LLM vs JEV

Same input, different output.

LLM

You give a prompt, and the model produces a natural language response.

Analyze this support ticket.

output

The customer appears to have **an urgent issue** that would likely require attention from the **billing team**, given the nature of the problem... The customer also appears to be **moderately frustrated** based on their language used and their repeated attempts to resolve it...



It explains and gives context... but I still need to work out the key fields myself.

JEV

You define the questions and output structure up front. The model evaluates each question and returns structured, probabilistic results.

Analyze this support ticket.

Pre-defined questions & output types

Is this urgent?	• True • False
Which team should own it?	• Technical • Billing • Support
How frustrated is the customer?	• Score 0: Not at all • Score 1: Slightly • Score 2: Moderately • Score 3: Very

output

Q1 Urgent?
TRUE | 94% confidence

Q2 Owner?
BILLING | 91% confidence

Q3 Frustration?
2.8 of 3 | 88% confidence



I get the exact structure I need, with confidence scores!

What does this look like in the real world?

AI-Powered Journal Entry Review

From manual journal entry review to structured exception handling.

1 Input: journal entry + context

Journal entry
(GL account, amount, description, date, etc.)

Supporting documents
(invoices, contracts, approvals, etc.)

Accounting policy / guidelines
(internal policies, accounting standards)

Entity & account information
(entity, business unit, account, cost center, etc.)

2 JEV: structured review

Supporting documentation complete?	NO	84%
Complies with accounting policy?	NO	72%
Manual review required?	YES	87%
Risk level?	HIGH	82%

↑ *Illustrative outputs — the criteria are defined up front.*

3 Business logic / workflow

IF *risk level = HIGH*
OR *manual review required = YES*

THEN
→ *Route to exception queue*
→ *Flag for accountant review*
→ *Prepare review summary*

↑ *Structured outputs can trigger rules and downstream workflows.*

4 LLM: explain and summarize

Draft review summary

The journal entry is missing the required supporting documentation... The entry also does not appear consistent with the relevant accounting policy... The item is flagged for review because...

↑ *Use an LLM to turn structured signals into a clear review summary.*

5 Human: final review



Review the signals, read the summary, and consider business context.

- ☐ *Approve*
- ☐ *Request more information*
- ☐ *Escalate*

↑ *AI surfaces structured signals; the accountant makes the final judgment.*

Key takeaways

LLM

Flexible, but less bounded.

- Open-ended, natural language output
- Great for analysis, explanation, and synthesis
- Can produce incorrect or off-track outputs
- Needs additional parsing to use in software

JEV

Bounded, structured, confidence-aware.

- Predefined output types (yes/no, choice, score)
- Optimized for decision tasks and code integration
- Type-safe outputs with calibrated confidence
- The model doesn't decide what matters.
  You define what matters.