



SYNAPTIX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

We tested *JEV* across our Accelerators

Four decision points. Real numbers. Here's what we found.




01 / Document classification

100%

accuracy on 78 classifiable financial documents



Metric	Gemini	JEV
 Accuracy	-	100%
 Median latency	~9 sec / doc	~1 sec / doc
 Median cost / call	\$0.0046	~\$0.00015
 Repeat stability	not deterministic	100%



82 documents tested, 78 had usable OCR text

Manually verified against Gemini on the same set

— 02 / Model routing

A decision we had never automated before

91.3%

overall accuracy picking the right sized model,
across 92 realistic test questions



98%

on easy calls



87%

on medium calls



85%

on hard calls



Half a second or less per decision
Around \$0.00004 per decision

100% repeat stability

— 03 / SynSights intent & query understanding

100%

accuracy, matching our current pipeline exactly



Metric	Current LLM	JEV
 Accuracy	100%	100%
 Median latency	~1.46 sec	~0.50 sec
 Median cost / call	\$0.00110	~\$0.000022
 Repeat stability	not deterministic	100%




30-question golden set, grounded in the Panama Loyalty catalogue
Confirmed by a blind LLM judge and DeepEval as a separate check

04 / Lattice-level intent classification

85%

accuracy against the current classifier's 81%



Metric	Current Classifier	JEV
 Accuracy	81% (48/59)	85% (50/59)
 Median latency	~2 sec	~0.4 sec
 Median cost / call	\$0.53/1,000 Decision	\$0.058/1,000 Decision
 Repeat stability	not deterministic	100%



Scored on domain, purpose and modality per request
Domain remains the hardest axis for both models



SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

How was your experience with JEV?

Tell us in the comments. We read every reply

corp@synaptyx.ai

