

AKTANA

Context-Aware AI for Effective Decision Making

Agentic AI & Continuous Improvement



Context-Aware AI for Effective Decision Making: A View from Derek Choy

“ When our team at Aktana expanded our AI systems from pilots to global deployments, we found that scaling these initiatives was more complex than anticipated. Initial attempts relied heavily on manual configurations—rules and static models—which rapidly became cumbersome to maintain at scale. We lacked a **centralized orchestration mechanism** and a structured knowledge base that could dynamically evaluate constraints, and make value driven decisions. This gap led to inefficiencies and stalled adoption.

These experiences taught us the critical importance of having a centralized, dynamic knowledge base at the heart of our AI architecture. Today, Aktana’s multi-agent system leverages this knowledge base to **continuously refine its recommendations, rapidly learning what strategies succeed or fail**. By capturing real-world signals, such as engagement metrics and KPI attainment, it enables proactive adjustments rather than reactive corrections.

The Aktana Strategy Console drives these feedback loops by mapping strategic objectives to tangible KPIs, providing visibility into what’s working and why. For example, stakeholders no longer need to piece together insights from disparate monthly reports; they gain instant strategic clarity to questions like:

- “What is the current progress on our priority KPIs?”
- “Are our current omnichannel strategies effectively driving these KPIs, and why or why not?”
- “What adjustments can we make immediately to optimize performance?”

This enables teams to act swiftly, share best practices, and adapt strategies in real-time for continuous improvement.

Combining traditional machine learning and rules-based systems with newer, more dynamic AI techniques such as reinforcement learning, generative models and Agentic AI has become crucial. This balanced approach ensures recommendations are not only insightful and contextually relevant but also scalable and easy to manage.

Context-aware, continuously learning AI is essential for sustained global impact. At Aktana, we’ve embedded these lessons into our core design, empowering life sciences teams not only to execute today’s strategies more effectively but to dynamically evolve them as conditions change.”

Derek Choy
Co-founder and Chief Product Officer, Aktana

Frequently Asked Questions About **Agentic AI & Continuous Improvement**

Click on a question to jump to answer

Q1 Does Agentic AI use LLMs and NLP? What role do these technologies play?

Q2 Can Agentic AI solve optimization problems, such as resource allocation in pharma sales and marketing?

Q3 What are some real-world applications of Agentic AI in life sciences, particularly in pharma and biotech?

Q4 What makes a multi-agent system better suited for life sciences than a single-agent AI model?

Q5 What makes context-awareness the key differentiator in AI decision-making?

Q6 Commercial, marketing, and medical affairs teams often work in silos. How can AI-driven agents help them collaborate more effectively?

Q7 How does context-aware orchestration improve HCP engagement?

Q8 Everyone talks about AI improving sales force effectiveness, but what does that actually look like in practice?

Question 1

Does Agentic AI use LLMs and NLP? What role do these technologies play?

Yes, **Agentic AI** heavily uses **Large Language Models (LLMs)** and **Natural Language Processing (NLP)** for intelligent decision-making, especially in complex fields like life sciences.

Think of it this way: An Agentic AI system, like a human brain, needs to understand language, reason, and learn from experience.

- **LLMs and NLP (The Brain's Cognitive and Language Centers)**

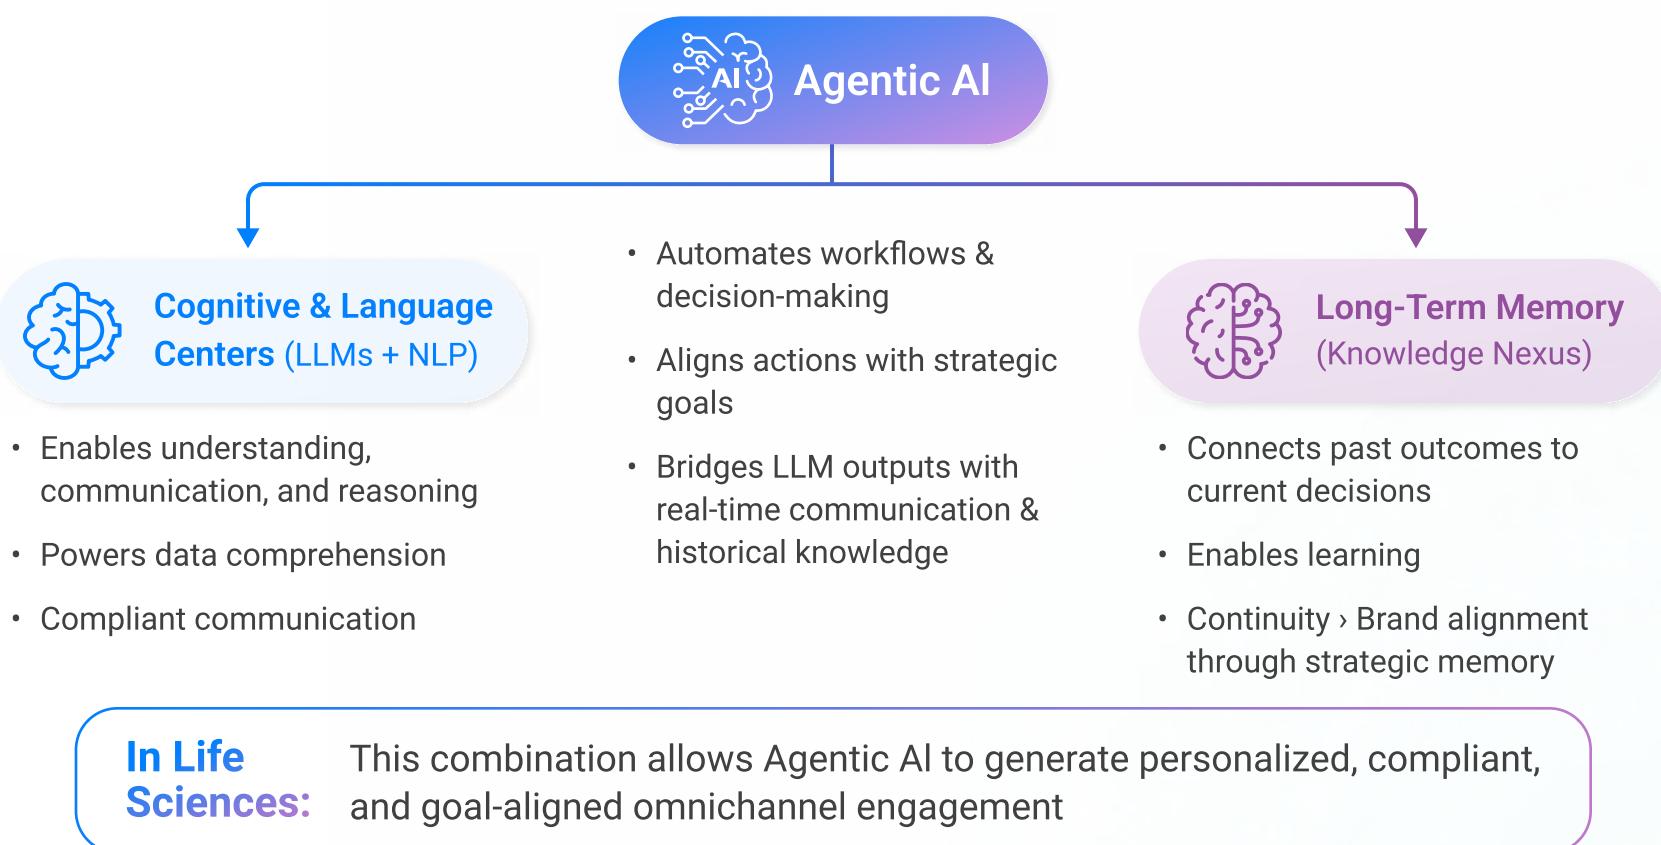
Together, LLMs and NLP form the core “brain” of Agentic AI. They enable the AI to process vast amounts of data, understand context, and generate recommendations. This allows for tailored communication, adjusting content and channels in real-time based on behavior and goals. NLP also helps the AI understand human input and communicate clearly, ensuring messages are compliant and relevant.

However, for true “agentic” behavior, these systems need long-term memory. This is where Aktana’s **Knowledge Nexus** comes in.

- **The Knowledge Nexus (The Brain's Long-Term Memory)**

The Knowledge Nexus acts as the Agentic AI’s long-term memory. It stores successful strategies and lessons learned, connecting current recommendations to past outcomes. This allows the AI to apply accumulated wisdom, ensuring brand strategies can be activated at a quicker pace and performance continuously improves.

In short, LLMs and NLP provide the cognitive power for understanding and communication, while the Knowledge Nexus adds the crucial element of memory and strategic learning. Together, they enable Agentic AI to deliver intelligent, personalized, and goal-aligned orchestration.



Question 2

Can Agentic AI solve optimization problems, such as resource allocation in pharma sales and marketing?

Yes. AI-powered agents, operating within an Agentic AI architecture, offer a significant advantage in optimizing resource allocation for pharmaceutical sales and marketing teams. By moving beyond static, rule-based approaches, they enable dynamic, data-driven decision-making that aligns commercial strategies with evolving market conditions and healthcare professional (HCP) preferences.

Here's how these agents address key optimization challenges:

- **Precision in Field Force Deployment:** AI agents analyze prescribing patterns, engagement trends, and competitive intelligence to guide the strategic deployment of sales teams—ensuring the right rep reaches the right HCP at the right time.
- **Adaptive Omnichannel Orchestration:** Agents facilitate the seamless integration and adjustment of digital and in-person interactions, dynamically sequencing channels to meet both brand priorities and HCP behaviors.
- **Data-Driven Marketing and Budget Allocation:** AI systems refine campaign strategies by identifying the most effective messaging, channels, and investment areas—while remaining compliant with regulatory and medical guidelines.
- **Continuous Learning and Agility:** Agents learn from real-time engagement data and user feedback, adjusting their recommendations as market dynamics and brand strategies evolve.

Through these capabilities, an agentic approach can empower pharmaceutical organizations to achieve continuous optimization at scale—driving improved efficiency, stronger HCP engagement, and greater commercial success. Plus, because every recommendation should be explainable and traceable, brand teams can retain full oversight and confidence in how resources are being allocated.

Question 3

What are some real-world applications of Agentic AI in life sciences, particularly in pharma and biotech?

Agentic AI is transforming how pharmaceutical and biotech organizations operate—automating compliance, accelerating R&D, optimizing engagement, and improving patient outcomes. Its modular architecture allows specialized agents to collaborate in real time, aligning decisions with strategy, regulation, and human oversight.

Here are some high-impact applications across the life sciences landscape:

- **Regulatory Compliance & Monitoring**

AI agents automatically flag potential compliance risks, track global regulatory updates, and support audit readiness.

Example: AI systems review promotional content pre-distribution to detect off-label language or outdated claims.

- **Personalized Medicine & Treatment Optimization**

AI analyzes clinical history, biomarker data, and prior outcomes to support precision treatment recommendations.

Example: In oncology, AI-driven platforms assist providers in selecting optimal therapies based on genetic profiles.

- **Drug Discovery & Development Acceleration**

Agentic AI accelerates R&D by identifying promising targets, optimizing molecular structures, and simulating trial outcomes.

Example: AI models help reduce preclinical screening timelines by rapidly narrowing down viable compound candidates.

- **Omnichannel HCP Engagement**

Agents orchestrate tailored digital and in-person interactions, aligning content delivery with HCP behavior and preferences.

Example: AI curates personalized digital experiences for HCPs, ensuring relevant, timely, and compliant messaging.

- **Sales & Marketing Intelligence**

AI enhances forecasting, territory planning, and rep targeting through real-time market signals and HCP data.

Example: AI-powered tools help sales teams prioritize high-value HCPs based on prescribing behavior and digital engagement patterns.

- **Real-World Evidence & Post-Market Surveillance**

AI continuously monitors real-world data to track outcomes, surface safety signals, and inform future strategy

Example: AI scans EHRs, claims data, and patient forums to detect emerging trends in adverse events for newly launched therapies.

Across these examples, Agentic AI enables dynamic, explainable, and context-aware decision-making - allowing life sciences teams to move faster without losing control.

Question 4

What makes a multi-agent system better suited for life sciences than a single-agent AI model?

In life sciences, decisions aren't linear—they're layered. Choosing the next-best action for an HCP requires balancing brand objectives, engagement history, regulatory boundaries, timing, field capacity, and more. These interdependent variables make the decision space highly contextual—and difficult to navigate with a single, monolithic AI agent.

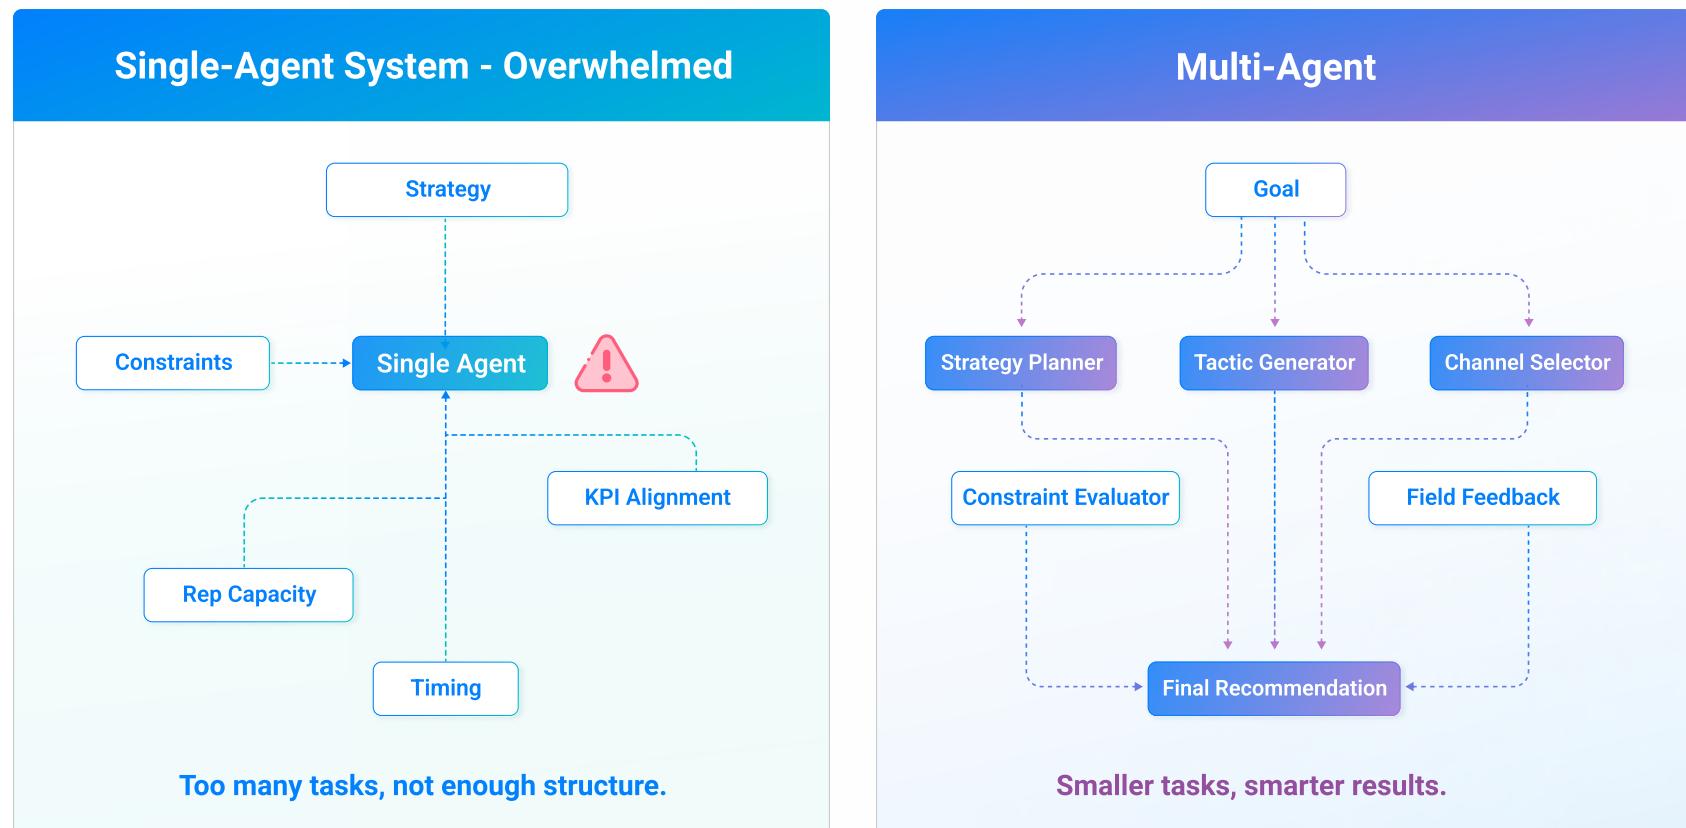
When one AI model is tasked with solving a complex goal like “determine the best action for this provider,” it lacks the structure to reason through competing priorities. The result is often oversimplified outputs or opaque logic, especially when constraints stack up.

That's why the most effective approach is a system of **specialized, collaborative agents**, each focused on a specific aspect of the decision process. One agent might interpret strategic objectives. Another simulates likely outcomes. Others assess channel constraints, timing, or prior engagement effectiveness. This division of labor—akin to **Chain of Thought Reasoning**—lets AI break down complex challenges into smaller, solvable parts while maintaining transparency, intent, and compliance.

Aktana's multi-agent architecture puts this into practice. By distributing intelligence across modular components—each aligned with a business function—our system ensures that recommendations are **explainable, strategic, and real-time**. Every decision reflects not just data, but brand goals, regulatory standards, and field realities.

In high-stakes, regulated environments like life sciences, this isn't just a design preference. It's the only practical path to scalable, trusted AI adoption.

Why Multi-Agent AI Works (and Single-Agent Doesn't)



Question 5

What makes context-awareness the key differentiator in AI decision-making?

In life sciences, static AI models can't keep up with evolving HCP behavior, shifting brand priorities, or real-time market dynamics. **Context-awareness** is what sets Agentic AI apart - enabling the system to adapt on the fly and deliver recommendations that reflect not just strategy, but the reality of the moment.

At Aktana, context-aware decision-making is powered by:

- **The Contextual Intelligence Engine (CIE)**

Our real-time engine continuously synthesizes live data—such as HCP preferences, field activity, and channel availability—to prioritize and deliver next-best actions across both field and digital channels. CIE ensures that every decision is immediately relevant, while still aligned with brand goals. (Read more: [Contextual Intelligence Engine](#))

- **The Knowledge Nexus**

Acting as the system's long-term memory, the Knowledge Nexus captures what's worked (and what hasn't) across brands, markets, and use cases. Tapping into 12+ years of contextual data to inform our agentic decisions, it provides ready-to-use blueprints to boost strategic effectiveness and continuous learning across deployments. (Read more: [Agentic AI](#))

Together, these components ensure that every AI-driven recommendation is timely, explainable, and aligned with brand strategy—even as that strategy evolves. **Because decisions are made transparently and in real time, teams can act with confidence knowing they retain oversight and influence at every step.**

Question 6

Commercial, marketing, and medical affairs teams often work in silos. How can AI-driven agents help them collaborate more effectively?

AI-driven agents bridge cross-functional silos by sharing context, coordinating actions, and adapting in real time. They operate with awareness of shared goals, incorporate feedback from all functions, and adjust based on what's actually happening in the field.

By maintaining a **dynamic memory** of past actions, content usage, and planned engagements, agents ensure teams stay aligned even if they operate within different systems or workflows.

For example, if a medical inquiry signals HCP interest in a new indication, an agent can capture that signal and surface it across functions—prompting marketing to update its content arc and commercial to prioritize follow-up. Each function remains informed and responsive, without relying on manual updates or coordination meetings.

Aktana's agentic architecture supports this level of synchronization through its **Optimizer** and Knowledge Nexus - two core components that ensure decisions are both coordinated and contextually intelligent. The Optimizer functions as a dynamic orchestration engine, continuously aligning actions across teams based on real-time priorities, constraints, and strategic goals. Meanwhile, the Knowledge Nexus acts as a shared memory system, capturing insights and outcomes across brand, field, and medical functions to inform future decisions. Together, they ensure life sciences teams engage HCPs with a unified, coherent voice—not three separate ones—while preserving the transparency and control each function needs to stay accountable.

Question 7

How does context-aware orchestration improve HCP engagement?

Omnichannel success depends on coordination. That's where AI-driven orchestration delivers the greatest impact. Instead of relying on disconnected campaigns or static plans, AI agents bring structure, timing, and personalization to every interaction.

Agents prioritize actions across field and digital channels based on real-time context: what an HCP has already seen, what they've engaged with, and what's most likely to resonate next. This ensures each touchpoint feels connected and relevant—not repetitive or out of sync.

For example, if an HCP opens a clinical update email, an agent can suppress similar content, queue up a rep visit with complementary materials, or escalate to a higher-value tactic. Every step is sequenced using a full view of the HCP—not just one team's inputs.

At the core of this orchestration is real-time learning in Agentic AI—a feedback-rich approach where every action informs the next. Aktana's multi-agent architecture powers this orchestration through dynamic segmentation, intelligent timing, and continuous learning. And because every action is transparent and traceable, teams can confidently scale engagement—delivering the right message, at the right time, in the right way.

Question 8

Everyone talks about AI improving sales force effectiveness, but what does that actually look like in practice?

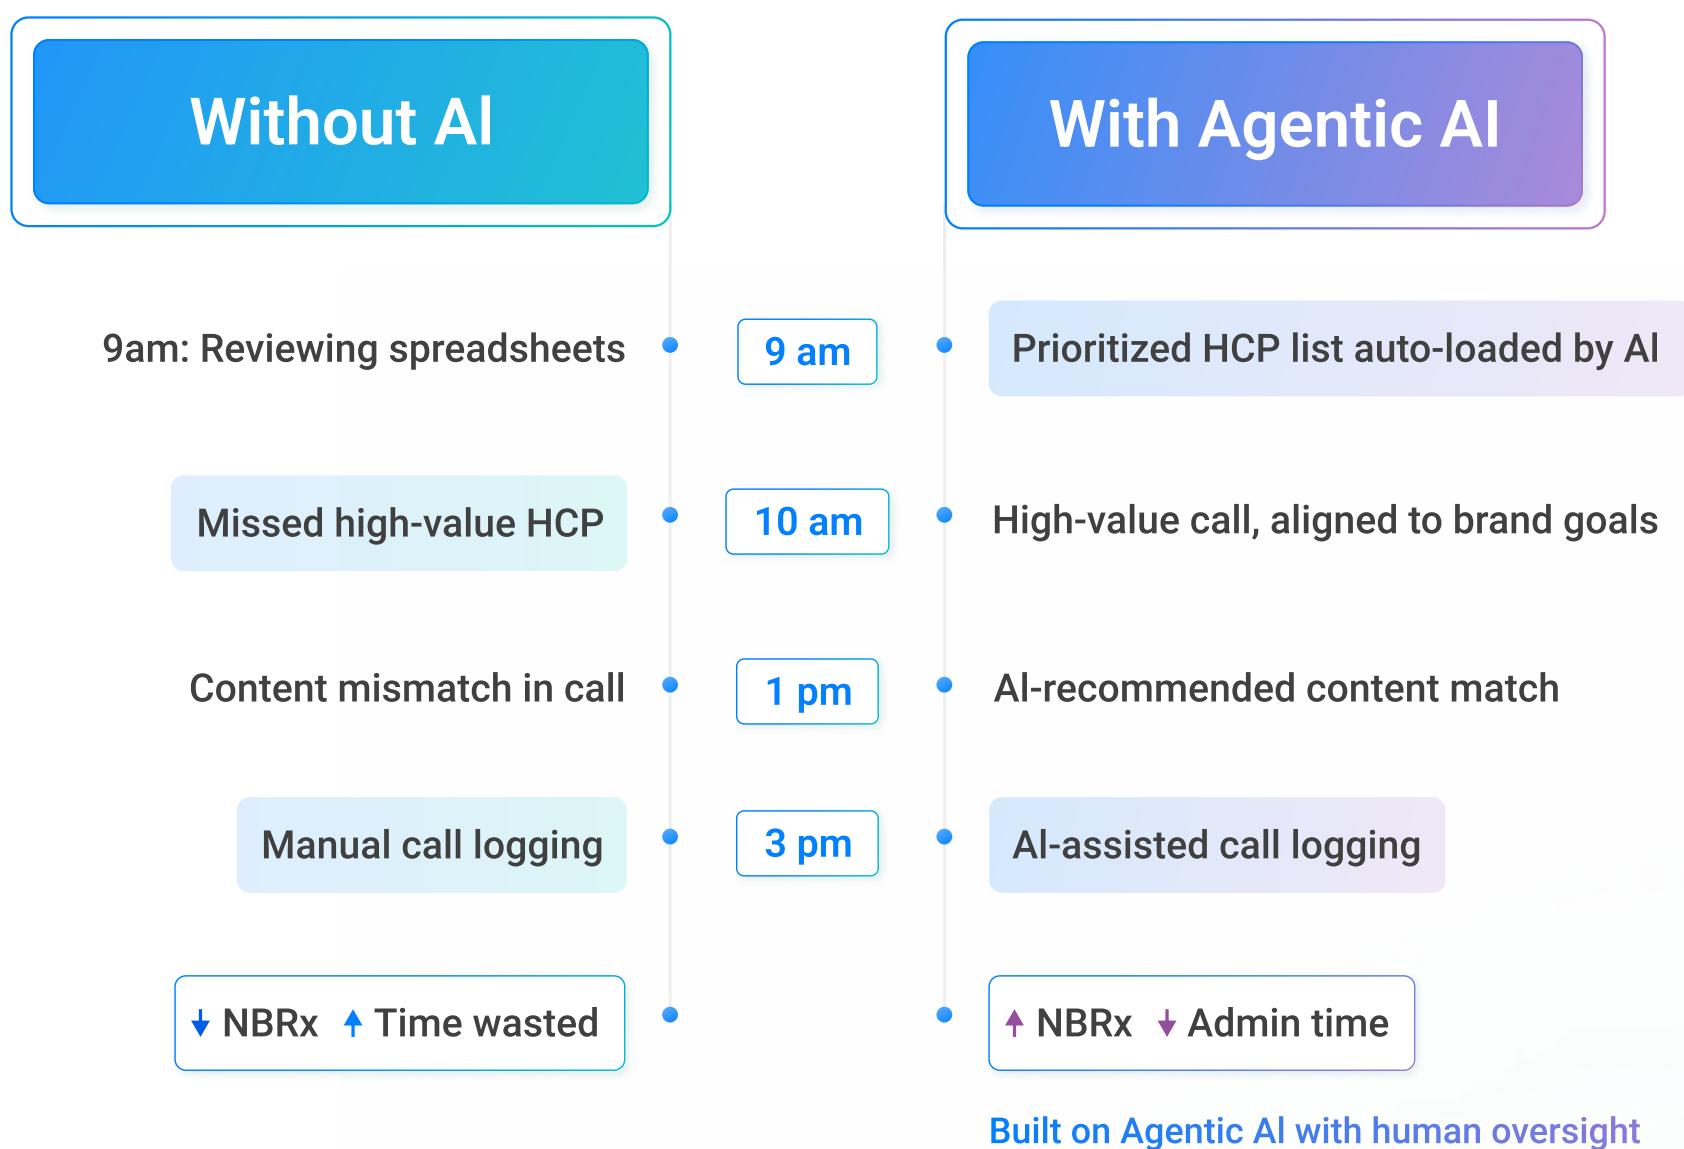
AI improves sales force effectiveness not by adding more tasks, but by helping reps focus on the right ones. From pre-call planning to territory coverage, AI agents streamline decision-making so field teams can spend more time engaging and less time guessing.

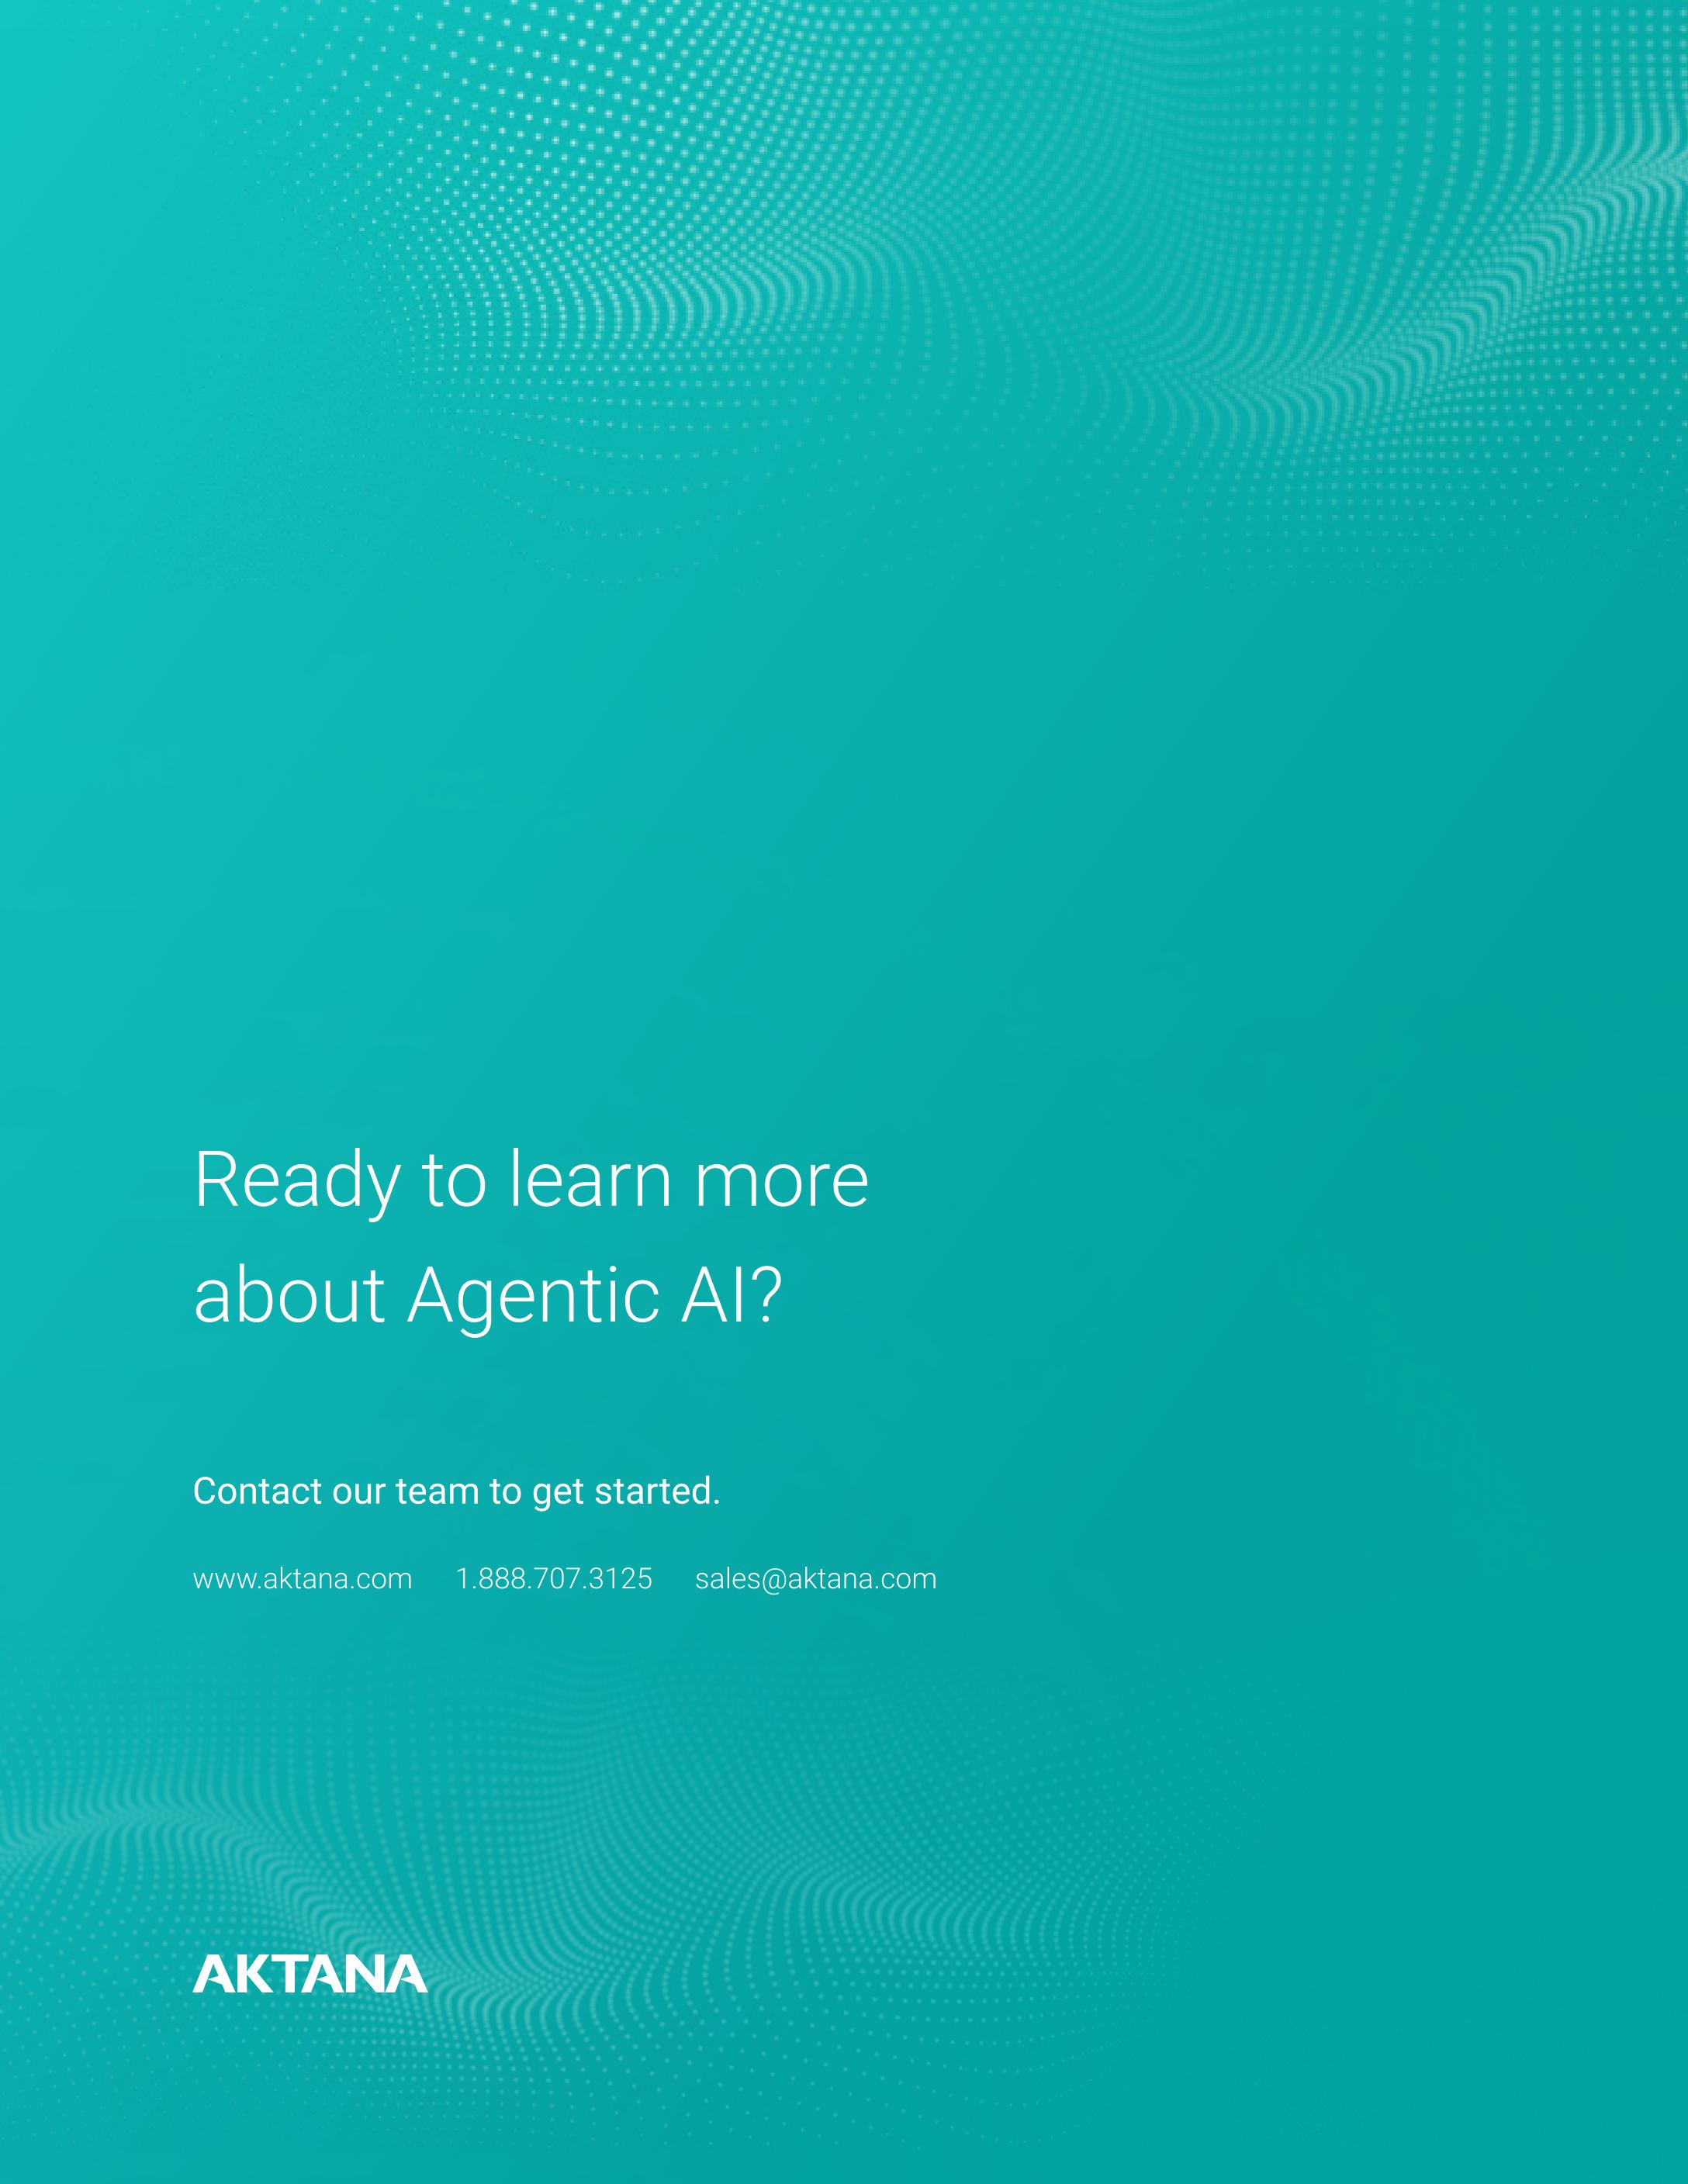
Instead of relying on static call plans or gut instinct, reps receive context-aware suggestions: who to prioritize, what to say, which channel to use, and when to engage. These recommendations reflect brand strategy, HCP behavior, rep availability, and compliance rules, ensuring every action is purposeful and aligned.

For example, if a rep has limited capacity this week, AI might prioritize HCPs showing rising engagement trends, suppress outreach to recently contacted providers, and surface pre-approved content tied to brand KPIs—all without requiring manual coordination.

With Aktana, these suggestions aren't just personalized—they're explainable, auditable, and continuously improving. Reps gain confidence, managers see measurable impact, and the entire commercial operation becomes more strategic, agile, and aligned.

A Rep's Day: Before vs After Agentic AI





Ready to learn more
about Agentic AI?

Contact our team to get started.

www.aktana.com 1.888.707.3125 sales@aktana.com

AKTANA