

Worked Example - Learning-Support Request Triage

This is an example for comparison, not an answer to copy. Your own project should use a different problem or independently justified evidence.

Problem

Learners send support requests through one form. Requests are manually read and routed to course, technical, account or wellbeing support. Common requests wait too long, while sensitive requests must be handled carefully.

User and outcome

- Primary user: learner requesting help
- Operational user: support coordinator
- Desired outcome: route routine requests faster without allowing sensitive or uncertain cases to bypass human review
- Baseline: required category selection plus keyword rules
- Unacceptable failure: a wellbeing or safety-related request is routed as a routine technical question

Ai suitability

Language varies and users may select the wrong category. Text Classification could be tested, but it must be compared against the rule baseline. Ai is not permitted to provide wellbeing advice or close support cases automatically.

Proposed workflow

```
Learner message
↓
Input validation and sensitive-term check
↓
Category and confidence prediction
↓
Low confidence or sensitive content? — Yes → Human priority review
↓ No
Suggested support queue
↓
Coordinator confirms or corrects route
```

Inputs

- New support-request text
- Approved category definitions
- Carefully reviewed historical examples, if lawful and appropriate

Sensitive personal content should be minimized, access-controlled and retained only as justified.

Output

- Suggested category
- Confidence or uncertainty indicator
- Reason for mandatory human escalation

The output supports a coordinator. It does not make a final sensitive decision.

Evaluation

- Overall category accuracy
- Recall for sensitive or priority requests
- False-negative review for sensitive categories
- Average routing time compared with baseline
- Coordinator correction rate
- Results by language style and relevant learner conditions

Responsible Ai controls

- Mandatory human review for sensitive terms and low confidence
- No automated advice or case closure
- Clear notice that Ai assists routing
- Access limits and protected logs
- Regular review of category errors
- Easy coordinator override

Decision

Pilot only with fictional or safely approved sample requests. Do not deploy until sensitive-case recall, privacy controls and coordinator workflow are tested.