INTERNET HYPE ENGINE BIRTHED A NEW ONE

Jev, the Decider.

A new type of model is hyped this week. Jev doesn't write text. It **decides** — typed yes/no answers from messy input.

Fast: under a second

Cheap: $42 per billion tokens

CONTENTS:

EXAMPLES
ALTERNATIVES
YOUR OPERATIONS
VERDICT

Use cases:

REFUND DESK
"€300 back — approve, or escalate to a human?"
yes · no · human

INBOX
"Is this a customer, a supplier, or noise?"
sorted before you look

LISTINGS
"Does this new listing break the rules?"
ok · broken

LEADS
"Real inquiry or spam — which rep gets it?"
routed in ms

Strict decision, thousands a day.
Life saver.

Alternatives:

PLAIN RULES — 1960s

"if price > X" — free, instant, breaks on messy input

NAIVE BAYES — 1990s

the 90s spam-filter classic — fast, free, needs training data

YOUR OWN TRAINED CLASSIFIER — 2010s

still wins when you own the data

LANGUAGE-NATIVE DECIDERS — 2025

Jev
Laya (421M, on your machine)
SemIf ($1k GPU)

> **The new trick:** reads like a chatbot, answers like a form — **can't ramble.**

To use in your project, run four questions:

a Plain rules enough? **Then skip AI.**

b Yes/no decisions from messy input? **Decider territory.**

c Do they need thought and nuance? **Chatbot — thinking prices.**

d Will volume 100× once it's cheap? **Pennies × millions. Architect first.**

Watch it work.

IDEA A · REVIEW FILTERING

Customer reviews flagged **genuine / spam** — yes/no from messy text (b) · no nuance (c) · every review, every day (d). → **Decider.**

IDEA B · CONTRACT DRAFTING

Needs reading, judgment, nuance (c). No typed answer exists.
→ Chatbot prices — or don't build.

Same test. Different tool.

THE RULING

Keep it in the toolbox.

Not a revolution — a tool. For special cases, **wonderful to have**.

PROOF IT'S REAL · 2,024 PRODUCTION DECISIONS REPLAYED

97.3% accuracy

151ms — the old way: 4s

$0.17 — the old way: $5.51

Which of your processes reduce to four options?