tenuto labs

Focus Frames

A Second Brain, For Smarter Work

The AI-Powered Wearable for the Knowledge Worker
Bringing enterprise intelligence out of the chatbot and into the physical world of work.

A Concept Prototype from Tenuto Labs

Read time: 8 minutes

The Friction: The Disconnected Knowledge Worker

We started with a question that kept surfacing across our client work: why does every enterprise invest millions in systems of record, only to have their best people make decisions without consulting them?

The data is there. Revenue figures in Salesforce. Cost structures in SAP. Customer escalations in ServiceNow. Workforce allocation in Workday. The enterprise has more business intelligence infrastructure than at any point in history.

And yet.

A VP of Sales sits in a quarterly pipeline review and hears a regional director claim coverage is 2.8x. She has no way to verify that in the moment. The actual number is in Salesforce, but nobody is going to pause the meeting, open a laptop, navigate three screens, and check. So the claim stands. The meeting moves on. Decisions get made on assertions, institutional memory, and whoever speaks with the most confidence.

A CFO walks out of a board prep session and realizes mid-hallway that a number they just agreed to present might not account for new lease obligations. The data is in SAP. But the moment has passed. She sends an email that takes three hours to resolve, or she lets it go and hopes someone catches it later.

A product manager has a hallway conversation where an engineer casually mentions that an integration partner changed their API timeline. That is critical information for a launch decision. But there is no system that captures hallway conversations, and by the time the PM is back at her desk, three other fires have consumed her attention.

The problem is not information access. Every enterprise has dashboards, BI tools, and now AI assistants embedded in their productivity suites. The problem is contextual information access: the right data, surfaced at the right moment, in the flow of how knowledge workers actually work.

And knowledge workers work in motion. In meetings. In hallways. Walking between conference rooms. On calls. Away from their screens.

We set out to design what it might look like to close that gap.

The Concept: Focus Frames

We prototyped a software and AI layer purpose-built for smart glasses hardware, designed exclusively for the enterprise knowledge worker. Not a consumer product adapted for business. A ground-up design for the way corporate decision-makers actually spend their days. We call it Focus Frames: enterprise glasses that function as a second brain for the knowledge worker.

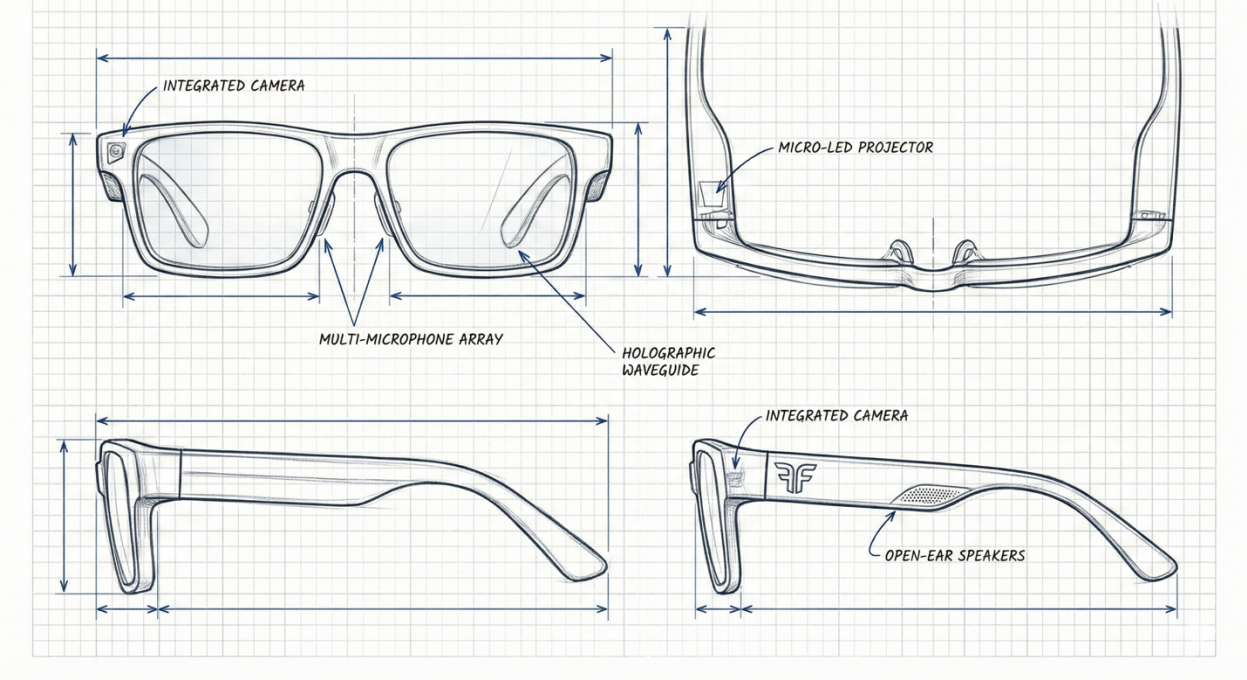

The Core Idea

The worker wears a pair of smart glasses that look like ordinary eyewear. During the workday, the glasses do three things that no existing tool does together.

Ambient audio intelligence. The glasses capture real-time audio during meetings, hallway conversations, whiteboard sessions, and informal interactions. This audio is processed by

generative AI to produce structured intelligence: summaries, action items, decision logs, and data verification flags.

Enterprise data on demand. The worker can query enterprise systems of record using natural language voice commands. "What was our customer acquisition cost in the Midwest last quarter?" The AI translates the question into the appropriate API call, retrieves the data, synthesizes an answer, and delivers it via voice response or pushes a formatted analysis to email.

Proactive contextual intelligence. In its most advanced mode, the system listens to the ambient conversation, identifies topics where enterprise data could add value, and surfaces relevant information without being asked. A silent analyst sitting in every meeting, cross-referencing everything said against everything the company knows.

Think of Focus Frames as carrying a brilliant, tireless chief of staff who has read every report, remembers every number, and has real-time access to every system in the enterprise. A second brain, for smarter work.

Why Now: The Generative AI Difference

Enterprise wearables have failed before. Google Glass Enterprise. RealWear. Vuzix. The common explanation is form factor and social acceptability. But the deeper reason is simpler: the use cases were not valuable enough.

"Look up the weather" or "scan this barcode" did not justify wearing a computer on your face. The information density was too low and the interaction model was too rigid. Workers were getting a small screen that did what their phone already did, less conveniently.

Generative AI changes this calculus in three ways that matter.

Natural language as the interface. Workers no longer navigate apps, screens, or structured queries. They speak naturally. The AI handles the translation from human question to system query. This means the glasses inherit all the knowledge embedded in enterprise APIs without the worker needing to know those systems exist.

Synthesis, not retrieval. Previous wearables could only retrieve pre-formatted data. Generative AI synthesizes information across multiple systems, applies reasoning, and delivers an insight rather than a data point. Not "your CAC is $142." Instead: "Your CAC in the Midwest was $142, 18% above plan, primarily driven by the underperformance of the digital campaigns you discussed in the Q3 review."

Ambient understanding. Large language models can process continuous audio, understand context, identify topics, and determine when data would be relevant. This transforms the glasses from a voice-activated search engine into an intelligent work companion. The system does not wait to be asked. It recognizes when it has something worth saying.

These three capabilities together create a product that previous hardware could never deliver. The form factor has not changed. The intelligence behind it has.

What It Looks Like: A Day with Focus Frames

To make the concept tangible, we designed a full day of interactions for a single persona. Sarah Chen is a VP of Commercial Operations at a mid-size enterprise. She manages regional sales directors, owns the customer success P&L, and reports to the COO. She is the kind of executive who remembers every number from last quarter but cannot recall where she parked her car.

She puts on her Focus Frames the way she puts on her watch. It is part of her daily toolkit.

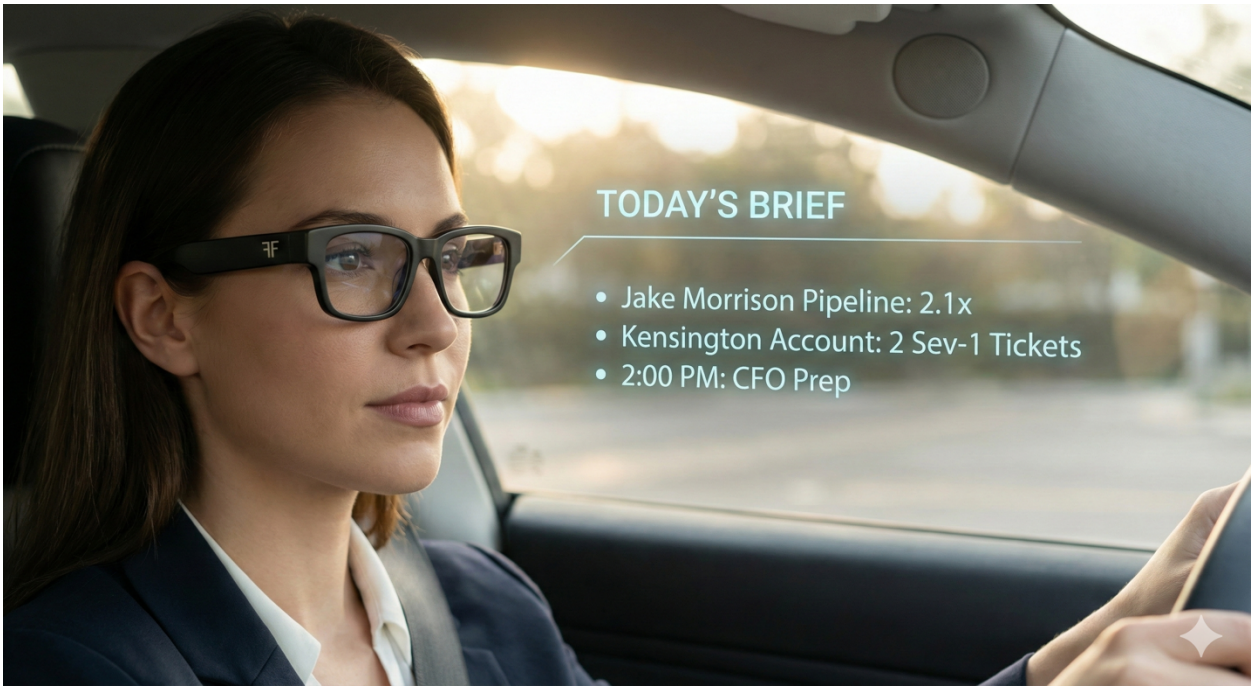

7:45 AM: The Commute Brief

On the drive in, Sarah says: "Brief me on today."

The glasses pull her calendar, cross-reference attendees against recent CRM activity and email threads, and deliver a 90-second audio summary. She learns that Jake Morrison's

Southeast pipeline coverage is at 2.1x, below her 2.5x threshold. She learns that the Kensington account has two severity-one tickets open in ServiceNow. She learns that the CFO prep session at 2 PM is confirmed and offers to send a pre-read to her email with the latest P&L data.

She says: "Send the pre-read. And flag Jake's pipeline number for me during the 9 AM."

Before she parks the car, she knows more about her day than she used to know after thirty minutes at her desk.

9:00 AM: The Pipeline Review

Eight people in a conference room. Regional directors presenting their numbers. Jake says: "Pipeline coverage is strong at 2.8x. We are in great shape for Q2."

Sarah hears a subtle tone in her ear, followed by a low-volume voice: "Jake stated pipeline coverage at 2.8x. Your Salesforce data shows 2.1x as of this morning. The discrepancy may be due to three deals he moved to Commit stage yesterday that have not been validated by sales operations."

Sarah now has a choice. She can address the discrepancy or follow up privately. Either way, she has the data. Without the glasses, that claim goes unchallenged.

She says, calmly: "Jake, can we double-click on that 2.8 number? I want to make sure we are counting validated deals only."

This is the moment we designed the entire product around. Not catching someone being wrong. Having the data to ask the right question at the right time. In the room. Not three hours later in an email thread.

10:30 AM: The Hallway Moment

Walking to her next meeting, Sarah runs into Marcus from Product. He mentions that an integration partner pushed their API delivery by three weeks. That affects the customer advisory board demo in six weeks.

Sarah touches the temple of her glasses and says: "Capture that."

The glasses save the last 90 seconds of conversation, generate a structured summary, create a draft task, and offer to notify the CAB planning team. Sarah confirms. A casual hallway conversation is now a tracked action item in the system of record. No email written. No note taken. The glasses captured it.

Half of the most important information in any company travels through hallways, not email. The problem was never that people did not share. The problem was that nothing stuck.

11:00 AM: The Escalation Meeting

The Kensington account review. The account manager says: "The client is happy. No major issues on our end."

Sarah says: "What about the two severity-one tickets that are still open? One has been open for eleven days."

The room goes quiet. The account manager did not know Sarah had that data. She did not pull up a laptop or check her phone. She simply knew.

This scene illustrates something important about the product design. The AI is completely invisible to everyone else in the room. Sarah appears to be an extraordinarily well-prepared executive. The technology is ambient. The human impact is what people experience.

12:30 PM: The Working Lunch

Over a sandwich, Sarah asks: "Give me a comparison of our top ten accounts by revenue growth versus support ticket volume over the last two quarters."

The AI queries Salesforce for revenue data, ServiceNow for ticket data, runs a correlation, and sends a formatted analysis to her email with a chart and summary. Fifteen seconds. She reviews it on her phone, forwards it to her team with a note.

That analysis used to take half a day. Pull data from two systems, clean it, build a spreadsheet, make a chart, email it around. She asked for it over lunch and had it before she finished her sandwich.

2:00 PM: Board Prep with the CFO

The CFO asks: "What is our net revenue retention, and how does it compare to the SaaS benchmarks we referenced last quarter?"

Sarah says quietly: "What is our NRR from the Q4 close data, and pull the SaaS NRR benchmarks from the KeyBanc report we reviewed in December."

Within seconds: "Your NRR is 112%. The median SaaS benchmark from the KeyBanc report was 110%. You are above median but below top quartile at 118%."

Sarah turns to the CFO: "We are at 112. Benchmark median is 110, top quartile is 118. We are above average but have room to grow."

Nobody left the room. Nobody opened a laptop. Nobody sent a Slack message to an analyst and waited forty-five minutes. The answer was there, in the conversation, when it mattered.

4:30 PM: End of Day

Sarah asks for her end-of-day summary. The glasses compile everything: six meetings, 4.5 hours, key decisions made, open items requiring action by specific dates, and two items that were discussed but never formally assigned.

She confirms the creation of draft tasks for the unassigned items.

Her day's institutional knowledge is captured, structured, and actionable. Nothing fell through the cracks. Nothing relies on her memory alone.

The Design Philosophy: Augment, Don't Interrupt

The hardest design constraint was not technical. It was behavioral.

A knowledge worker in a meeting is already managing multiple cognitive streams: listening, evaluating, preparing their next point, reading the room. Adding information to that context is only valuable if it reduces cognitive load rather than adding to it.

We designed around three principles.

Voice-first, screen-second. The primary interaction is voice in, voice out. The worker speaks a question, the AI responds in their ear. For complex outputs like charts, multi-system analyses, or anything that benefits from visual review, the system pushes to email or Slack. We do not try to cram a dashboard into a heads-up display. The right output channel depends on the complexity of the answer.

On-demand before ambient. The most immediately buildable and trustworthy interactions are explicit queries: the worker asks, the system answers. Ambient proactive intelligence, where the system listens and surfaces data without being asked, is higher value but requires more sophistication. We designed the product in horizons, with on-demand capabilities first and ambient intelligence as the system learns the worker's patterns and context.

The handoff is the product. The value is not in the AI output itself. It is in how the AI output changes the human workflow. The pre-meeting brief changes how Sarah walks into a room. The post-meeting synthesis changes how her team follows up. The hallway capture changes

whether informal knowledge enters the system of record. We designed around workflow transformation, not feature demonstration.

How It Works: The Technical Architecture

The system is designed in four layers, each with clear boundaries.

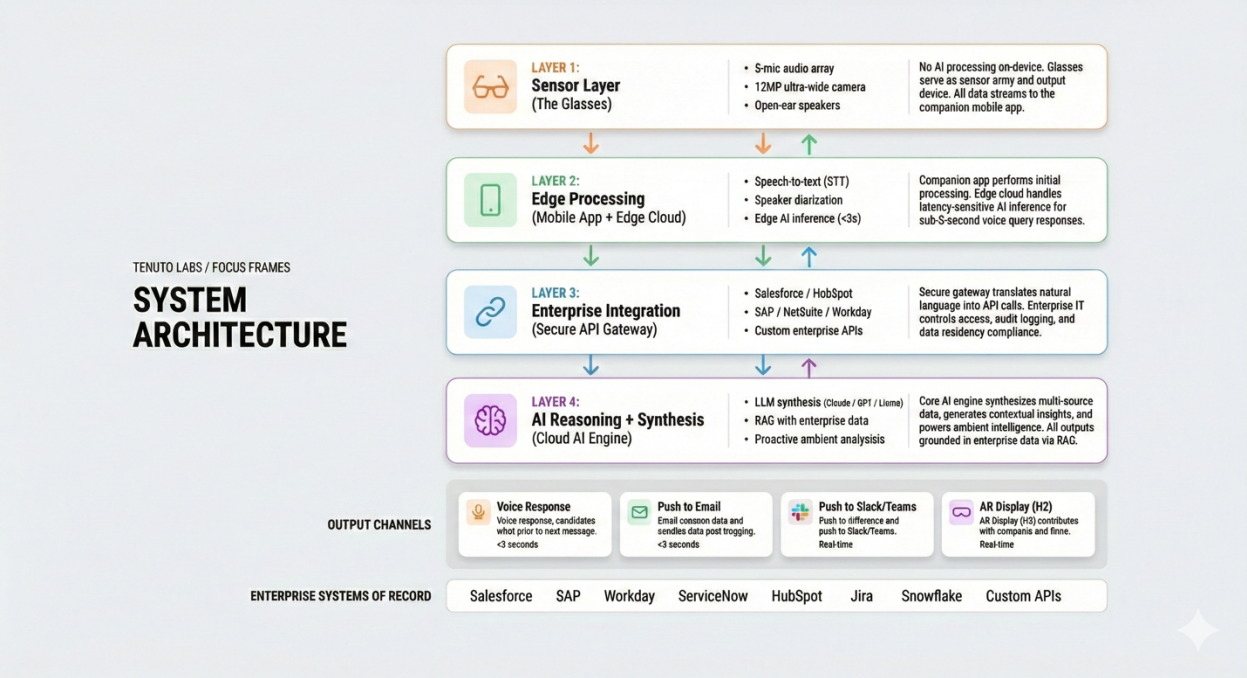

Layer 1: Sensor Layer (The Glasses). The glasses are a sensor array and output device. They capture audio via multi-microphone array and visual data via integrated camera. They deliver output through open-ear speakers. No AI processing happens on the glasses themselves. All sensor data streams to the companion mobile application via the Meta Wearables Device Access Toolkit.

Layer 2: Edge Processing (Mobile App + Edge Cloud). The companion app receives raw sensor data and performs initial processing: speech-to-text, speaker diarization, and intent classification. For latency-sensitive operations, edge cloud infrastructure handles AI inference closer to the user, targeting sub-three-second response times for voice queries.

Layer 3: Enterprise Integration (API Gateway). A secure gateway translates natural language queries into structured API calls to Salesforce, SAP, Workday, ServiceNow, and other platforms. This layer is where enterprise IT maintains control: which data sources are accessible, which users have access, what audit logging is required.

Layer 4: AI Reasoning and Synthesis (Cloud). The core AI engine synthesizes data from multiple sources, generates contextual insights, produces meeting intelligence, and powers proactive ambient analysis. This layer uses large language models with retrieval-augmented generation (RAG) to ground all outputs in actual enterprise data.

Output Channels

The system delivers intelligence through the channel that matches the use case. Voice response through the glasses for quick queries and ambient alerts. Push-to-email for complex analyses and visualizations. Push-to-Slack or Teams for collaborative items. And in the future, AR in-lens display for glanceable data overlays on hardware that supports it.

The Platform: What Exists Today

We are not building hardware. We are designing the AI and software layer, Focus Frames, that makes enterprise glasses useful for the knowledge worker.

Meta launched the Wearables Device Access Toolkit in developer preview in late 2025. The SDK supports Ray-Ban Meta (Gen 1 and Gen 2) and provides access to the camera, microphones, and speakers. Developers build companion apps on iOS or Android that process sensor data and return output to the glasses.

This gives us enough to prototype the core experience today: voice-activated data queries, pre-meeting and post-meeting intelligence, hallway capture, and end-of-day synthesis. Audio in, AI processing, voice and email out.

What the SDK does not yet support is sending visual content to the in-lens display on Meta's newer Ray-Ban Display hardware. That capability will likely open to developers over time. But our near-term design intentionally does not depend on it. Voice and email are powerful enough output channels to demonstrate the full value proposition.

We designed the product in two horizons.

Horizon 1: Buildable now. Voice-in, voice-out. Push-to-email. Pre-meeting briefs, on-demand data queries, conversation capture, post-meeting synthesis, end-of-day digests, ambient CRM updating, cross-system analysis, whiteboard capture, communication coaching, and customer visit intelligence. Nine use cases on existing hardware and SDK.

Horizon 2: Future capabilities. Proactive ambient intelligence (the system listens and surfaces data without being asked), AR data overlays in the lens, live translation with enterprise terminology, negotiation support with real-time data surfacing, and onboarding

acceleration for new employees. These require either display API access from Meta or more advanced ambient processing pipelines.

The SaaS Platform Opportunity

One of the more interesting angles we explored is not selling to enterprises directly, but partnering with the SaaS platforms whose data the glasses surface.

Consider the incentive structure from Salesforce's perspective. Their customers pay significant license fees but frequently underutilize the platform. Sales reps avoid updating CRM records. Executives rely on exported spreadsheets rather than live dashboards. The data in Salesforce is only as valuable as the frequency and context in which it is accessed.

Focus Frames represents a new access pattern for Salesforce data: continuous, contextual, ambient. If a VP of Sales can query Salesforce by speaking a question during a meeting, the platform becomes dramatically more valuable. If CRM records update themselves based on captured conversations, adoption friction drops to near zero. Salesforce becomes more sticky, more utilized, more essential.

The same logic applies to SAP, ServiceNow, Workday, and every other enterprise SaaS platform. The pitch is straightforward: we are not a competitor. We are a new modality for your platform. We bring the user experience. You bring the data. Together we create a new category of enterprise interaction.

What We Are Still Thinking About

This is a concept prototype. We have not built production software. We have not conducted enterprise pilots. What we have done is demonstrate a design direction and show that the technical components exist today. Honesty about the open questions matters as much as the vision.

The trust problem in ambient listening. If three people in a meeting are wearing glasses that capture audio and cross-reference it against enterprise data, the meeting dynamics may change. People may become guarded. The value of spontaneous, informal knowledge exchange could degrade precisely because it is being captured. This is not a technical problem. It is a cultural design problem, and it requires thoughtful organizational change management, not just good software.

Latency versus real-time. The proactive ambient intelligence scenario requires a multi-step pipeline: audio capture, speech-to-text, intent detection, query formulation, API call, data

retrieval, AI synthesis, and delivery. Done well, that takes 10 to 30 seconds. In a live meeting, the conversation has moved on. Our design accommodates this by delivering ambient insights as "after the fact" nudges rather than pretending they arrive instantaneously.

Enterprise security. Getting enterprises to authorize a wearable device to access production systems of record is a multi-month IT security review. Audio data flowing to cloud AI services must comply with data residency requirements. SOC 2, ISO 27001, and similar compliance frameworks will be expected. This is solvable, but it requires building the compliance architecture from day one, not as an afterthought.

The Microsoft question. Microsoft has Copilot, enterprise data integrations through Microsoft Graph, Teams, and an early partnership with Meta's SDK. If they ship a "Copilot for smart glasses" that covers 60% of what we describe, the differentiation question becomes urgent. We believe the opportunity is in vertical depth for specific personas, not horizontal breadth, but this risk deserves honest acknowledgment.

Why We Built This

Tenuto Labs is an Applied AI studio. Most of our work is software: enterprise applications, intelligent workflows, decision support systems. We built this concept prototype to demonstrate something we believe strongly.

The next interface for enterprise AI is not a chatbot. It is not a sidebar in Microsoft Teams. It is not a dashboard with a query bar. It is ambient intelligence, delivered in the physical world, in the moments where decisions are actually made.

The hardware is shipping. Meta's Ray-Ban glasses have sold millions of units. The SDK is open. Generative AI can synthesize data across enterprise systems in seconds. Edge computing can deliver sub-three-second response times. Every component exists. What has been missing is the design thinking to connect them for the enterprise knowledge worker.

That is what we do. We design applied AI systems that remove friction from critical decision loops. Focus Frames shows what that looks like when the loop extends beyond the screen and into the conference room, the hallway, and the commute.

Tenuto Labs is an Applied AI studio that designs and builds bespoke AI applications. We work with organizations navigating the space between prototype and production, in software and beyond.