Support Ticket Routing Baselines

A reproducible comparison of two deterministic routing baselines on synthetic English support tickets.

Version 1.0.0 | 2026-11-25 | BestAgent AI Research Lab

Review: Reproduction against the published fixture

Methodology

We wrote a fixed fixture of 16 English support tickets, with four examples each for billing, account access, technical issues, and other requests. A prediction is correct when its single category matches the expected label.

The first baseline always predicts other. The second normalizes text into lowercase tokens, checks category keyword lists, and predicts other when zero or several categories match.

Results

The always-other baseline classified 4 of 16 tickets correctly. The keyword baseline classified 14 of 16 correctly. The results file includes every ticket, expected category, and prediction.

Reproduce the comparison

Download the Python runner. Run python3 support-routing.py results.json and compare the output with the published JSON. The runner uses the Python standard library and the fixture is included.

Limitations

The fixture is synthetic, small, and authored with task-specific labels. It is useful for checking a reproducible scoring setup. Real support logs, model-based agents, latency, and cost were outside the scope of this comparison.

Conflict disclosure

BestAgent AI built and published the fixture. The comparison used locally implemented reference rules. No vendors supplied results or paid to be included.

Sources and materials

BestAgent Research Methodology

Download fixture and results

Download Python runner