Quiq

VUX.

Beyond the Buzz:

How to identify and evaluate Agentic AI platforms for CX

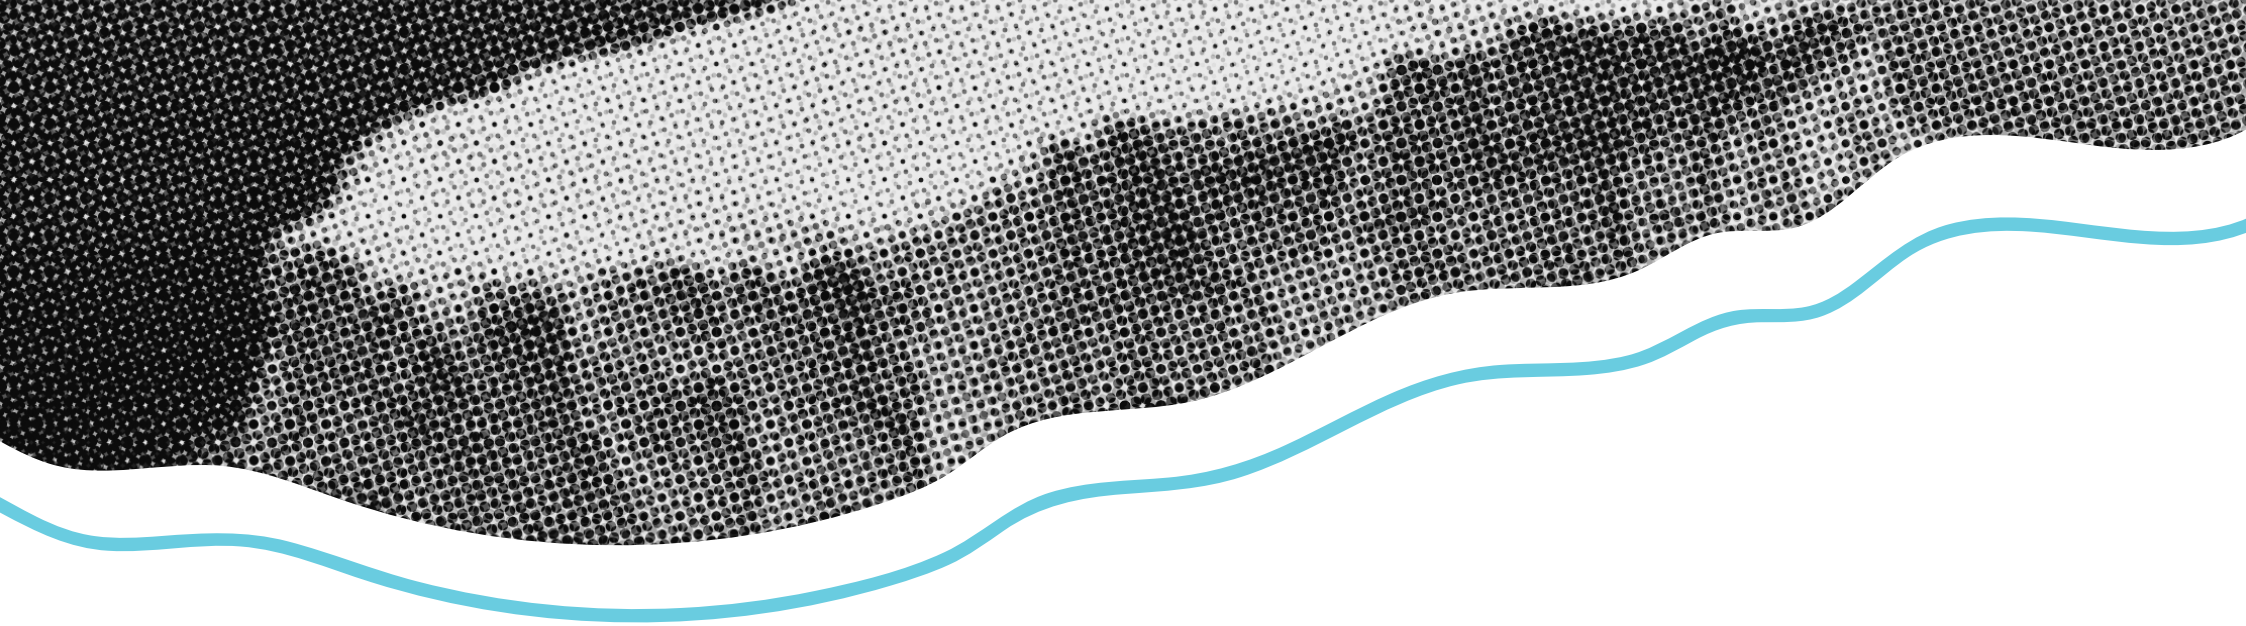

The adoption of agentic AI is not just a trend, but a strategic imperative.

Scripted chatbots and basic automation have served their purpose over the last decade or so. However, they have inherent limitations that place a ceiling on success. Large Language Model-based (LLM) generative AI solutions, orchestrated into the shape of AI agents, have the potential to surpass that ceiling and redefine the kind of outcomes that can be achieved in AI automation programs.

The adoption of agentic AI is not just a trend, but a strategic imperative. According to a Frost & Sullivan report1, agentic AI is among the top trends set to dominate 2025. AI is forecast to autonomously resolve 80% of common customer service issues, and create up to 30% reduction in operational cost by 2029, according to Gartner2.

The problem is that, today, everything is marketed as agentic. For business leaders, how can you understand what agentic AI is and how it applies to your business goals? How can you identify the kind of technology you need and the specific benefits you can achieve through leveraging these agentic capabilities?

This whitepaper seeks to solve that problem for you. It delves into the concept of agentic AI, exploring its defining characteristics, real-world applications and practical guidance for evaluating agentic AI solutions. By understanding and leveraging agentic AI, CX leaders and technical stakeholders can position their organizations at the forefront of customer experience innovation.

Defining agentic AI

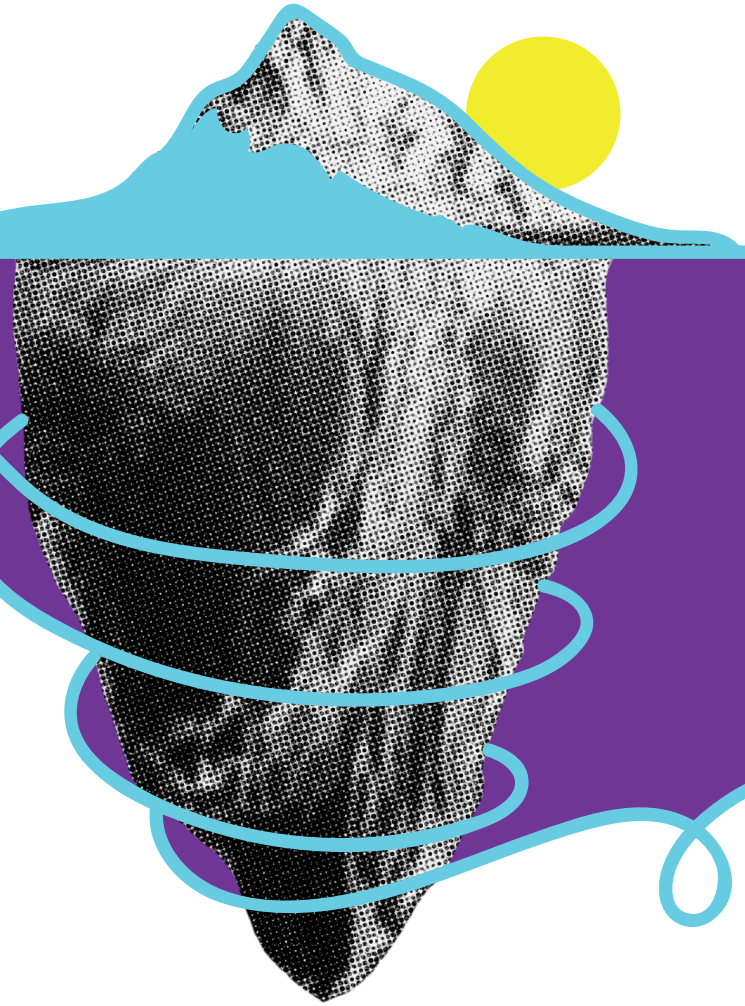

The generally accepted definition of an AI agent comes from the definition of software agents in distributed artificial intelligence and multi-agent systems. Franklin and Graesser3, in their landscape-defining literature review, summarize and define an agent as:

> **A system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future."**

In Michael Wooldridge's 2009 book, An Introduction to MultiAgent Systems4, he builds on the above definition and provides a simple and memorable framework for what agents do: Agents Sense, Decide and Act within their **Environment** to fulfil a given agenda, learning from their actions so as to inform what it senses.

The degree to which a system can operate on its own, and decide how to progress through this **Sense**, **Decide** and **Act** loop could be seen as a measure of its 'agency' or 'agenticness'.

Applying this agentic AI framework to an enterprise AI platform

Let's walk through this framework with the lens of an enterprise AI platform, to help you understand how each of these elements manifest into something that would enable you to leverage this technology in the real world. The platform we'll use for this example is Quiq, an enterprise agentic AI platform that enables businesses to build and orchestrate their own agents for the purposes of CX automation. The outcome of that action is fed back into the Sense, Decide, Act cycle when the customer responds with, for example "Actually, that's no good. Can you do Thursday?"

> At a high level, we're not specifying exactly what an agent should do. We're not necessarily defining the intents that the agent should have. We're defining an environment that the agent should work inside of, and then the agent has agency inside of that environment." **Mike Myer, CEO, Quiq**

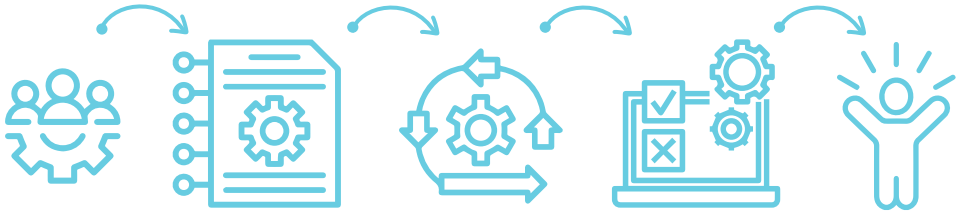

Environment:

Setting up your environment includes the channels that your agent will be accessed via (phone, chat etc), as well as the goals that the agent has to accomplish, the rules it should follow in accomplishing them and the tools that it can use to do so.

For example, an agent that helps customers to arrange an appointment would be given that as its goal, along with some specific rules (or guides). For example, "only offer appointment slots within business hours, avoid double-booking by checking real-time calendar availability via an API."

Sense:

The agent then continuously monitors its incoming stimuli (customer utterances) to understand the customer's intent and context. In this case, it detects that the customer wants to book an appointment by parsing natural language input such as "I'd like to schedule a time to talk" or "Can I book a slot for next week?"

Decide:

Based on what it has sensed, and the guidance that you've provided, the agent evaluates what it should do next, using its Cognitive Architecture (more on that later). This could lead to a decision wherein, the best thing to do, is to find the next available appointment slot by using a specific tool.

Act:

Once a decision is made, using the tools and APIs you've granted it access to, the agent takes action. In this scenario, it might include making an API call to retrieve the next available appointment and presenting that back to the user.

This whole Sense, Decide and Act loop will unfold at a turn-by-turn level throughout the conversation, and can also describe the holistic, high-level process of an agent.

VUX.

What's happening under the hood of an AI agent?

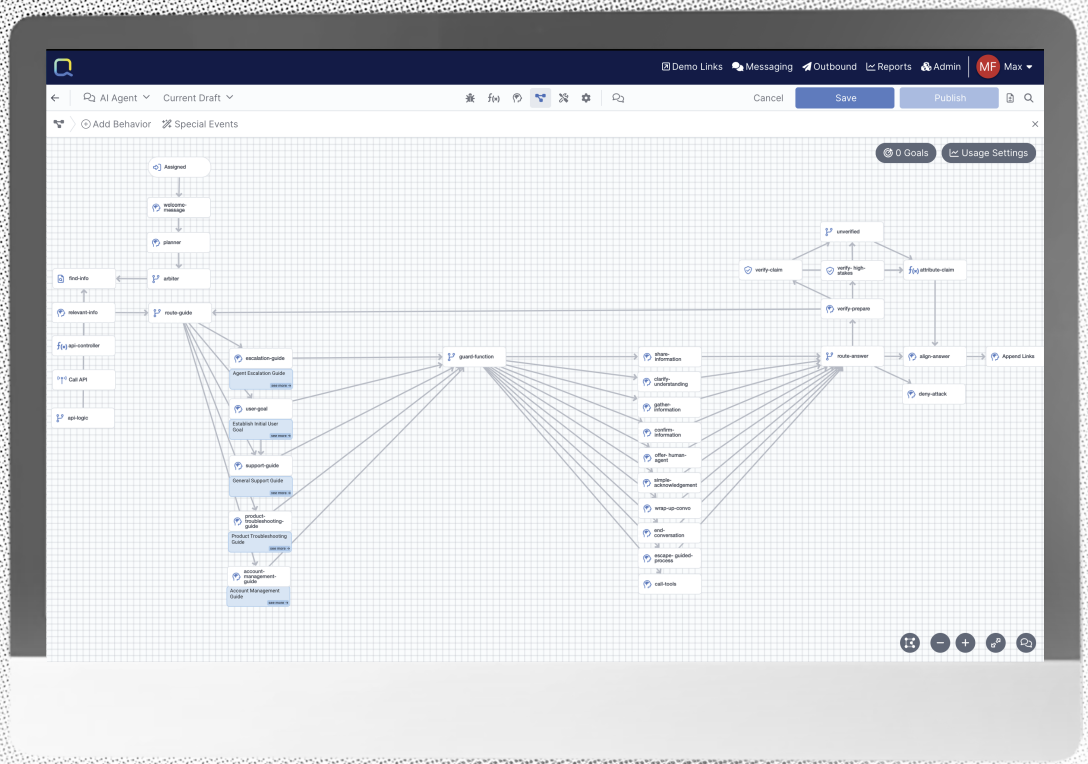

> Cognitive Architecture is like an employee who is capable, has been trained, follows operational procedures, can access relevant systems and data to carry out a task, and has been given guidance on proper utilisation of the various systems they use." **Mike Myer, CEO, Quiq**

To continue using our example of Quiq, the platform leverages what it terms a Cognitive Architecture.

This senses each input event, such as a message from the user, a file upload, a web page being loaded or any related event triggered from a third party system. Each event is sensed by the system and analyzed by the agent's knowledge, context, and capabilities. Based on this, the agent will choose, switch or continue executing a multi-step guided process.

In Quiq terms, these guided processes are akin to a standard operating procedure for a human agent. Guides provide the agent with a set of heuristics and tools that clarify intent, gather data, retrieve knowledge, perform API actions and generate responses.

In practical terms, the Cognitive Architecture provides the AI agent access to the context, guides, reasoning and tools a human agent would have when serving as an employee of a company.

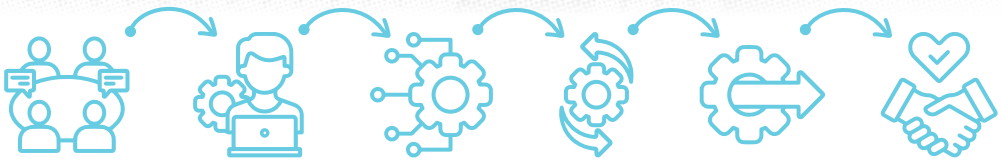

Key components of Cognitive Architecture

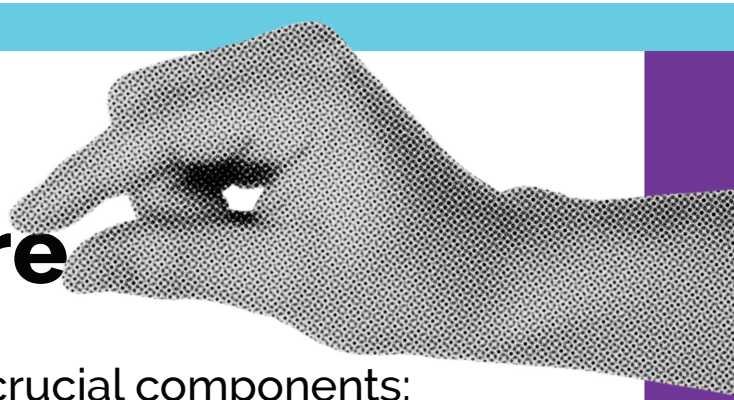

The Cognitive Architecture encompasses several crucial components:

Planning and Decision Making:

Language models analyze user input within the context of the conversation to determine the next appropriate step, select the relevant guide or identify required data retrieval.

Consulting Standard Operating Procedures:

Human-readable guides provide the AI with general instructions on how to handle different types of inquiries, outlining process steps, tools and decision-making criteria.

Information Retrieval and Distillation:

The system retrieves pertinent data from various sources and then refines it to extract the essential information.

Utilising Core Messaging Tools:

These tools enable the AI to share information from knowledge bases, ask clarifying questions based on guides and conversation state, and perform other standard communication functions.

Employing Guide-Specific Tools:

These tools include vector search, data gathering and API calls to access backend systems that are specific to the human readable guide chosen in the planning phase.

Safety Checks and Guardrailing:

Evidence is collected to support the LLM's claims. A specialized AI model then verifies if this evidence substantiates the LLM's assertions, helping to prevent inaccuracies and ensure appropriate responses.

This architecture facilitates more adaptable and context-aware interactions when compared to rigid flow-based systems. For example, in the appointment setting example we used below, after the AI agent had offered some appointment times, the customer may ask "Is there an additional charge for a weekend appointment?". With a flow-based system, this question wouldn't be able to be answered because the agent was in the midst of an appointment flow and was expecting the customer to confirm a time. With Quiq's Cognitive Architecture, this question can be handled and then the appointment booking process is continued.

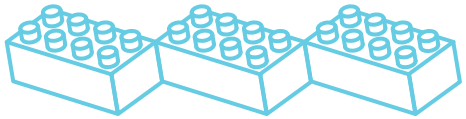

Due to the fact that every user input is fed through the entire Cognitive Architecture, development teams can ensure that they have suitable checks and balances in place so that users aren't misunderstood, mishandled, or given wrong or hallucinatory responses.

VUX.

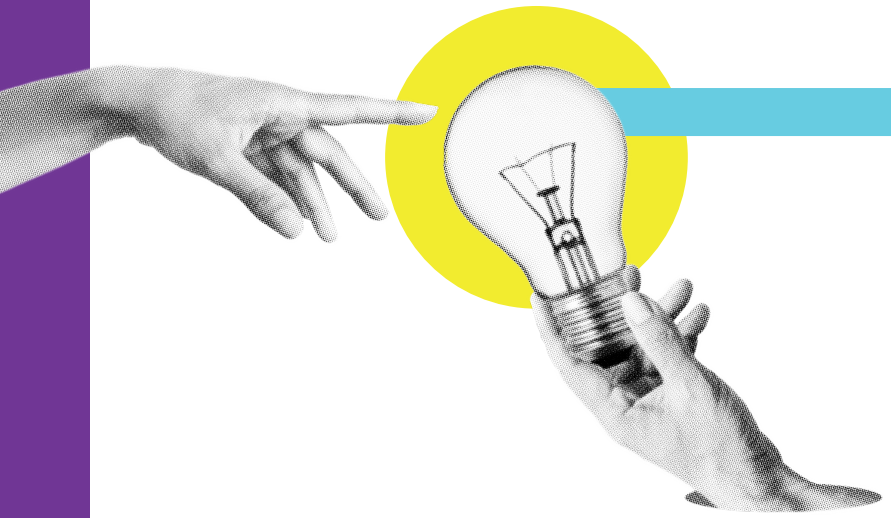

Monitoring, Debugging and Analyzing Agentic AI

Critical considerations for quality control and analysis of agentic AI systems include:

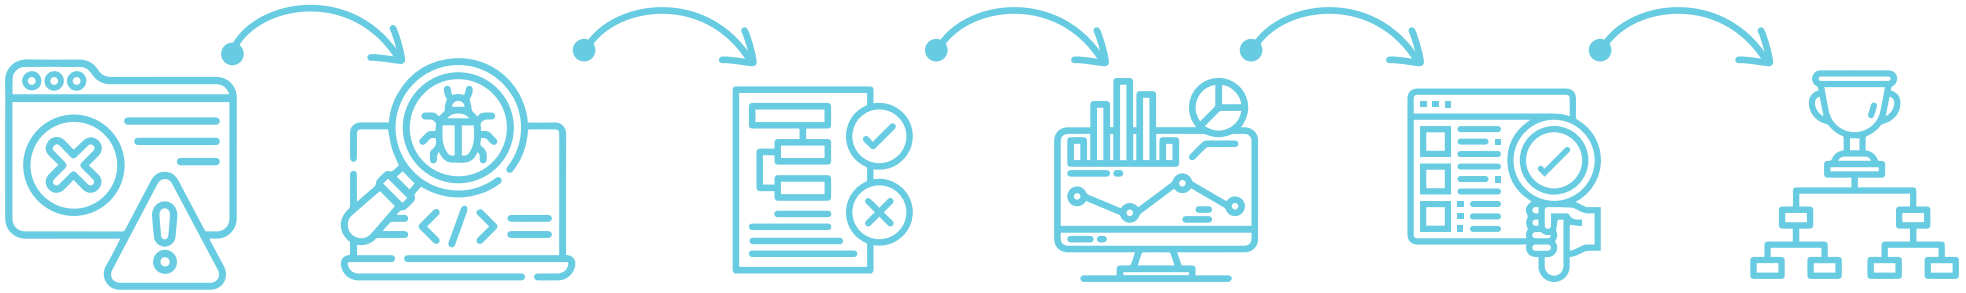

Hard and Soft Errors:

Agentic AI systems must detect and respond to failures—soft errors notify users and offer escalation paths, while hard errors trigger immediate handoff to human agents.

Debug Workbench:

Platforms should provide tools to track, analyse, and review errors at the prompt-level, throughout the pipeline, so that you can properly assess and debug exactly what's happening when, and why.

Conversational Grounding Errors:

Misunderstandings in conversation (e.g. incorrect interpretations or hallucinations) can be debugged by replaying and analysing interactions to improve prompt design or model selection.

Testing and Regression:

Past conversations can be saved as test cases with expected outcomes, enabling automated regression testing and validation of new updates before deployment.

Claim Verification:

Agents should be able to verify their own outputs using supporting data—separate models can confirm whether the AI's responses are backed by evidence.

Post-Conversation Analysis:

After conversations end, language models can assess effectiveness, helping teams identify issues, surface insights, and continuously improve agent behaviour.

VUX.

Agentic AI in Action

Now that we've taken a look at the architecture of an agent, let's see one in practice. Consider the following practical application of Quiq's Cognitive Architecture in managing a customer inquiry regarding home delivery:

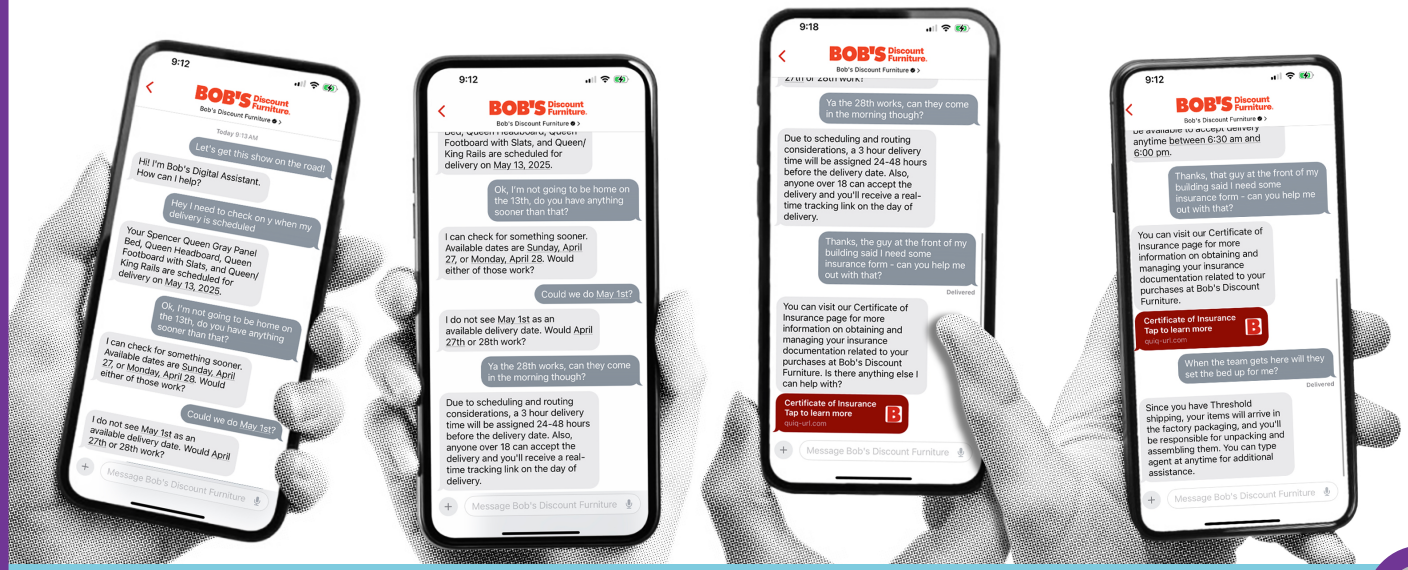

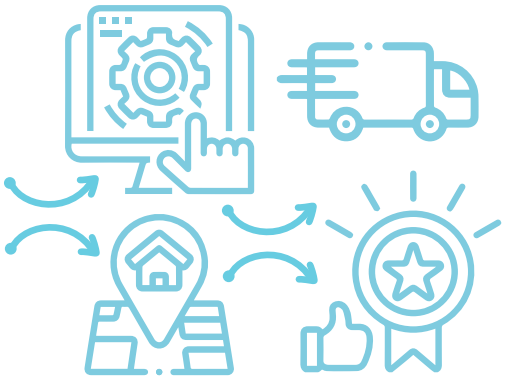

Customer: Asks to reschedule.

Agent: Instead of following a pre-defined script, the agent, operating within the Cognitive Architecture, processes this input and selects a guide that outlines how to handle scheduling. Thisincludes tools for API integration to access the customer's order details.

Customer: Requests an earlier delivery time.

Agent: The agent processes this new input by first determining whether the prior guide still applies or if a different guide should be selected based on the customer's evolving needs in the multi-turn conversation. The guide provides instruction to the agent how handle a request for a specific time.

Customer: Requests proof of insurance.

Agent: The agent processes this new input by first determining the most appropriate guide to use. In this case, the customer is asking a question that falls into the general information category, so the agent is instructed to retrieve relevant information from a knowledge base.

Customer: Inquires about furniture setup.

Agent: The agent determines that the delivery information guide should be used. This guide includes instructions for the agent to look-up order-specific information to see if "white glove delivery" (including setup) was purchased.

The AI utilises its access to the complete platform toolkit to manage every user request within a unified conversational thread. It begins with the context of the previous conversation (the customer's order of a Queen size bed), and adapts to the currently evolving conversation.

Rather than building this exact flow in a deterministic way, utilising Quiq's Cognitive Architecture allows for agents that can handle evolving multi-turn flows gracefully, with the usage of APIs for tasks such as calendar lookups and bookings, and without losing context.

VUX.

About VUX World

The front door to the world of AI-powered customer experience, helping businesses transform CX with conversational and generative AI technologies.

About Quiq

Quiq's customer-centric AI solutions are for enterprise brands that need to improve CX outcomes and want to leverage AI to do so while empowering both AI assistants and humans to play their strongest roles.

References:

1. AI: Top 10 Predictions, 2025, Future Growth Opportunities for Innovation and Transformation in the AI Industry, Frost and Sullivan, 2025
2. Agentic AI Set to Transform Customer Service & Support Landscape, Reshaping Inbound Interactions and Forcing Service Teams to Embrace Automation, Gartner, 2025
3. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents, Stan Franklin and Art Graesser, 1996
4. An Introduction to MultiAgent Systems - Second Edition, Michael Wooldridge, John Wiley & Sons, May 2009