

OUR APPROACH

Ontology-First AI for Small Business

How sig.ai builds decision systems that know how a specific business actually runs — and keep people in charge of every decision.

The problem with plausible text

Most AI products are built to generate. Ask a question, get fluent prose. For drafting an email, that is enough. For running a business, it is not: fluent prose about your pricing rules is worthless — or dangerous — if it does not match how your pricing actually works.

Small businesses feel this gap more than anyone. An owner-run company's most valuable operating knowledge rarely lives in documents an AI can summarize. It lives in judgment: when a discount is allowed, which customer needs special handling, when the owner quietly steps in, which vendor exceptions are real. Generic AI cannot see any of that — so it guesses, confidently.

We take a different route. Before any system reasons about a business, we build a faithful model of how that business works, grounded in its own records. We call the approach **ontology-first**, and it rests on three definitions simple enough to state in a sentence each.

Knowledge = Ontology + Facts

An **ontology** is a working map of how a business actually operates: its jobs, customers, vendors, pricing rules, exceptions, escalation paths — and how they connect. Not an org chart or a folder of SOPs, but a structured model in which “a callback on a warranty job routes differently than a new install” is a first-class statement a system can check.

Facts are evidence attached to that map. Each fact carries a source reference: the invoice, ticket, CRM note, or record it came from. A fact without a source is not a fact in our systems — it is a guess wearing a suit.

Together, ontology and facts are **knowledge you can audit**. Anyone can ask “why does the system believe this?” and get a real answer: here is the claim, here are the records behind it, here is who confirmed it.

Insight = Reasoning over Knowledge

Insight is what reasoning produces when it runs over that audited knowledge — never over statistical vibes. This has a discipline attached, and the discipline is the product:

- **No source, no claim.** Every claim a system serves must trace to evidence. Claims that cannot are not served.
- **Confidence is earned, not asserted.** Claims carry risk classification and confidence scores, and higher-risk claims face higher bars.
- **Refusal is a feature.** When the evidence was never captured, the honest answer is “we don’t know yet” — and the system says exactly that, instead of improvising.

The practical consequence: correction beats generation. It is far more valuable to put a specific, source-linked claim in front of the person who knows the truth — “is this threshold right?” — than to ask them to write a manual from memory. People are excellent editors of concrete claims and poor authors of exhaustive documentation. Our systems are built around that asymmetry.

Action = Insight × Goals + Constraints

Insight alone does not tell anyone what to do. Action requires two more ingredients that belong to the business, not to the AI: **goals** (what you are optimizing for) and **constraints** (budgets, policies, regulatory boundaries, risk tolerances — the lines that must not be crossed).

The same insight yields different actions under different goals and constraints, which is why we insist those live explicitly in the system rather than implicitly in someone’s head. When constraints are code, they are enforced every time, reviewable by anyone, and changeable on purpose instead of by accident.

$$\text{Decision} = f(\text{ontology, facts, goal, constraints})$$

The function is evaluated by people

The equation above has a quiet but non-negotiable property: **a person evaluates it**. AI agents in our systems draft, reconcile, flag, and propose. They do not decide. Every consequential action is a human-

owned proposal; ambiguity escalates to a person by design; and the riskier the question, the more deliberately human the path becomes.

This is not a limitation we tolerate — it is the design. Owners and operators hold context no record captures, accountability no model can carry, and judgment their businesses were built on. The system's job is to make their decisions better-informed and faster to reach, not to make the decisions.

In practice: knowledge is assembled from real records and verified with the people who hold the context; insight is served only with sources attached; action happens inside explicit goals and constraints; and a person signs every decision that matters.

Built with the SMB ecosystem

Small businesses are not small enterprises. Their records are messier, their processes more personal, and their margins for error thinner. Software built for enterprises — or trained on the average of the internet — routinely misreads them.

So we do not build alone. We work inside the ecosystem that small businesses already trust: owners, operators, brokers, advisors, lenders, and the professionals who move with them through moments of change. They keep the ontology honest — it must reflect how these businesses really run, not how software wishes they ran — and they are why a human is at the core of every system we ship.

Where this runs today

- **Data engineering & measurement** — building the evidence layer teams need before automating anything: pipelines, definitions, and metrics people can actually trust.
- **TransferOS** — key-person diligence for owner-dependent business transitions. An owner's operating knowledge is reconstructed from records, verified through targeted correction, and delivered to a buyer as auditable claims rather than folklore. (sig.ai/transferos)
- **Text2Knowledge** — the ontology layer as a service: turning unstructured business records into structured, auditable knowledge (text2knowledge.com), with T2K Studio (studio.t2k.ai) as the hands-on workspace for building and refining ontologies.

What this is not

Honesty about boundaries is part of the method. Our systems provide operating knowledge and decision support. They are not legal, tax, HR, valuation, or quality-of-earnings advice, and they do not

replace the professionals who provide those. A reconstructed claim is labeled as reconstructed; an uncertain claim is labeled as uncertain; and no score or system output is a substitute for a person's judgment about their own business.