


Master-Thesis, Master of Science in Engineering

# Trustworthy AI-Supported Data Suggestions

What human reviewers need to judge AI suggestions in enterprise workflows

## Sequential two-phase qualitative design



INCOMING ASSET

Gas-insulated switchgear, ELK-04 (245 kV)

ASSET ID

AST-2026-04803

SERIAL NUMBER

GIS-24-3308-1F

ORDER

ORD-44603

CUSTOMER

CUST-30811

SHIP TO COUNTRY

Germany

SHIPPED

2026-03-13

Standard Stakes

AI-SUGGESTED TARGET SITE

Lippendorf Substation

SITE-DE-1102 · Germany · 50Hertz · 380 kV, 220 kV

SITE ADDRESS

Industriestrasse 11, 04565 Lippendorf, Germany

Approximate location, illustrative only.

UNCERTAINTY

Salesforce and SAP disagree on ship-to country.

WHY THIS SITE

SAP ship to country and Installed Base prior assets point to Lippendorf Substation.

Thinking process

9 steps

SUPPORTING EVIDENCE

3 items

SOURCES THAT AGREE

SAP · Installed Base · Service Records

SOURCES THAT DISAGREE

Salesforce

EVIDENCE

What the system relied on, by source.

| Source          | Field            | Value                                              | Freshness                  | Explanation                                                             | Source                      |
|-----------------|------------------|----------------------------------------------------|----------------------------|-------------------------------------------------------------------------|-----------------------------|
| SAP             | ship_to_country  | DE                                                 | • 2026-03-13 (2 weeks ago) | SAP ship-to country indicates Germany.                                  | <a href="#">View source</a> |
| Salesforce      | contract_country | AT                                                 | • 2026-03-12 (2 weeks ago) | <b>CONFLICT</b> Salesforce contract country conflicts with SAP ship-to. | <a href="#">View source</a> |
| Installed Base  | prior_assets     | Prior asset of same model recorded at SITE-DE-1102 | • site active              | Installed base supports the German site.                                | <a href="#">View source</a> |
| Service Records | last_visit       | Most recent visit to customer site in DE           | • 2026-02 (2 months ago)   | Service records support the German site.                                | <a href="#">View source</a> |

On hold

Send to expert

Reject

Modify

Accept

**Detail view:** one suggestion under review, with source-attributed evidence and the reviewer’s accept, reject, or modify decision.

## What reviewers need to judge a suggestion

- The reviewer stays the accountable decision-maker.
- Evidence is traceable and source-attributed.
- Explanation supports judgement rather than persuades.
- Uncertainty is shown qualitatively, not as a single confidence score.
- Control is preserved on a workable review surface.
- Trust depends on conditions beyond the single screen.

## Main results: from these six needs to validated guidance

- 6 requirement themes, validated across literature and both phases
- 37 user-centred requirements · 7 marked open
- 3 stages of design guidance · foundations · review surface · operation



Scan to explore the work:  
the interactive prototype, a 15-minute audio walkthrough, and the complete thesis

## Study at a glance

28 in-depth interviews (11 experts, 17 reviewers)

≈1,280 coded interview segments

18 design recommendations

14 scripted review cases

## Problem statement

Enterprises increasingly use AI to prefill, recommend, or validate data. Consider a common case: an AI system proposes which site a newly imported asset record belongs to, and a person must accept, reject, or change that proposal before it takes effect.

The reviewer often cannot see the suggestion's basis, its reliability, or its limits. This pushes people toward two failures: over-relying on a system they cannot check, or distrusting one that is right. Both produce wrong decisions.

Here, "trustworthy" does not mean the AI is always correct. It means a reviewer can judge a suggestion in an informed, calibrated way, supported by evidence, reviewability, honest uncertainty signals, and preserved human control. What reviewers need to assess them reliably is not well understood, and this thesis addresses that gap.

## Solution concept

The study follows a two-phase qualitative design. A literature review and 11 expert interviews identified what reviewers need. These findings shaped a low-fidelity clickable web prototype of an asset-to-site review interface, with scripted data and no real AI model.

Seventeen evaluative sessions then tested it. The prototype shows each suggestion with source-attributed evidence, a sober explanation, uncertainty cues at points of conflict, and preserved control.

## Results

The thesis turns what reviewers need into guidance an enterprise team can apply: a validated set of user-centred requirements and the design and implementation recommendations that follow from them. Across all of them the reviewer stays the accountable point of judgement, supported but never replaced by the system.

The set was not only proposed but tested. A separate group of prospective reviewers worked through the prototype and confirmed the requirements, refining them where needed.

Because information is layered, experienced and less experienced reviewers can both start from a clear overview and open deeper detail only when needed. The guidance is framed to transfer across comparable enterprise review settings, with local validation advised before reliance.

## Manuel Engler

Supervisor:  
**Prof. Dr. Christine Grimm**

Expert:  
**Dr. sc. techn. William Gizzi**