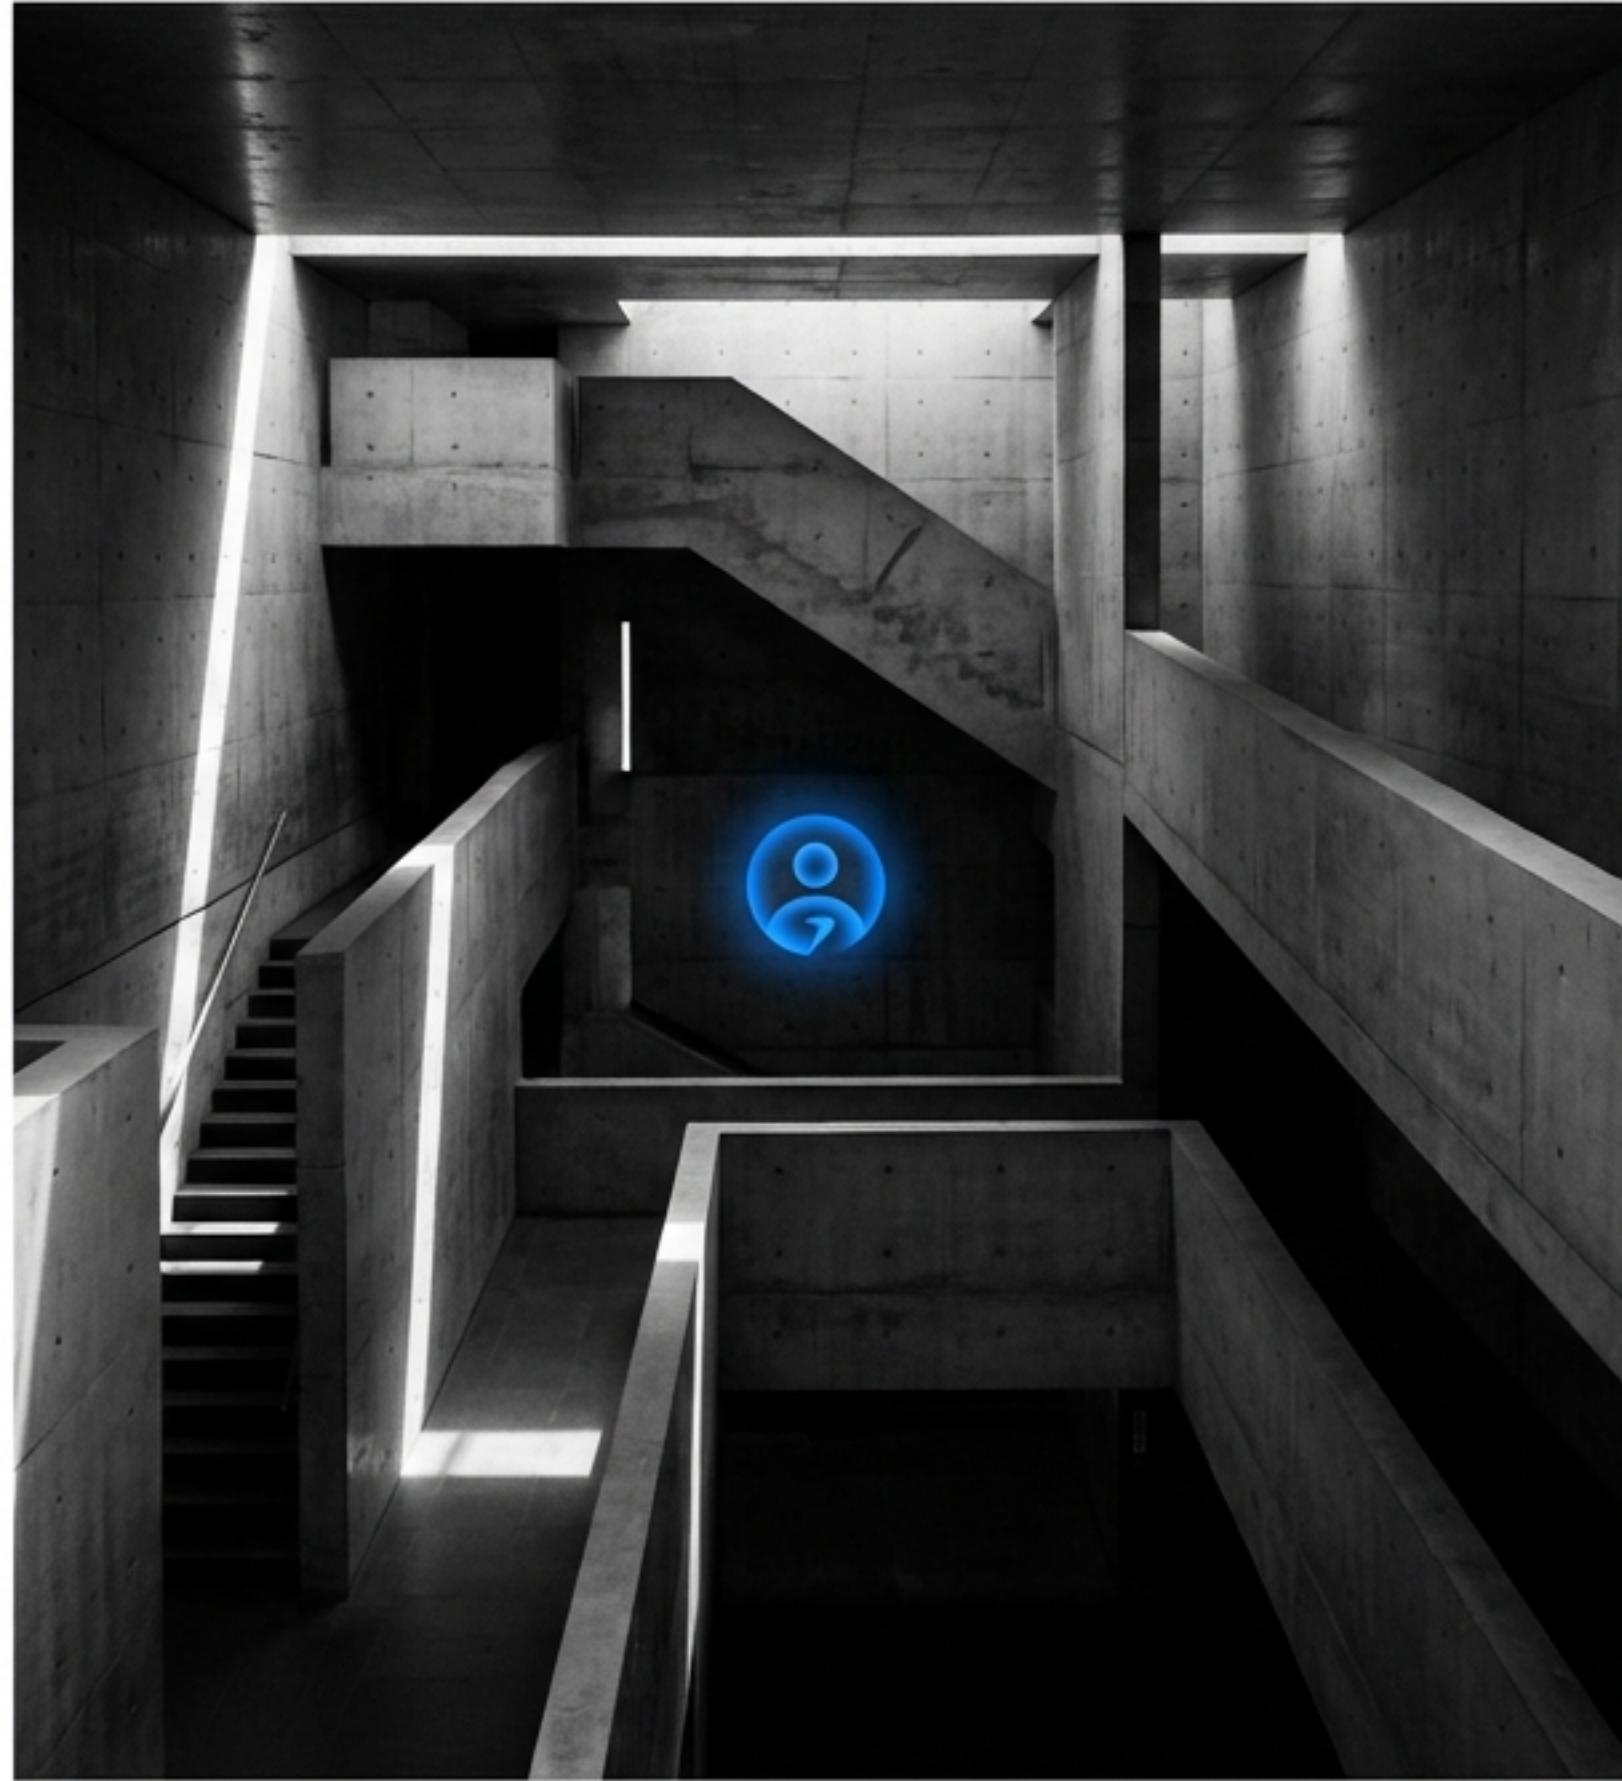


Orchestrating Intelligence

A Facet-Based Taxonomy
for Contextual AI Assistance

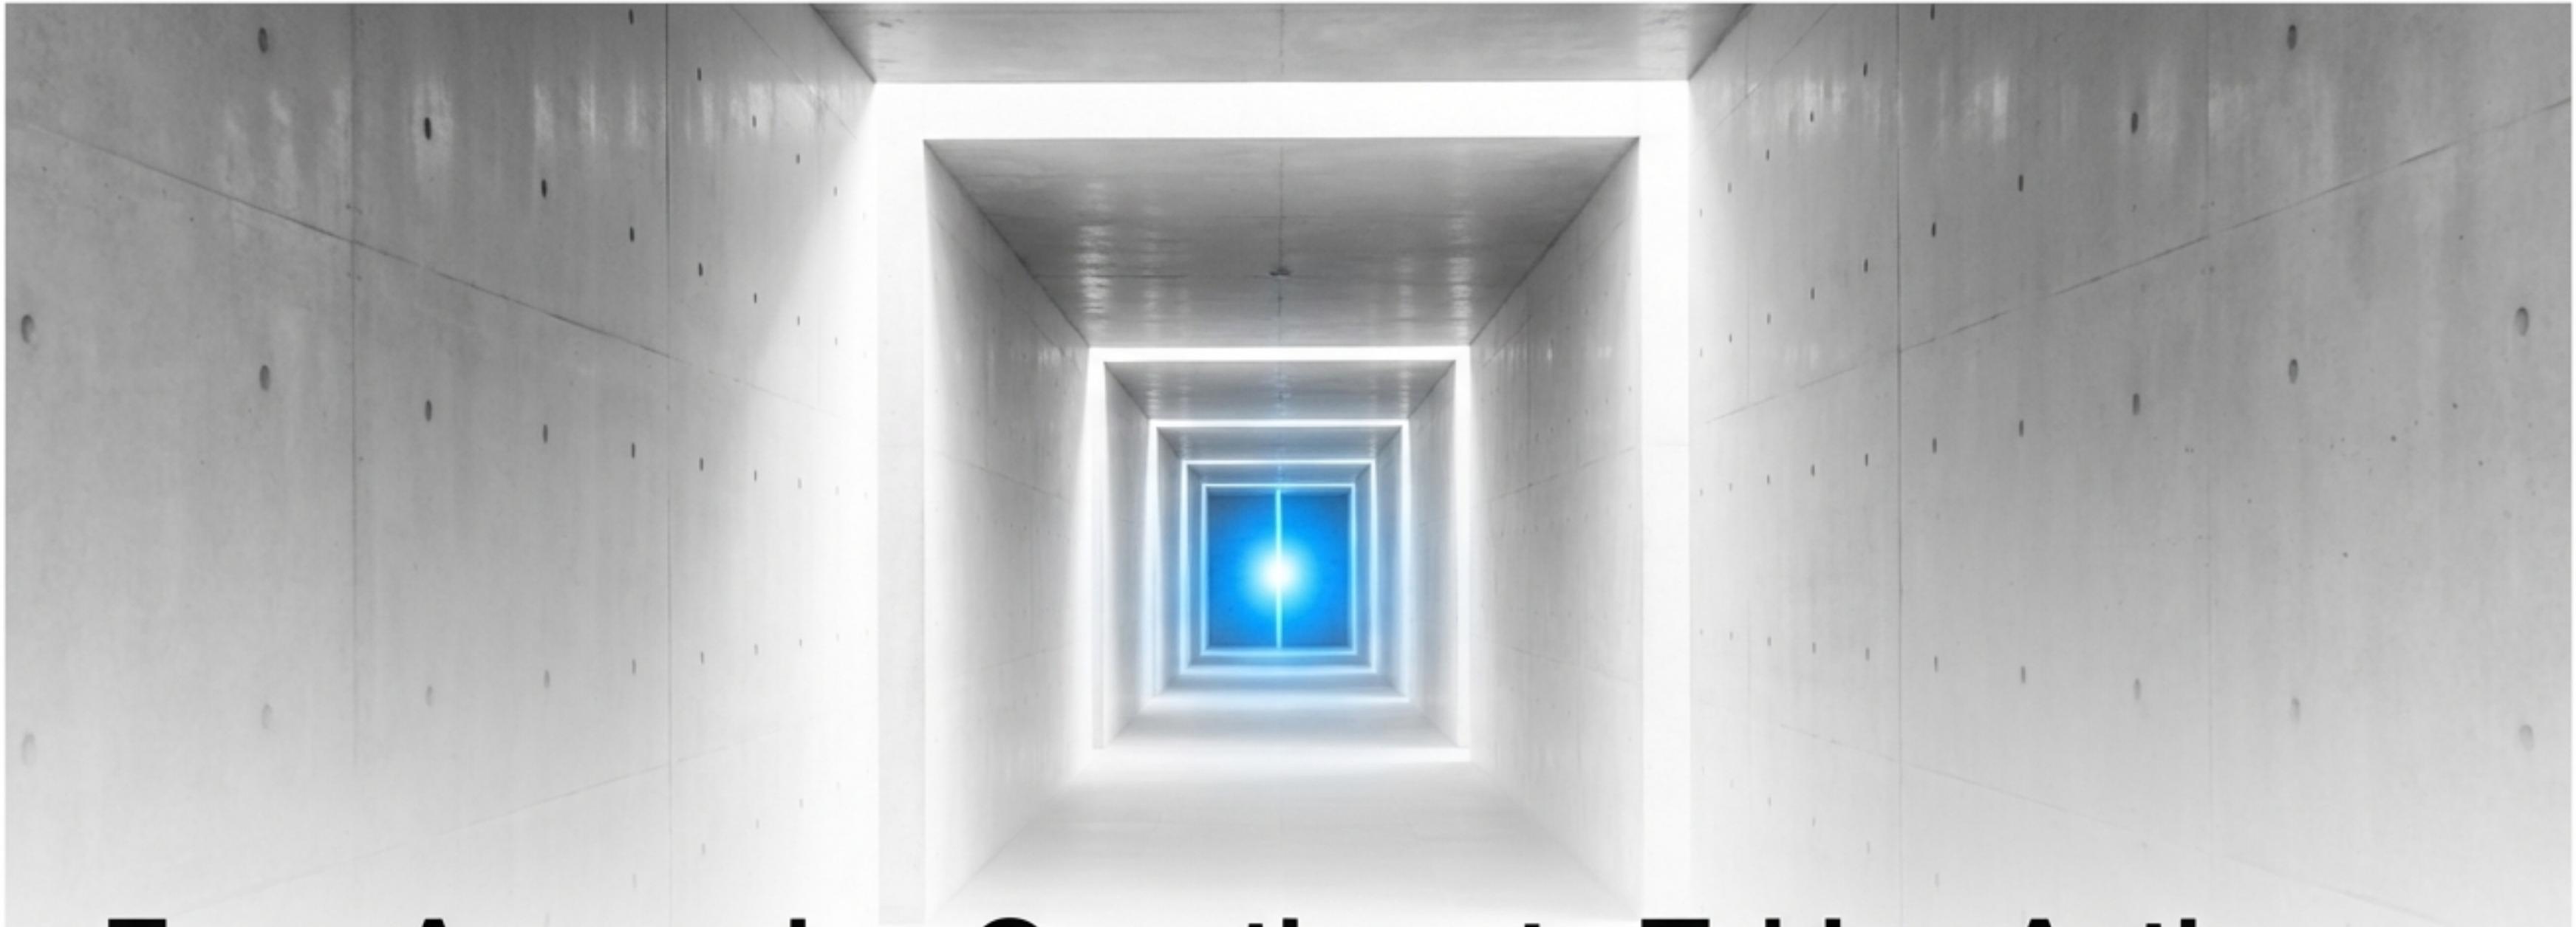


The Failure of Generic Help

Current AI assistants **lack true context**. They cannot distinguish between a novice accountant and an expert founder, a simple form and a complex workflow.

This leads to **user frustration, friction**, and **abandonment**.

They offer **one-size-fits-all answers** in a world of unique needs.



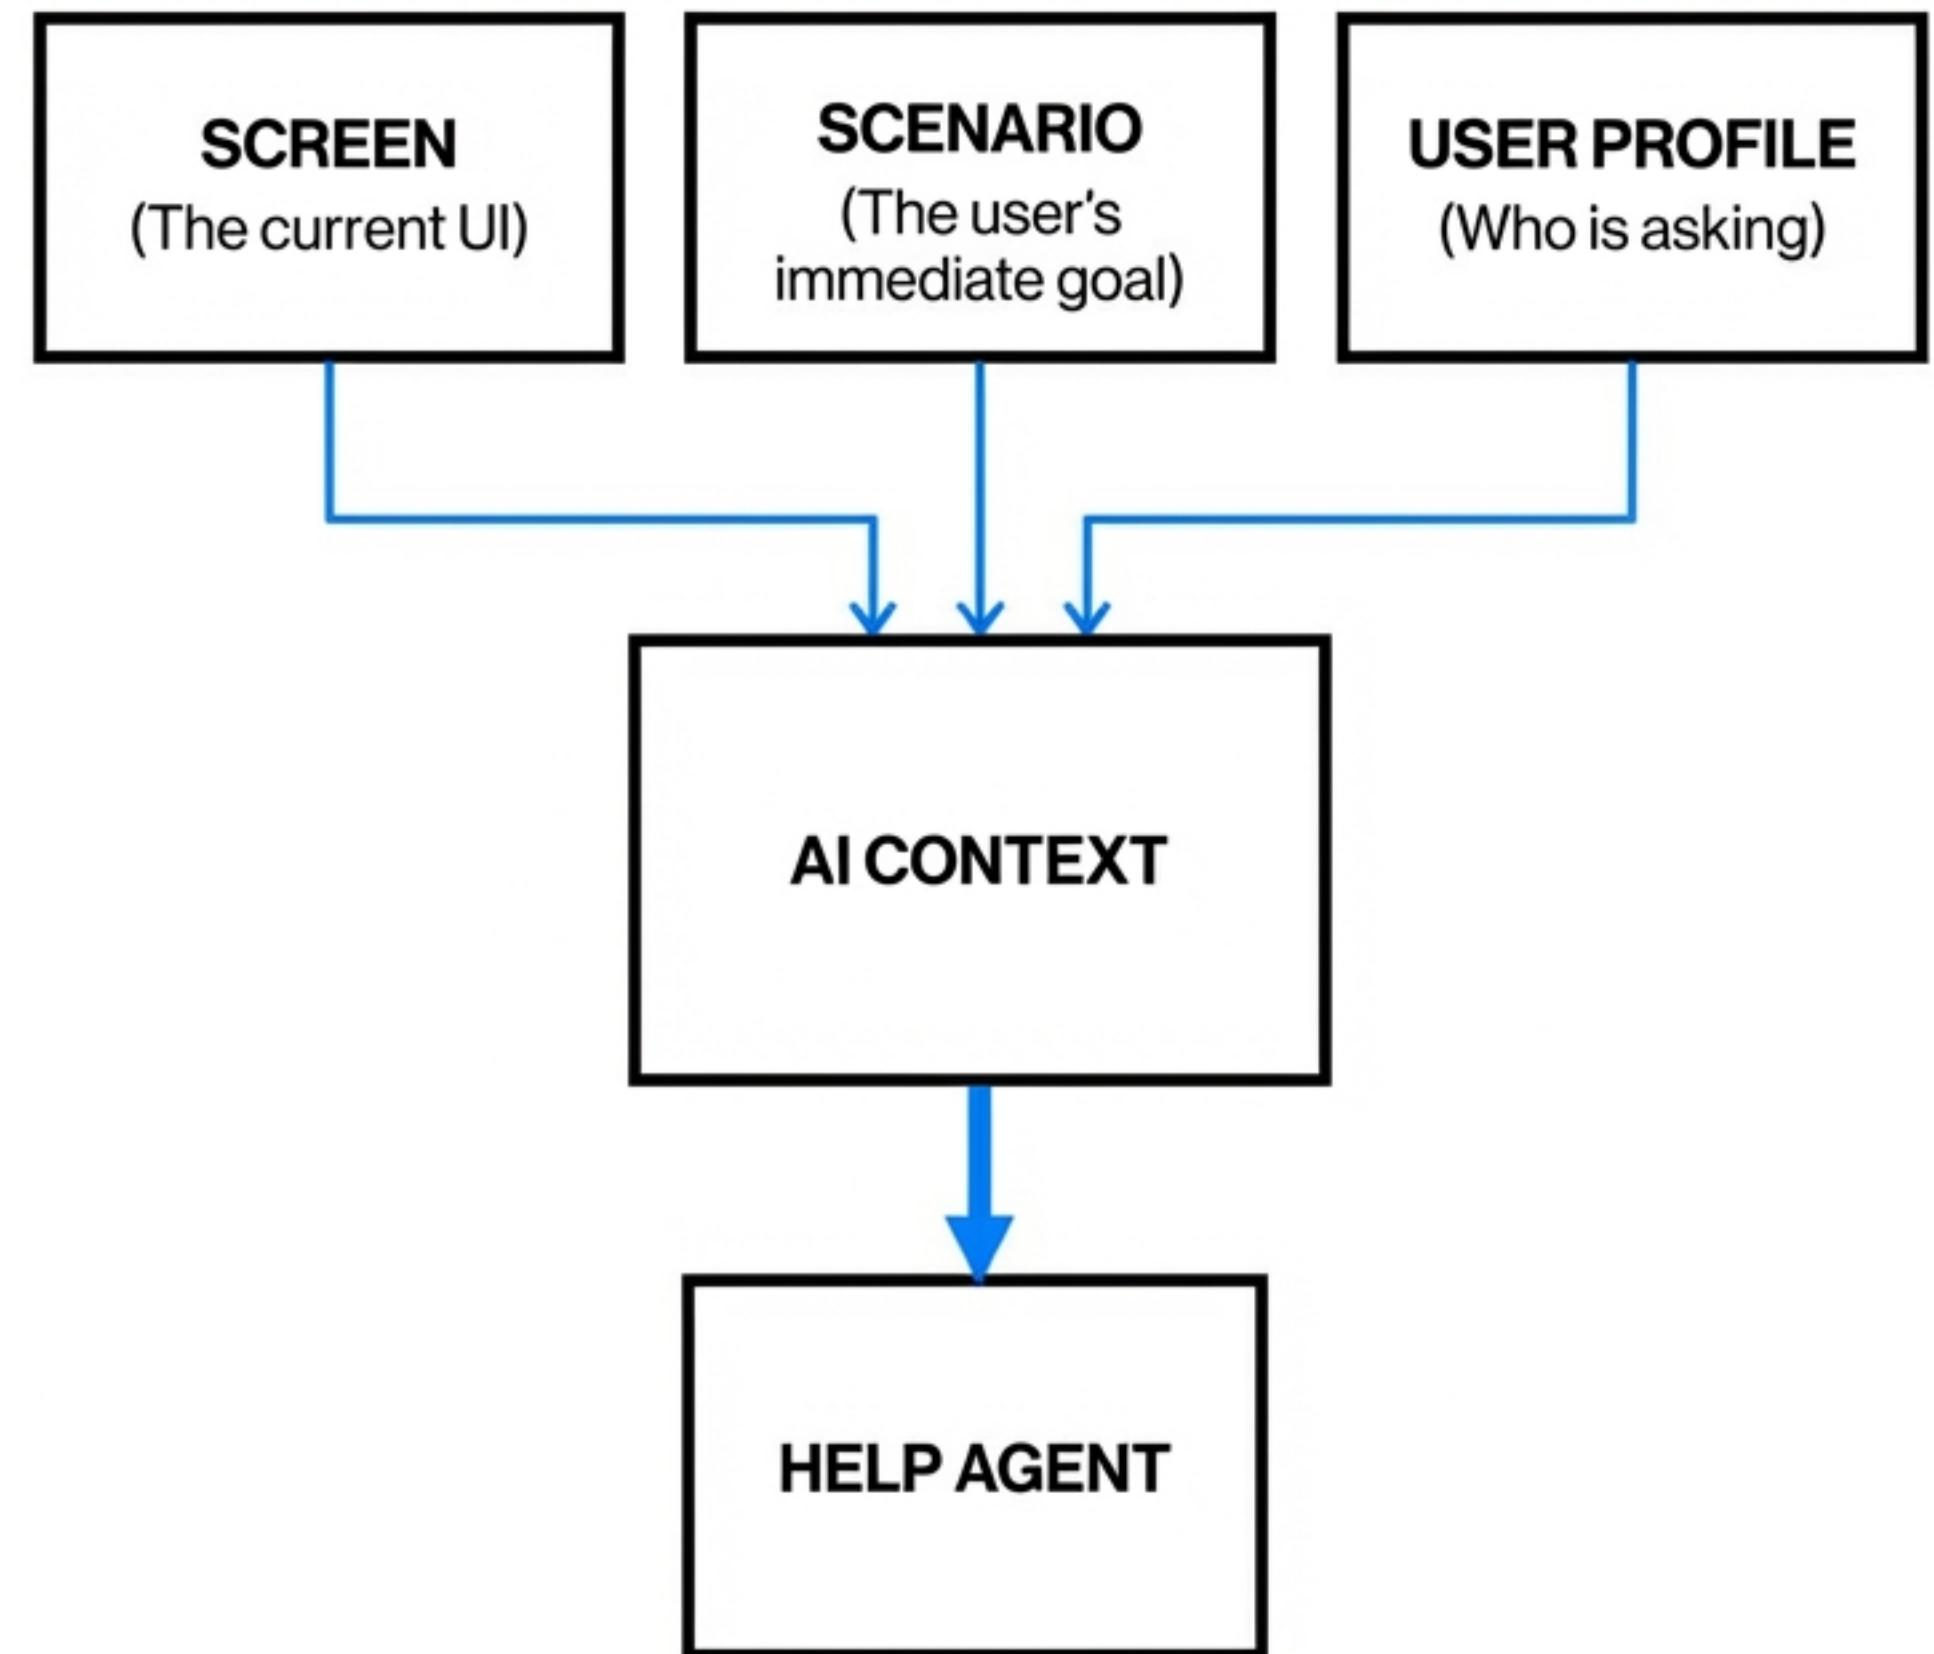
From Answering Questions to Taking Action

The vision is an AI that doesn't just **show**, it **understands**. It comprehends the screen's anatomy, the user's intent, and their unique profile. This deep understanding allows it to provide precisely the right assistance, at the right time, across three levels: **See**, **Go**, and **Do**.

The Blueprint for Understanding: A Facet Taxonomy

We achieve this vision not with a black box, but with structure. A taxonomy of ‘facets’—descriptive tags from three core sources—allows us to systematically map the entire UX landscape and every possible user context.

It is the architecture of empathy.



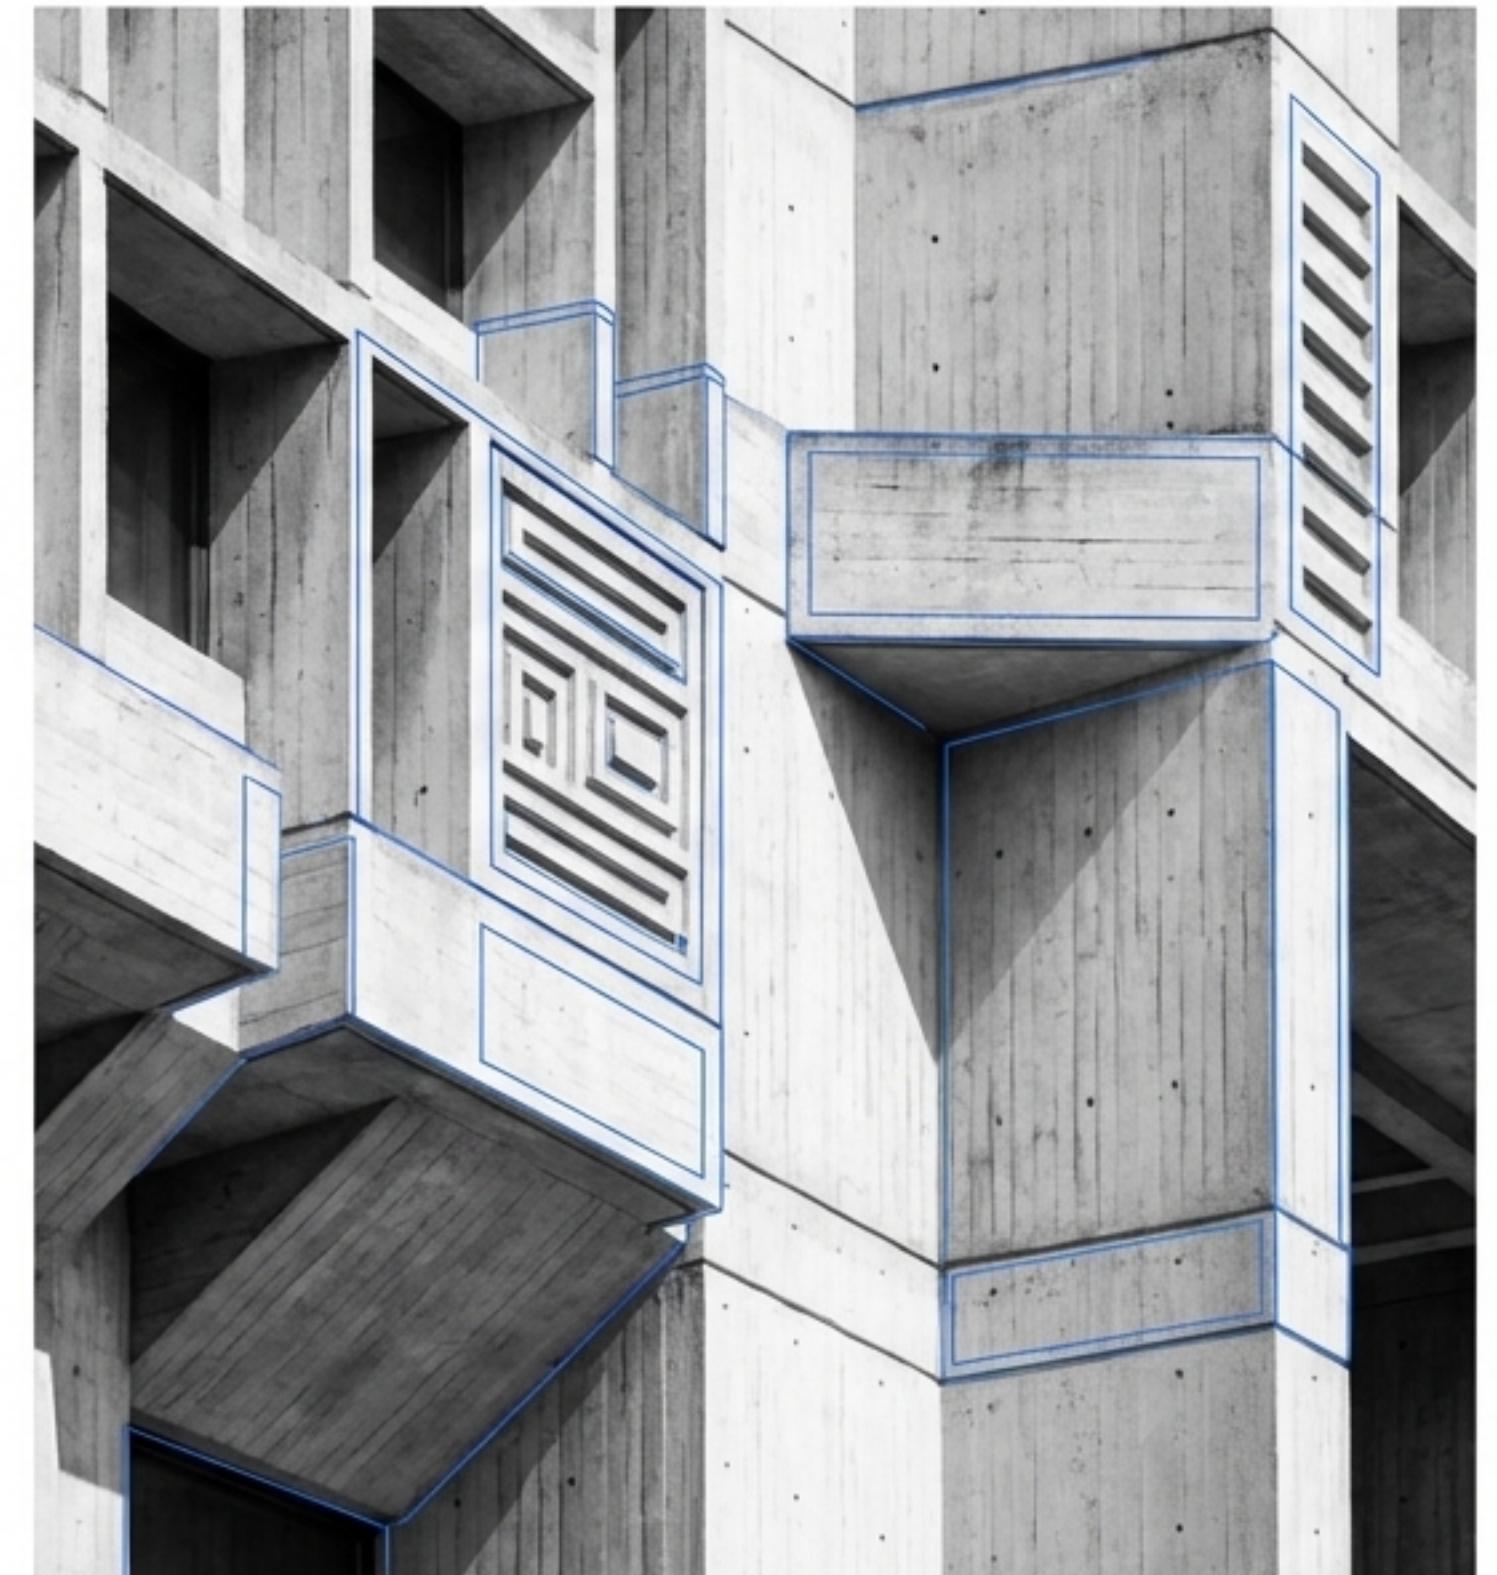
Pillar I: The Anatomy of the UI

First, the AI analyzes the screen itself. It deconstructs the interface into a structured set of functional, interactional, and technical attributes.

Functional: `Module` (Billing, CRM),
`Page Type` (Form, Dashboard, List),
`Main Action` (Create, Update)

Interaction: `UX Pattern` (Wizard, Master-Detail),
`Complexity` (Simple, Intermediate, Complex)

Ergonomics:
`Information Density` (Low, Moderate, High),
`Target Device` (Desktop, Mobile)



Pillar II: The User's Intent

Next, the AI deciphers what the user is trying to achieve *right now*. It analyzes their query, recent actions, and position within a workflow to understand the job-to-be-done.

Key Facets:

Goal: Job-to-be-done (Create invoice, Follow up with client),
Flow / Journey (Onboarding, Month-end closing)

Context: Immediate Context (Error, Blockage, Discovery),
Urgency (Low, Medium, High)

Risk: Expected Risk (Low, Medium, High),
Result Type (Draft, Final action)



Pillar III: The User's Profile

Finally, the AI considers who is asking. It draws on the user's role, history, and preferences to tailor the form and depth of assistance.

****Identity****

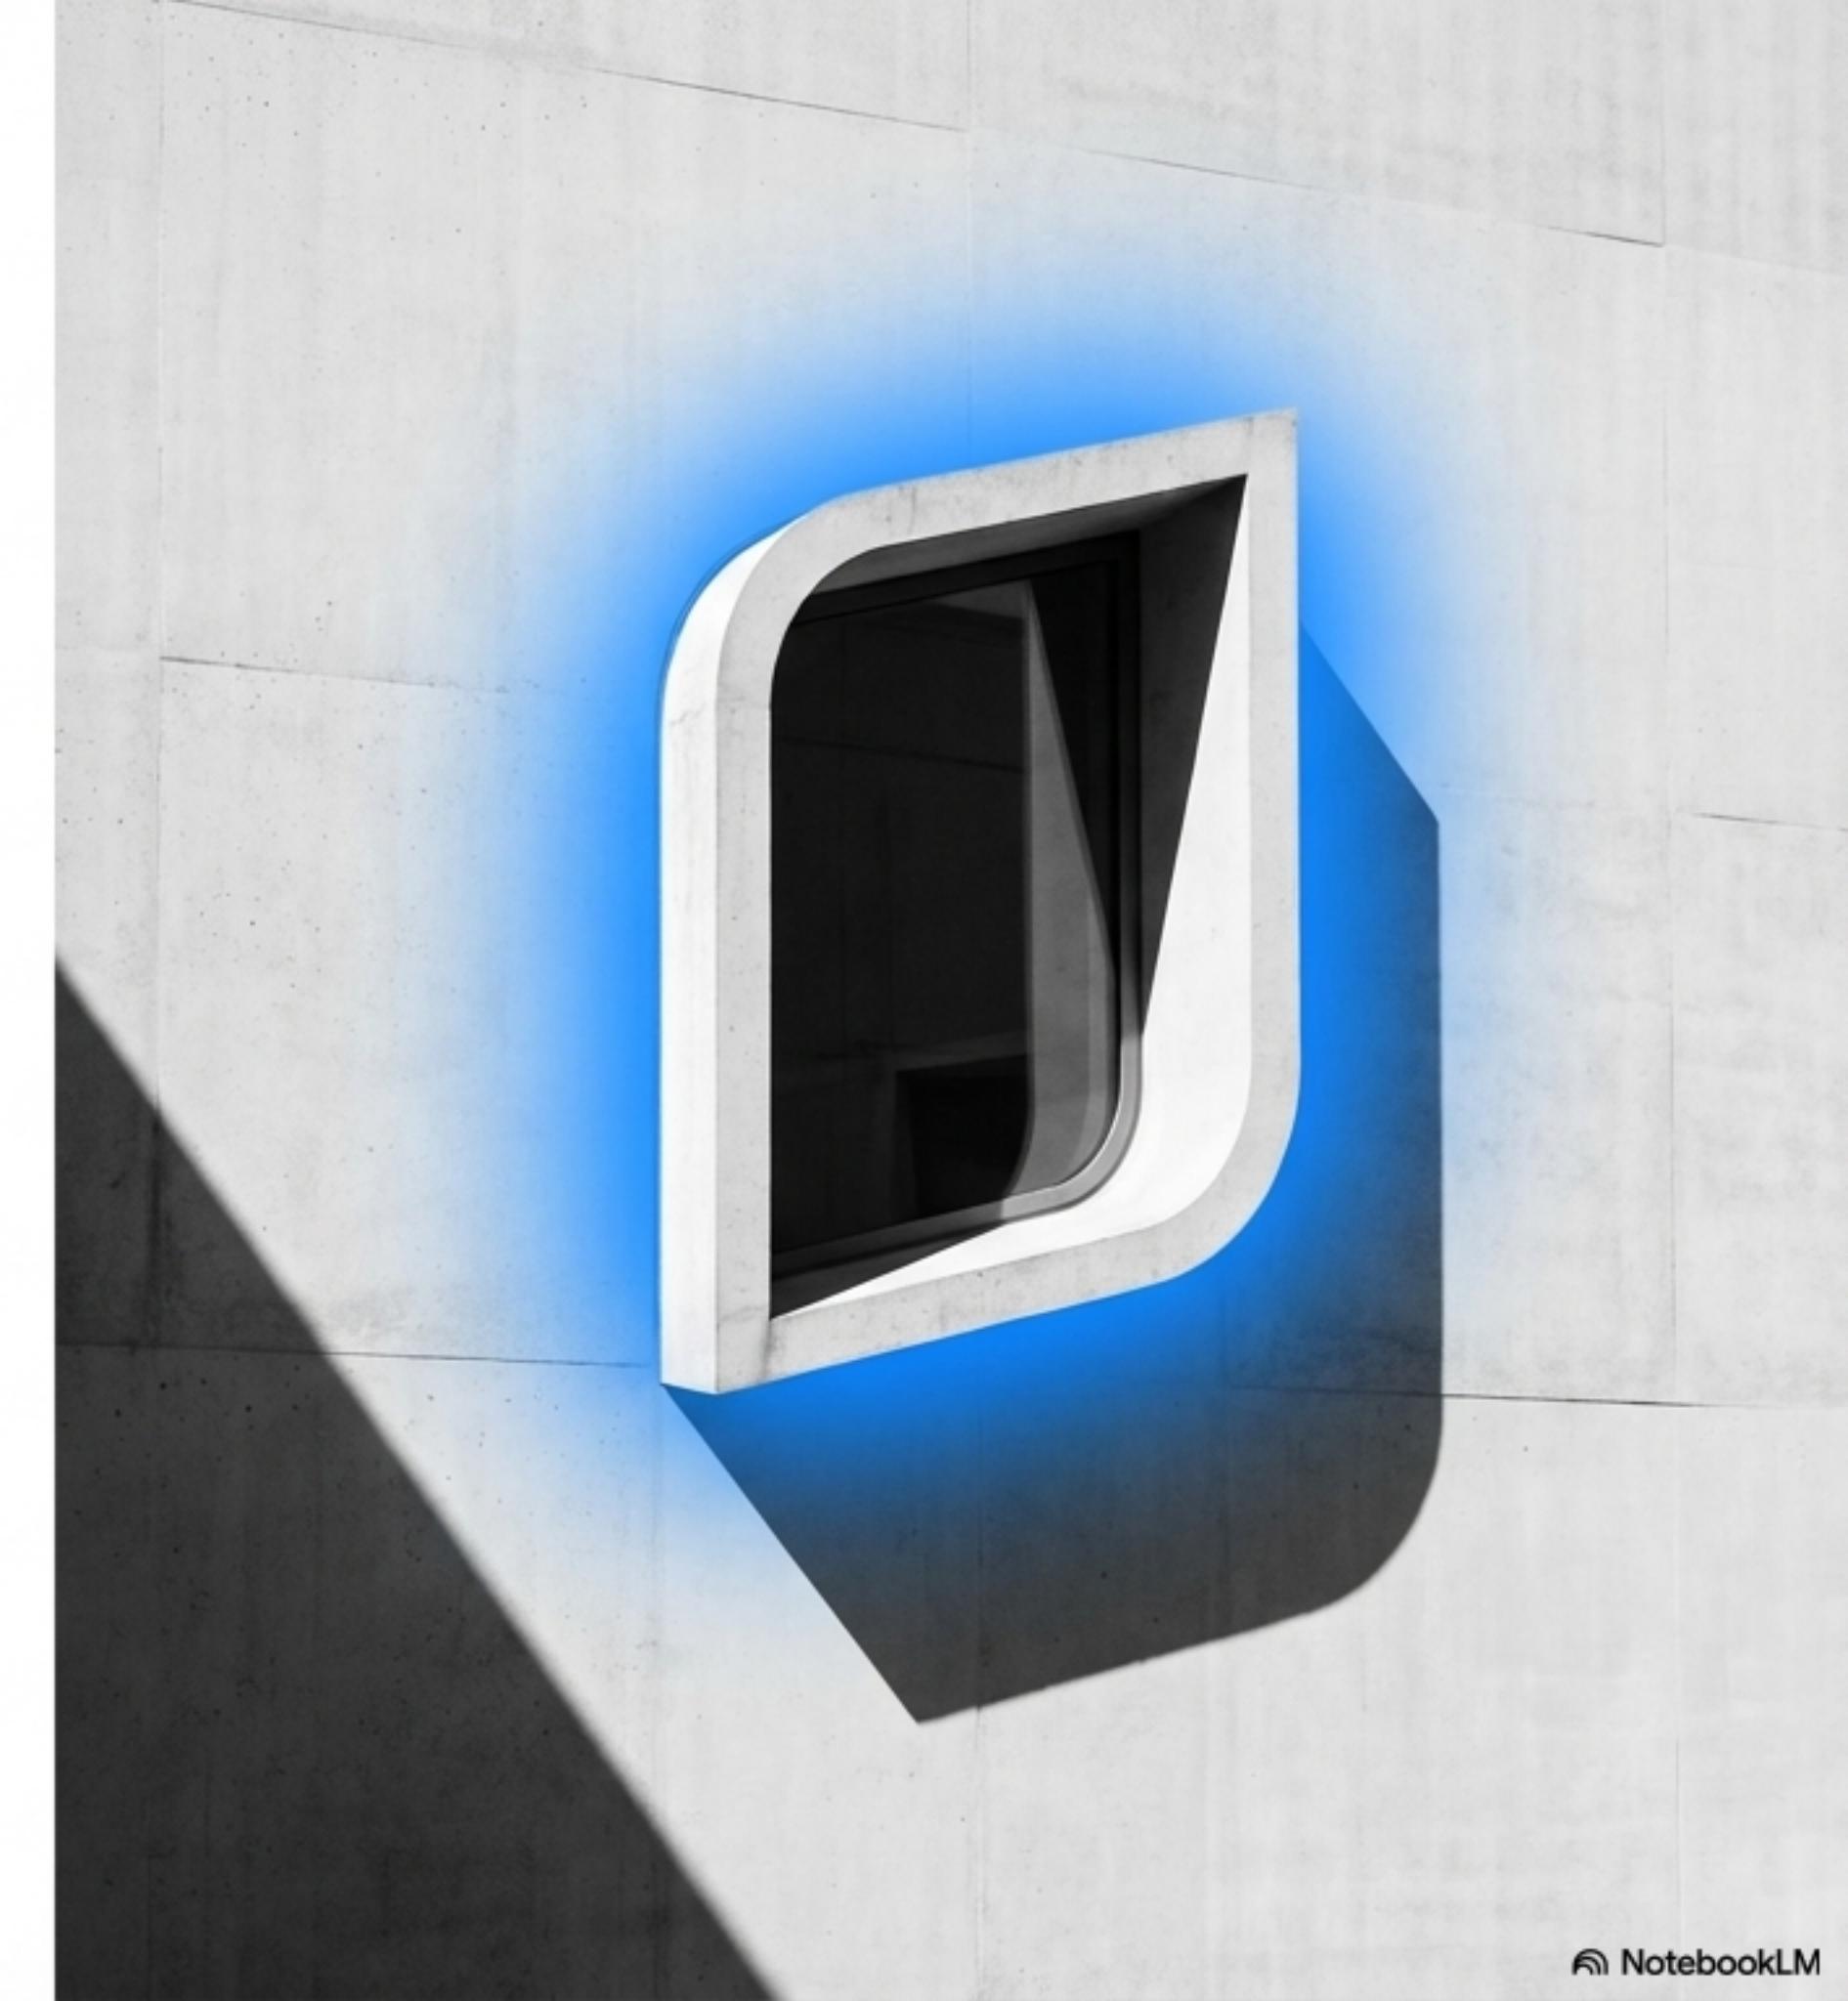
- Role: 'Accountant, Founder, Sales'
- Expertise Level: 'Novice, Intermediate, Expert'

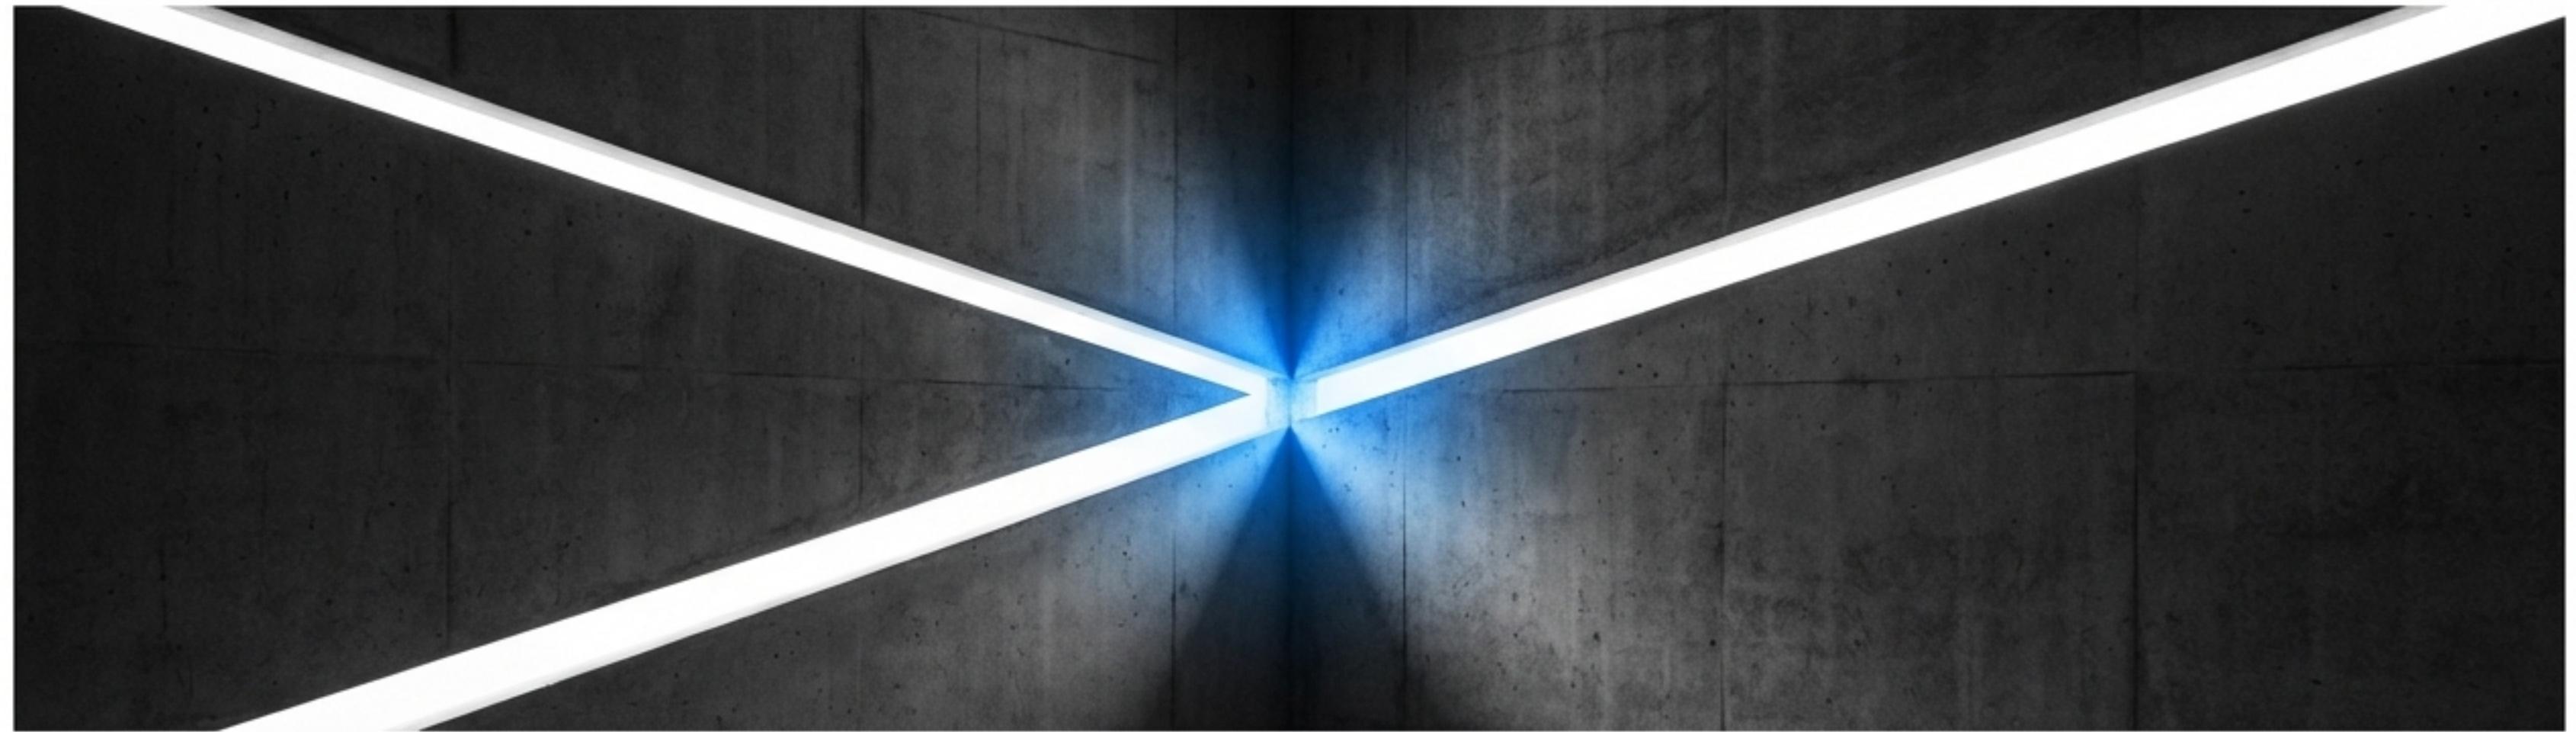
****Behavior****

- Usage Context: 'Daily, Training, Demo'
- Past Friction: 'Frequent errors, 'Aborted flows'

****Preference****

- Help Preference: 'Simple chat, 'Video, 'Auto-fill'
- Risk Tolerance: 'Low_risk_only, 'Full_trust'





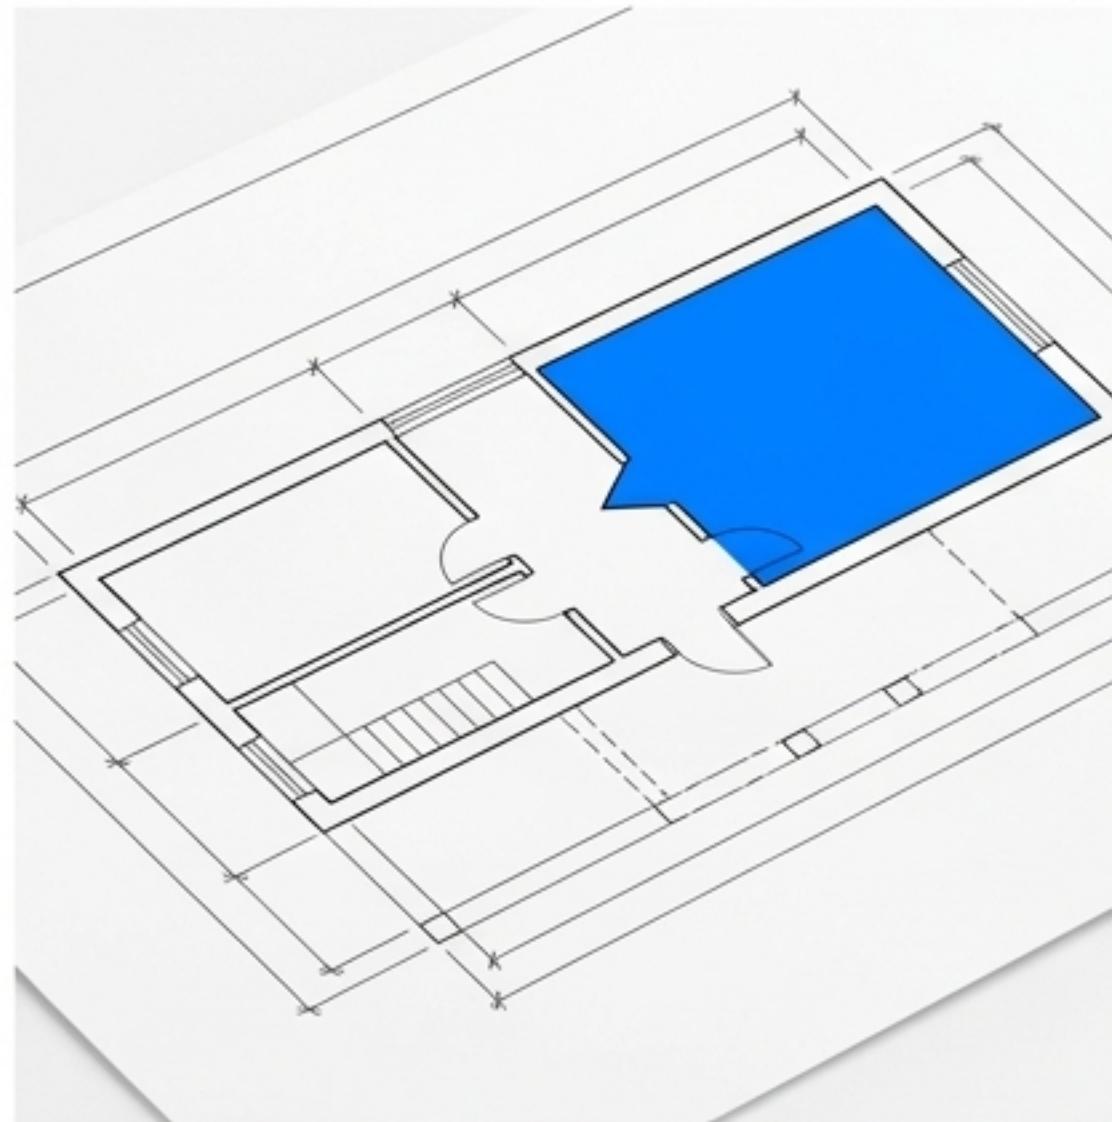
Synthesis: How the AI Achieves Clarity

The AI agent combines these three streams of facets—Screen, Scenario, and User—to create a rich, structured understanding of the precise moment. This context becomes a highly specific query that informs its decision.

Show all `intermediate complexity creation forms` (`Screen``) for a `novice accountant` (`User``) who is trying to `issue their first invoice` during `training` (`Scenario``).

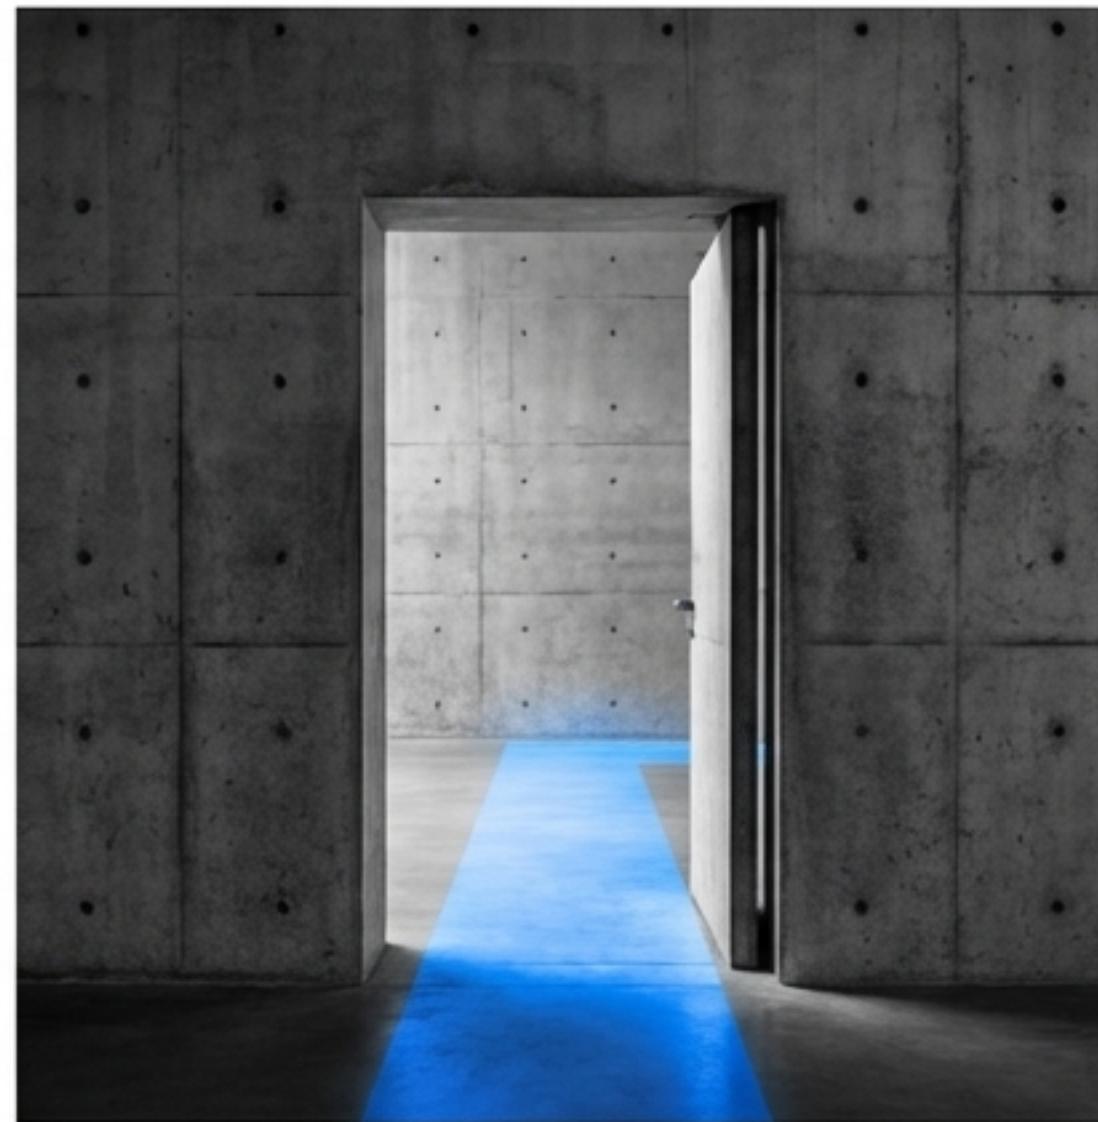
From Insight to Action: The Three Levels of Assistance

This deep context allows the agent to move beyond simple chatbots and offer three distinct levels of help, orchestrated perfectly to the user's needs.



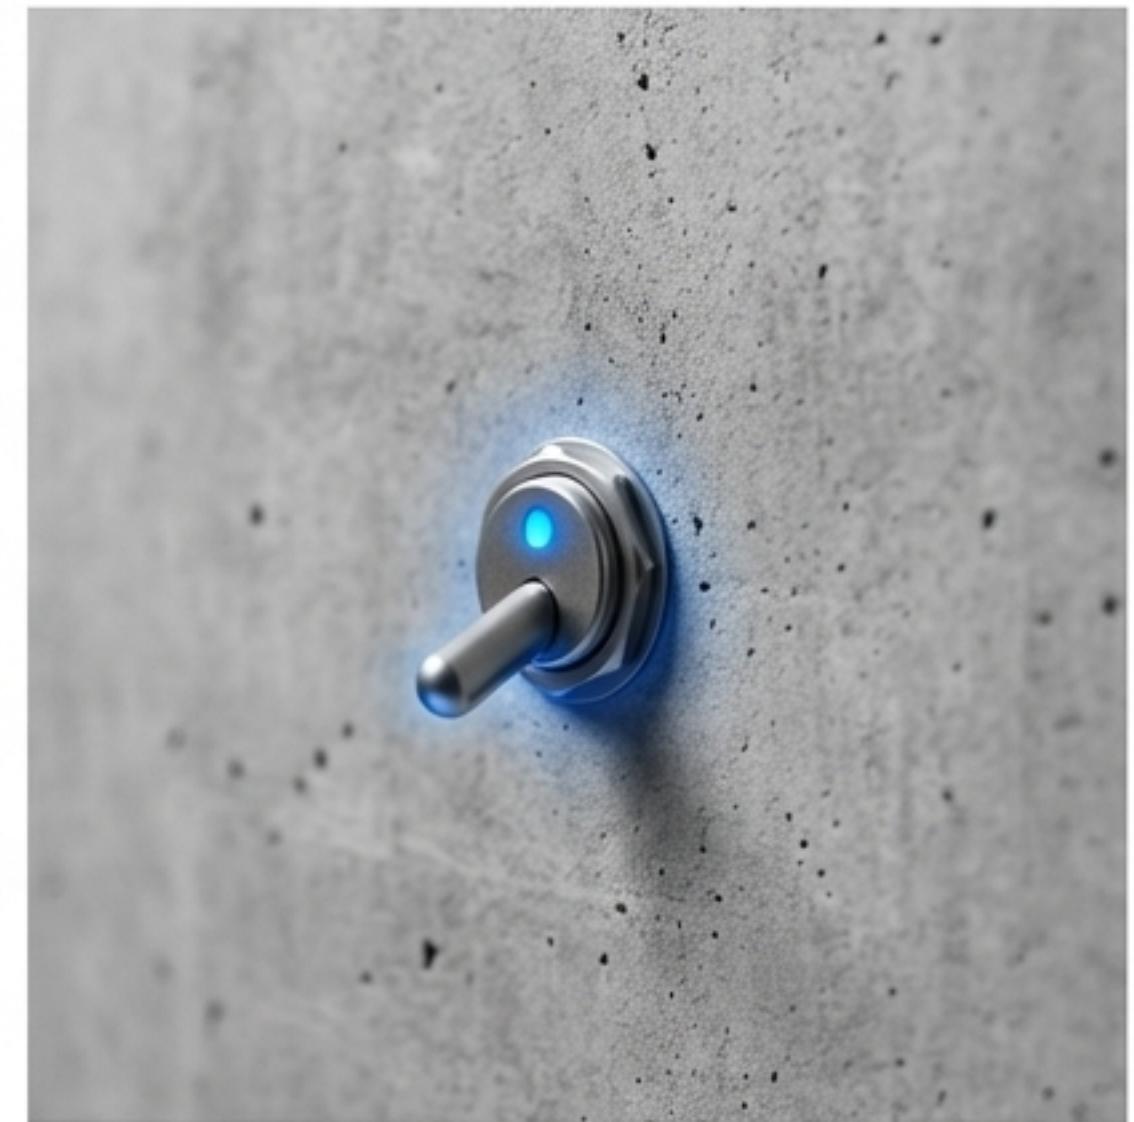
SEE

Provide dynamic visual guidance.



GO

Execute the task on the user's behalf.



DO

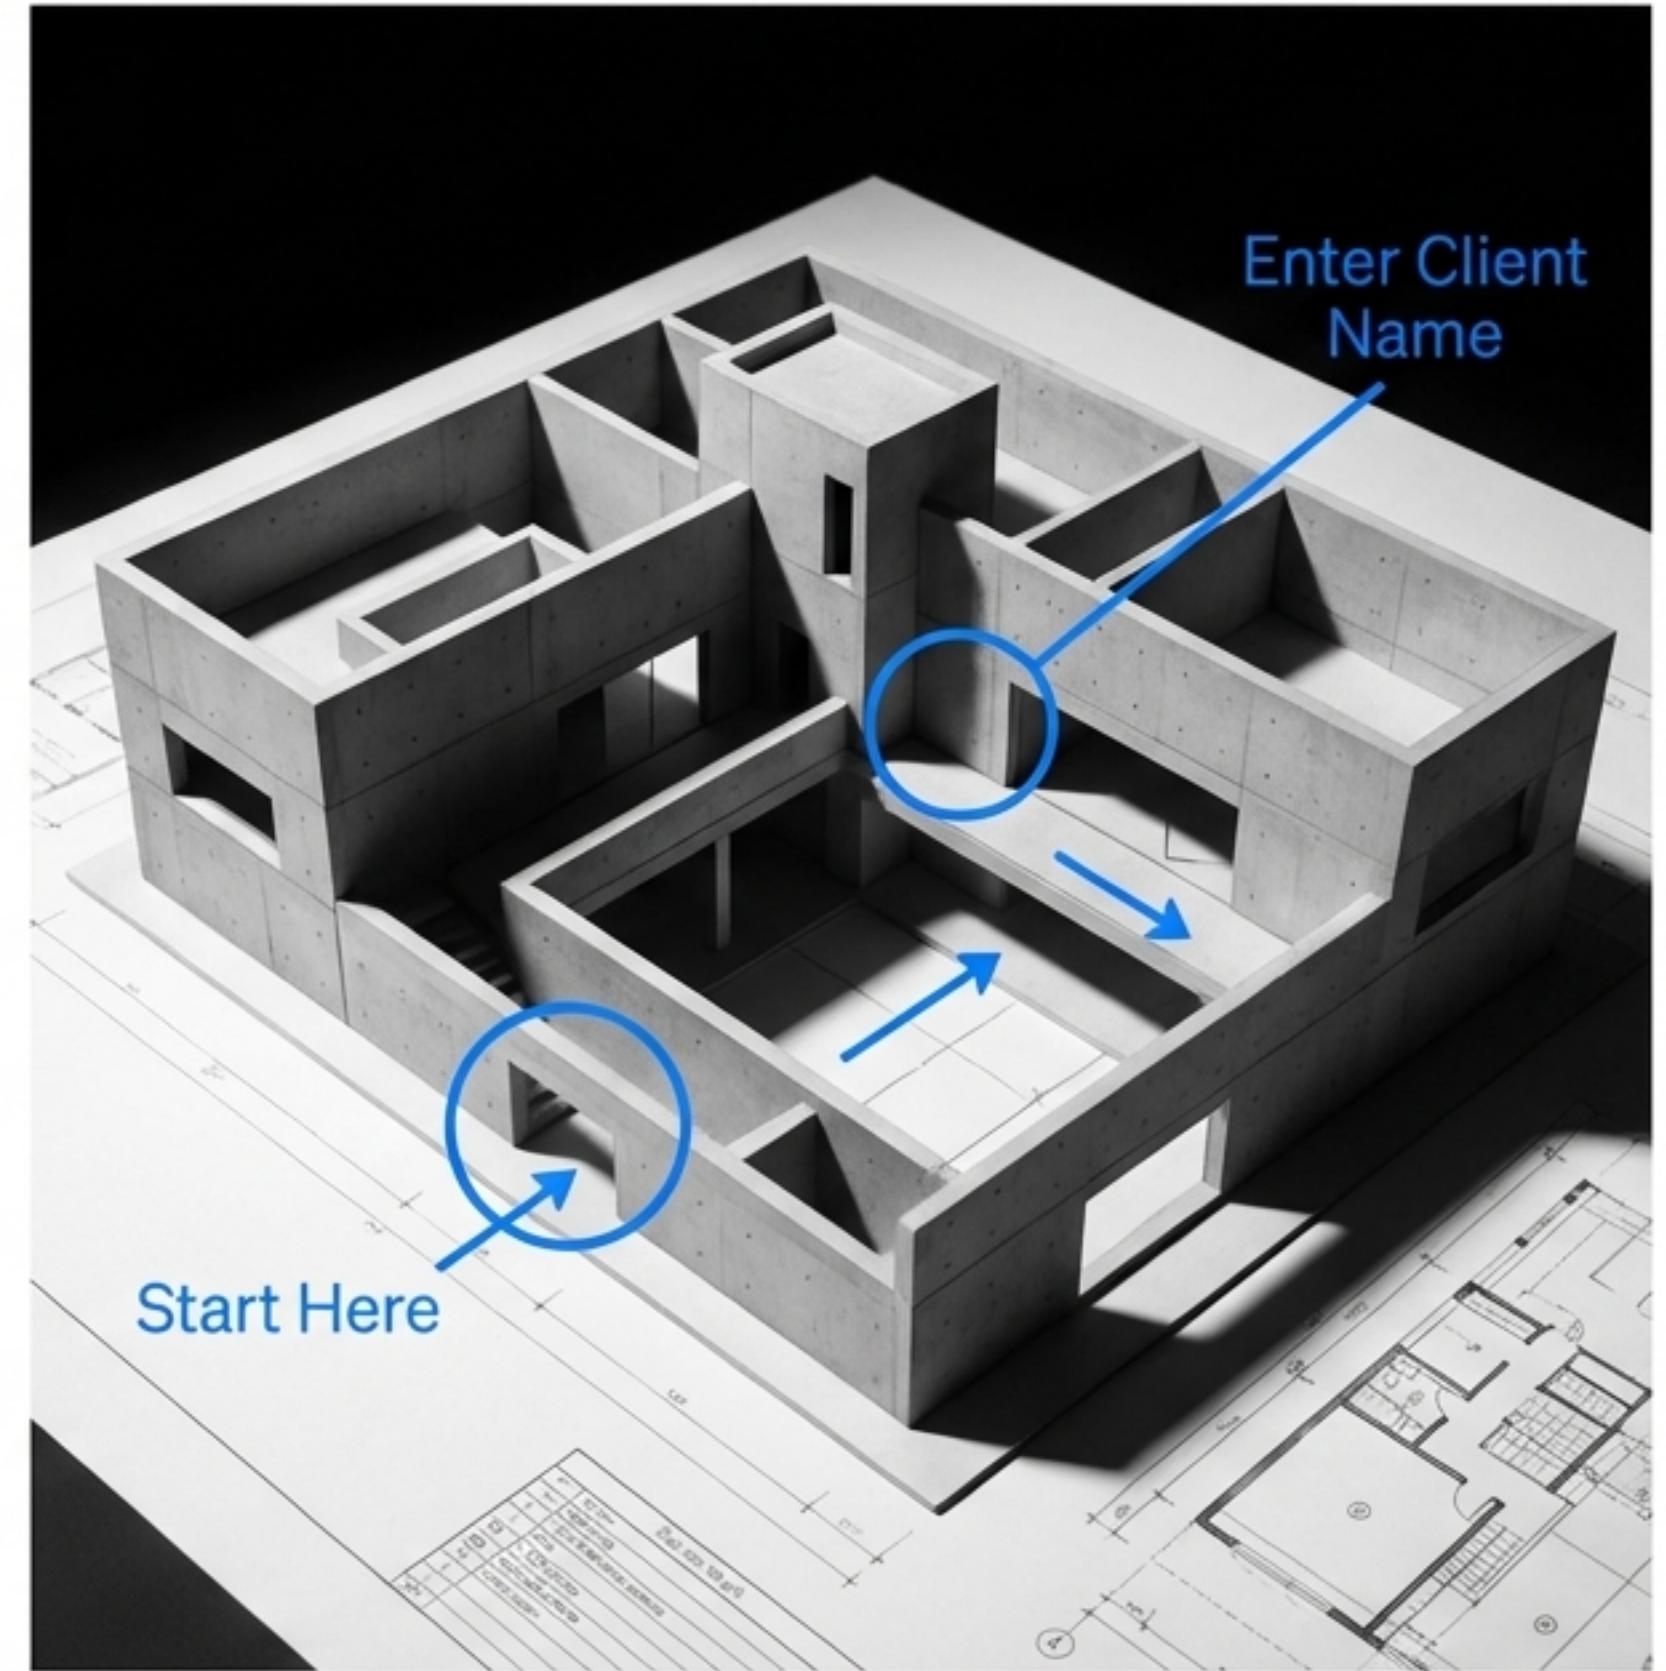
Execute the task on the user's behalf.

Level 1: 'SEE' - Dynamic Guidance for the Novice

Persona: Alice, the new accountant.

Scenario: On a blank screen, she asks, "How do I add my first contact?"

AI Decision: The combination of facets ('persona=novice', 'trigger_context=empty_state', 'help_preference=video') indicates a need for clear, visual instruction. The AI assembles a dynamic 'how-to' guide with annotated screenshots and proposes a short video.



Level 2: 'GO' – Smart Navigation for the Skilled User

Persona: Bob, the intermediate sales manager.

Scenario: He is slowly browsing and appears unable to find a specific deal view he needs.

AI Decision: Facets (`friction_score=high`, `feature_discovered=false`, `time_on_screen > 60s`) suggest the user knows **what** they want but not *where** it is. The AI proactively offers a button: '[Take me to the filtered deal pipeline.](#)'

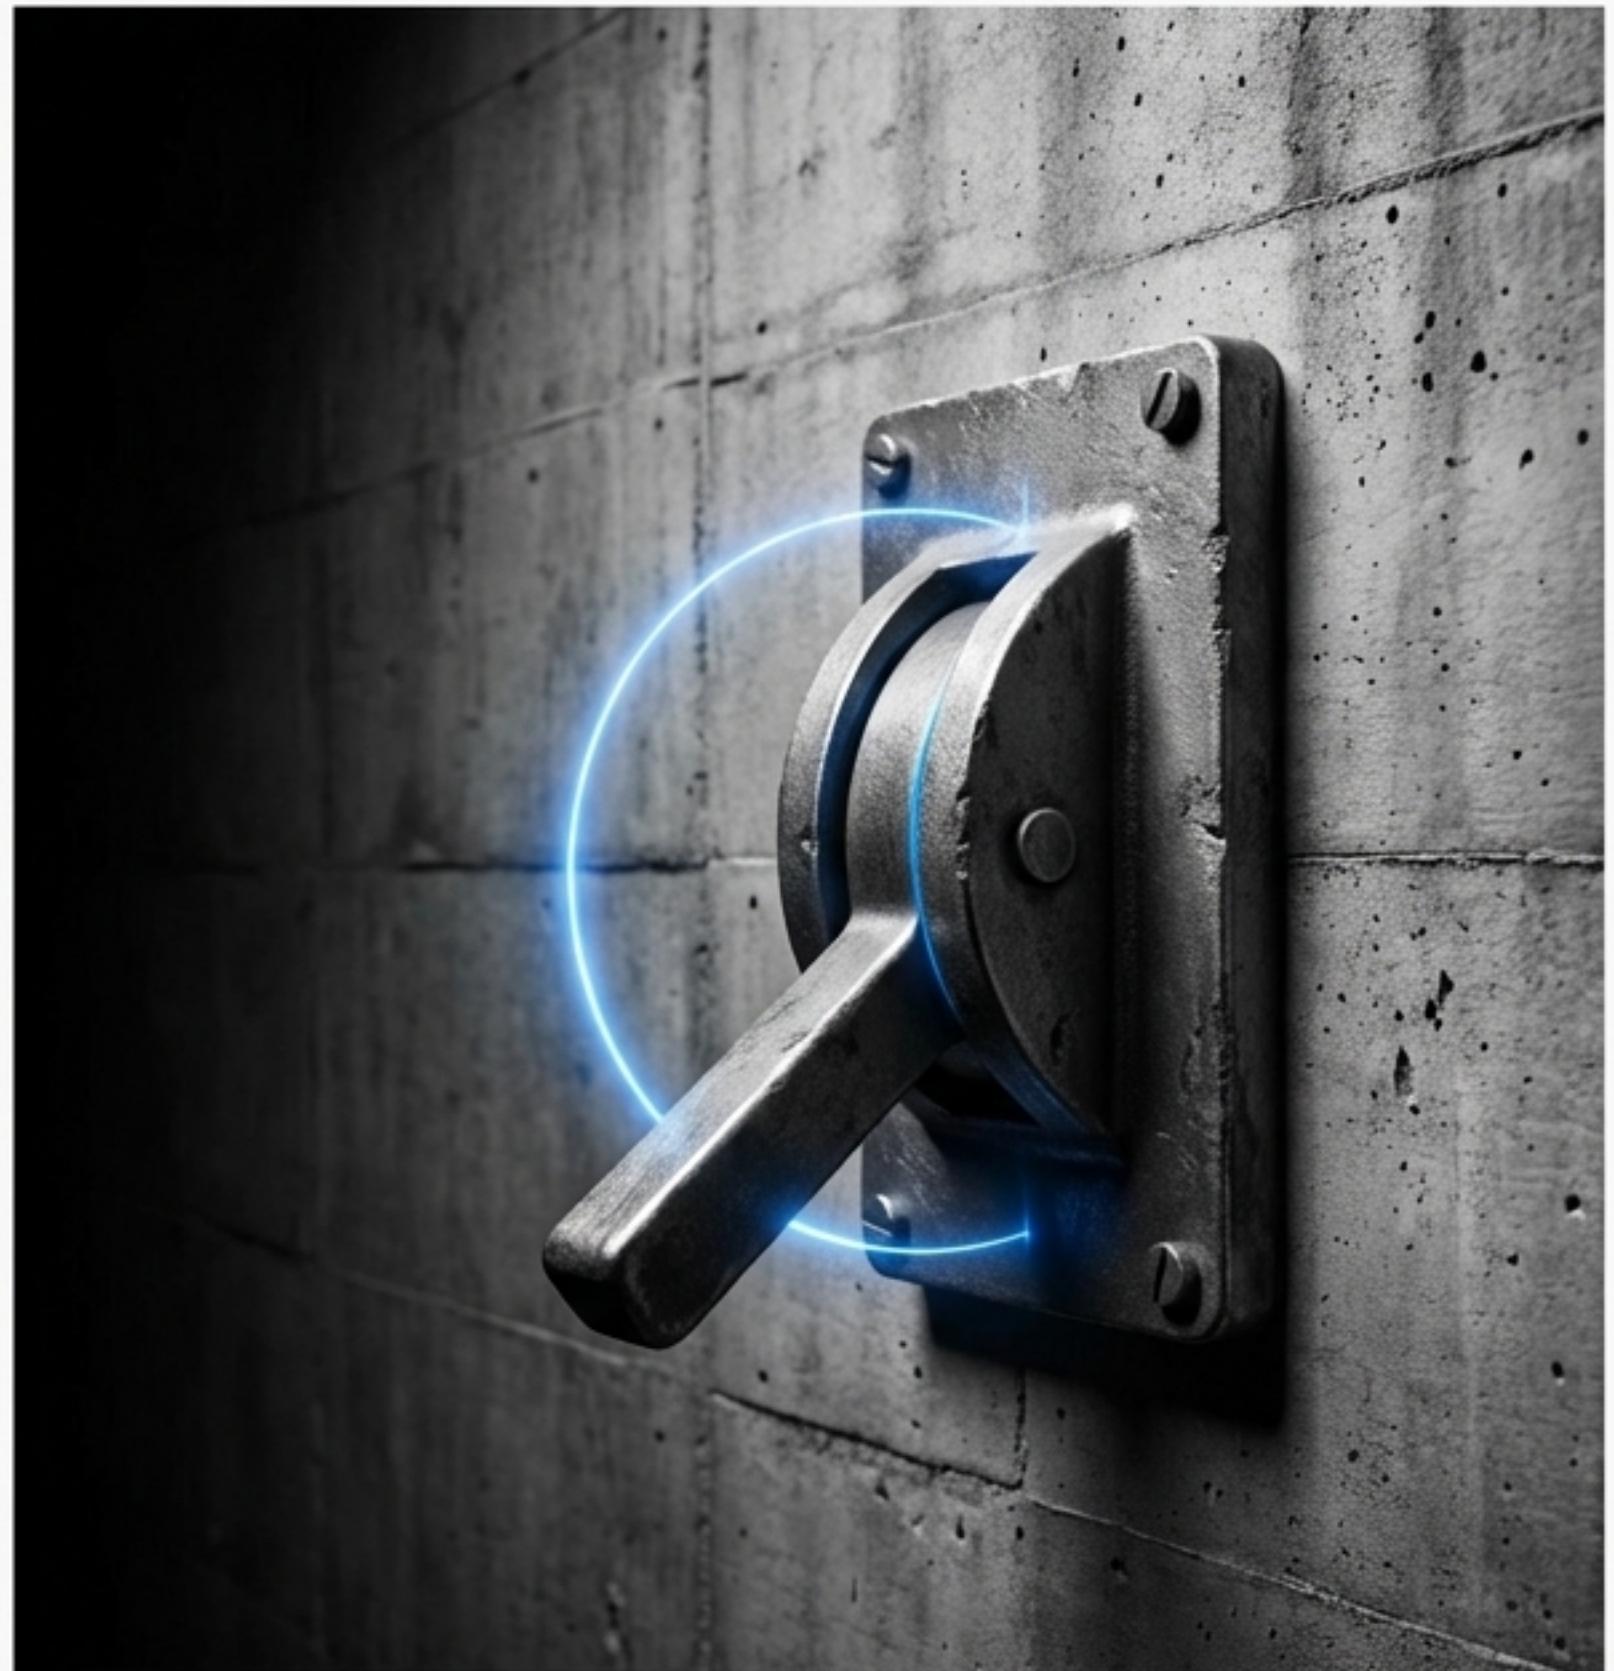


Level 3: 'DO' - Direct Execution for the Expert

Persona: Carol, the expert founder.

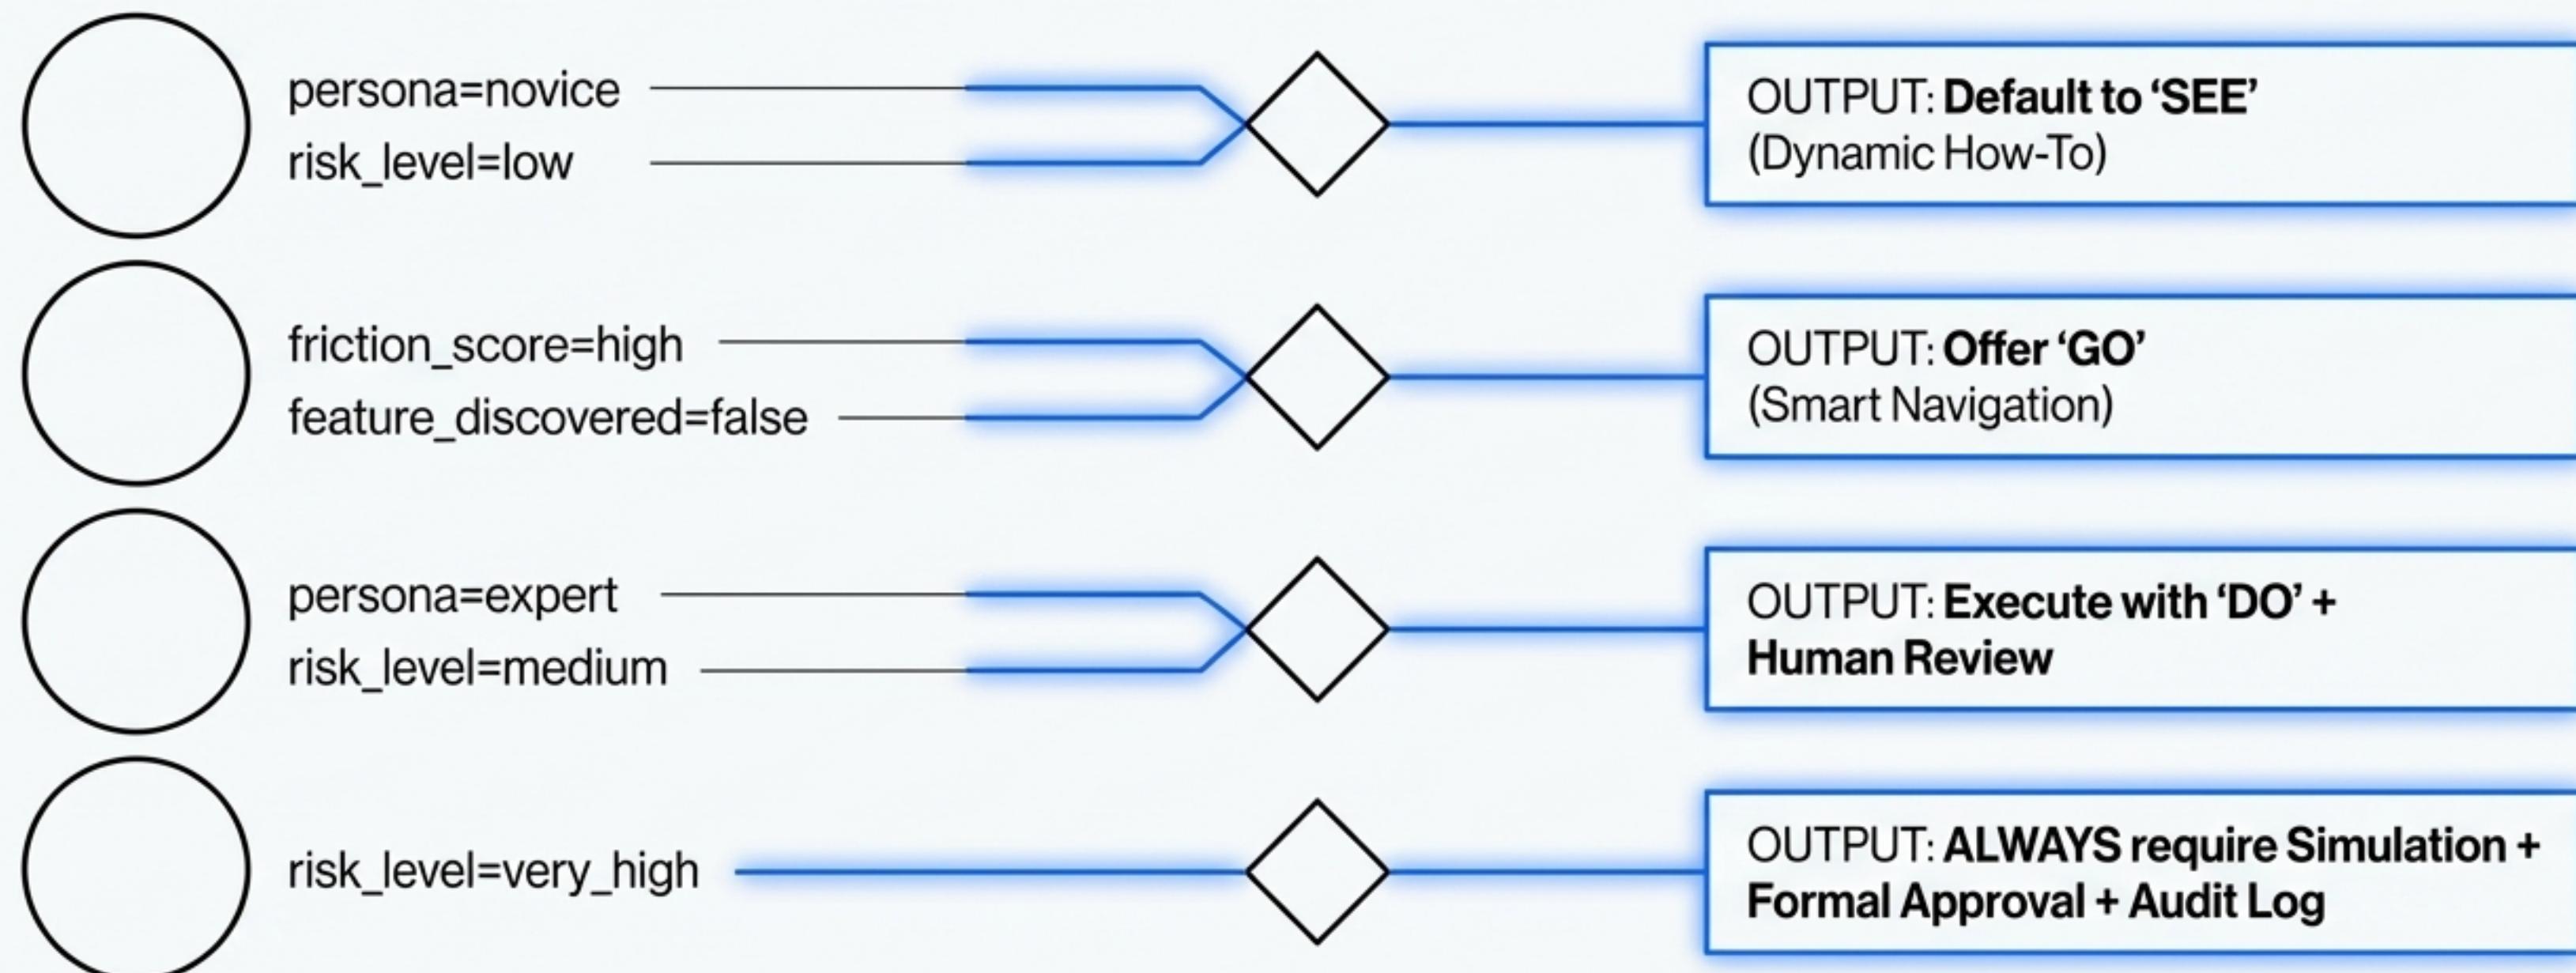
Scenario: She types a direct command: “Create a new Q4 sales pipeline for the enterprise team.”

AI Decision: Facets (`persona='expert'`, `'risk_level='medium'`, `'requires_approval='true'`) indicate a user who values speed but is performing a significant action. The AI performs the action, presents a summary for one-click approval, and logs the audit trail.



The Orchestration Logic

The choice of assistance level is not random. It is determined by a decision model that weighs a combination of facets to match the response to the context and risk.



Evolving the Blueprint: Facets for Action

To enable true 'Do' capabilities, the taxonomy is extended with action-oriented facets. These describe what the agent is allowed to do on each screen, transforming it from a documentation tool into an action engine.

Key Action Facets:

- **can_navigate**: The specific route or URL pattern for the screen.
- **can_autofill**: `true/false` plus DOM selectors or an `api_endpoint` for direct data injection.
- **risk_level**: `low | medium | high | destructive`.
- **required_permissions**: The specific roles or scopes needed to act.
- **webhook_topic**: An event key to trigger the agent proactively.



Structured Empathy

By replacing guesswork with an elegant, **structured taxonomy**, we build an AI that doesn't just respond—it understands and anticipates.

This is not artificial intelligence; it is orchestrated intelligence, designed for mastery and flow.