified routing accuracy across reviewed runs. Include the denominators. If the review sample oversamples suspected failures, report results as conditional on that selection unless the sampling design supports an adjusted population estimate. For example, 45 correct routes among 50 reviewed tickets and 100 eligible runs yield 90% reviewed accuracy and 50% review coverage. The other 50 outcomes remain unknown.

Measure elapsed time from ticket arrival to acceptance by the correct queue. Define how unresolved tickets are handled at the observation cutoff, and report their count. Record median latency and, when the sample supports it, the 95th percentile. Report manual reassignments, retries and staff time separately. These measurements can be retained as controlled categories, intervals, and pseudonymous run identifiers while ticket contents remain in the support system.

Data and code availability

The synthetic fixture, expected labels, predictions, and standard-library Python runner are available in the version 1.0.0 benchmark package. The original fixture and baselines were produced by BestAgent AI Research Lab. The operational protocol in this note is a proposal for future evaluation.