



LLMs for Workflow Automation: Soft Logic That Works

Why Soft Logic Beats the AGI Race

You don't need AGI to reclaim your calendar. You need a way to turn intent into execution without losing control. Soft logic does exactly that inside Google Workspace.

I used to spend three hours every Monday morning pulling data from five Google Sheets, summarizing project updates from Gmail threads, and writing status reports with the same template. The breakthrough wasn't a new app or an assistant, it was treating large language models as soft logic controllers that sit between my intent and the tools, much like PLCs do in factories.

In short, LLMs act as soft logic for task automation: they handle semantic interpretation and decisioning, pair well with Google Scripts, and let you automate digital labor while monitoring for drift so reliability stays high.

The Hidden Cost of Manual Digital Work

Every time you format a report, categorize emails, or extract insights from spreadsheet data, you're doing logic a machine can do. The cost isn't just minutes, it's the context-switching that drags down creative and strategic work. Traditional automations cover triggers and actions but stall when a task needs interpretation, so you keep doing the thinking: what matters in a thread, which data points to elevate, how to phrase a response that matches your intent.

How Soft Logic Actually Works

Think of a PLC in a plant: inputs in, logic applied, outputs out. PLCs deal in binaries and numbers; knowledge work needs flexible interpretation. That's soft logic, models that read natural language, weigh context, and make bounded decisions. LLMs excel here when you treat them as controllers, not co-authors.

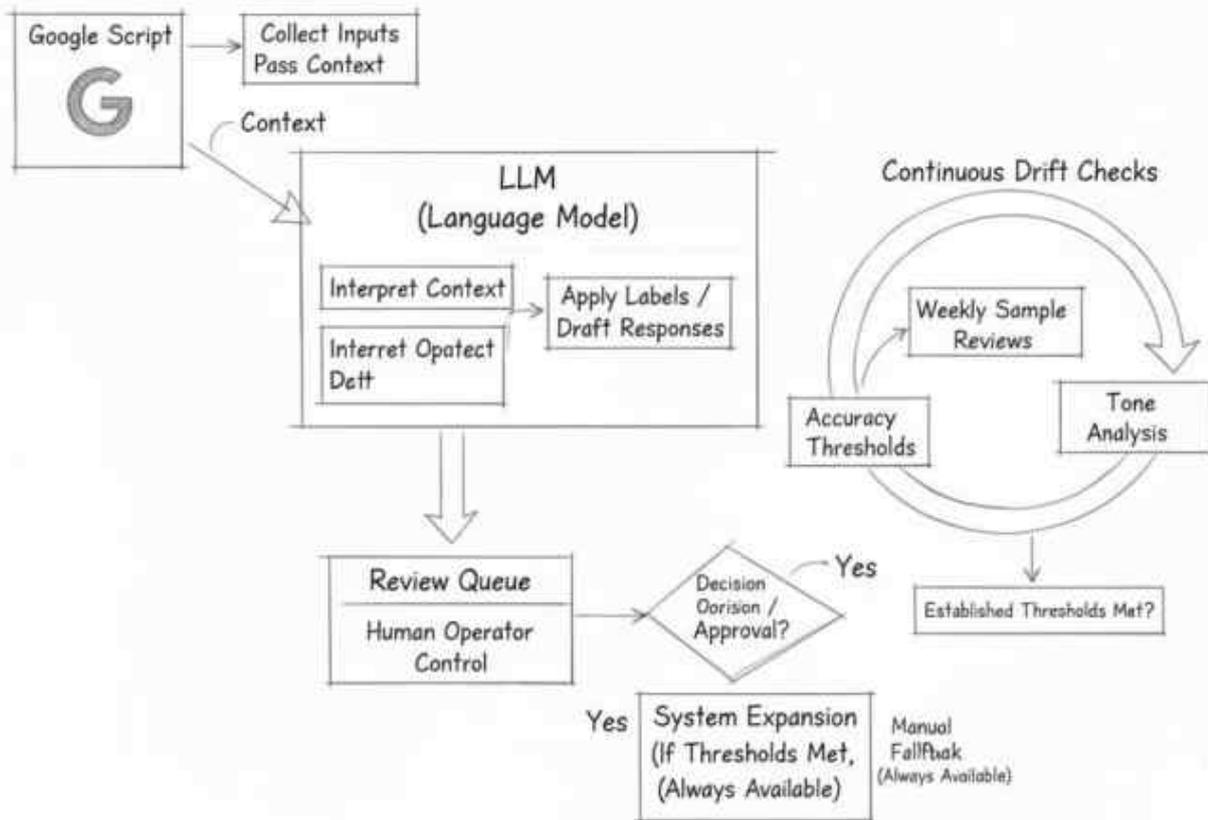


Treat the model as a controller that interprets intent, not as an author you debate with.

Here's a concrete example: a Google Script monitors a project spreadsheet, flags overdue tasks, drafts follow-up emails based on task context and recipient relationship, and queues them for my review. The LLM handles interpretation (urgent vs. routine) and language (tone and phrasing). I keep control through intent tracing, checking that outputs consistently map back to the goals I set.



LLMs for Workflow Automation: Soft Logic That Works



You want time back without hidden risks. The friction is brittle click-ops that fail on nuance. The belief: LLMs as soft logic interpret intent while you set guardrails. The mechanism: Google Scripts plus an LLM with intent tracing and drift metrics. The decision conditions: reversible pilots that clear simple thresholds on accuracy, time saved, and tone fit before you scale.



Measuring Drift Before It Derails You

Catastrophic failure is rare; slow drift is common. Build measurement into every process. For the follow-up workflow, I track response rate, tone feedback, and time saved. If any metric slips outside tolerances, I tune the logic.

Drift, not failure, is the enemy, so measure it and correct early.

I also run a simple audit: every Friday, I compare a sample of automated outputs with last month's manual versions. If outcomes would diverge more than 20% of the time, I adjust prompts, rules, or thresholds.

Why This Isn't AGI (And That's Good)

You don't need general intelligence to automate your business workflow automation processes. You need reliable, measurable automation of specific cognitive steps. AGI promises everything and delivers complexity. Soft logic delivers specifics: shorter intent-to-execution cycles, lower cognitive load, and observable performance.

Last month, a founder I work with automated her weekly investor update. Four hours on Sunday became 30 minutes of review. The system pulls metrics from her dashboard, summarizes progress against goals, and drafts in her voice. She keeps editorial control while stripping out mechanical work.

Your First Reversible Test

Start with a predictable, text-heavy task so you can see where interpretation matters. Here's a compact way to pilot without locking yourself in:

- Pick one workflow (e.g., email categorization) and define success metrics: accuracy, time saved, and tone fit.
- Wire a Google Script to collect inputs, pass context to an LLM, and apply labels or drafts while preserving a review queue.
- Instrument drift checks (weekly sample review, thresholds for accuracy/time/tone) and keep a manual fallback.



- Run for two weeks; only expand if it beats your thresholds with low maintenance overhead.

Soft logic gives you leverage today: clearer intent-to-output paths, smaller cognitive taxes, and feedback loops that keep the system aligned with your judgment. Small, measured pilots compound into dependable workflow gains while you stay in the loop.