


A SHORT FIELD GUIDE · 1 OCTOBER 2026

Next Big Thing in AI - SYSTEM ONE

When the useful output is a decision, why ask a giant model to write a paragraph?



By Viktor Zaika

Sometimes what you need is basically a very intelligent if statement.

01 / THE IDEA

A model that makes decisions.

TypeSafe calls Jev its first System One model. You give it state and a bounded question. It returns a typed choice, score, or yes/no probability. It is designed to sit inside software, where your code decides what happens next. [TypeSafe's announcement](#) ↗

Request



Jev




Decision



Action

Think of the expensive LLM call currently doing nothing more than answering “suspicious?”, “which queue?”, or “should a person review this?” Those are the calls worth examining first.



TypeSafe's real workflow diagram: triage, disposition, containment, playbook. Company illustration; [original and explanation](#) ↗

The small decision is a layer in the product, not the whole product.

02 / WHERE IT FITS

One architecture, three jobs.

Route Choose a queue, model, or next step from known options.

Score Rate urgency or severity against a defined scale.

Gate Estimate whether a bounded statement is true; send uncertain cases to review.

Liquid's [d1 documentation](#) calls these Choice, Score and Noul. That gives this approach a second implementation to inspect.



TypeSafe's own accuracy-versus-cost comparison across four company-selected workflows. [Read its methodology and caveats](#) ↗

A [Reddit user reported](#) that moving an initial “anything of concern?” check from an LLM to Jev reduced that first step from 3–6 seconds to milliseconds. This is an anecdote, not a controlled benchmark.

03 / USE JUDGMENT

Which calls should move?

Keep a generative LLM for writing, explanation, transformation and summaries. Consider a decision model when the final useful output is only **true**, **“sales”**, or **0.83**. The [Hugging Face community explainer](#) describes this bounded-output distinction.

Start with one repeated, low-risk decision. Compare latency, cost, false positives and false negatives against the current route. Keep a fallback for uncertainty. Do not mistake a well-typed answer for a correct answer.

Jev is new. Its striking published benchmarks are TypeSafe’s own, and the [discussion is mixed](#). More independent testing is needed.

Look at your LLM bill. How much of it exists just to make tiny decisions?

Sources

[TypeSafe · Introducing System One Models & Jev ↗](#)

[Liquid · Decision Models documentation ↗](#)

[Reddit · User anecdote and discussion ↗](#)

[Hugging Face · Community Jev explainer ↗](#)

The diagrams and evaluation chart are TypeSafe graphics. Claims and comparisons are attributed to their sources.