Overview of LLM AI at Quiq

Quiq

Overview of LLM AI at Quiq

Introduction.

This paper describes AI in customer-facing, AI agents, and human-agent assistance features that Quiq builds on third party LLMs. We refer to these features as AI solutions in this paper. None of the opinions in this paper are intended to describe the perspectives of the third-parties, including OpenAI, Anthropic and others, who are the developers of the AI models that are the source(s) of any LLM(s) used in Quiq's AI solutions.

At Quiq, we understand the profound impact that artificial intelligence (AI) can have on revolutionizing customer experience and engagement. As we innovate and implement AI solutions, we remain steadfast in our commitment to responsible and ethical AI solutions. This document delineates Quiq's principles for the responsible development of AI solutions, emphasizing alignment, transparency, privacy, and quality.

Quiq leverages Large Language Models (LLMs) for customer and agent-facing AI solutions, for orchestration of language models grounded in domain-specific CX knowledge and processes. Our platform is designed to support a wide range of LLMs, allowing our customers to select the most appropriate model for each specific task. Quiq does not build its own LLM. Quiq licenses LLMs from existing third-parties so that customers may select the best LLM to meet the needs of their Agentic AI. Quiq can also use LLMs provided by customers. A full list of Quiq's LLMs is found in the product and is updated frequently to incorporate the latest models.

In partnering with third party AI models for LLMs, we prioritize data privacy and security. It's important to note that our customers' data is not shared with third parties to train models.

This approach allows us to harness the power of AI while maintaining data protection and brand standards, ensuring that our AI solutions are not only effective but aligned to our customers' brand guidelines.

Quiq Security Page

Quiq Security Whitepaper

Alignment.

As a flexible platform for building AI, our goal is to deliver the business outcomes desired by our customers while aligning to the brand’s persona and standards. Quiq's AI Studio is designed to promote brand alignment through a comprehensive approach involving deterministic and prompt-based data filtering of AI inputs and outputs. While these methods significantly enhance alignment with customer needs and brand standards, we acknowledge that achieving perfect alignment requires continuous monitoring and improvement.

Our approach includes:

- **Interaction classification**: Our Agentic AI perform several classifications—such as topic, user type, sentiment, and message sensitivity—specific to each company's needs. This classification process is crucial for ensuring alignment with business requirements and use cases.
- **Entity Extraction and Deterministic filtering**: We apply clarification and entity extraction to deterministically filter knowledge queries. For example, a question related to vehicle model will apply a model filter to ensure only relevant information is considered by the LLM.
- **Prompt-based filtering**: We also rigorously test prompts to ensure they focus large amounts of information to only what’s relevant to the conversation. This targeted information guides the LLM when generating responses.
- **Safety & Tone**: We use carefully designed prompts to guide our Agentic AI towards brand verified and aligned responses. This includes instructions that reflect the customer's brand voice and generated response claims verification.

To assess and improve alignment, we employ a multi-faceted approach:

- **Quantitative metrics**: We measure how closely AI solution outputs match predefined alignment criteria.
- **Qualitative analysis**: Our team conducts manual reviews of outputs to ensure they align with customer expectations and brand standards.
- **Ongoing monitoring**: We continuously track outputs for misalignments and regularly update our filtering methods.

Overview of LLM AI at Quiq

We recognize that customer-provided materials play a crucial role in alignment. To address this:

- We provide guidelines to customers on creating training data that aligns with their business goals and brand standards.
- Our system includes checks to flag content that may not align with the customer's stated objectives.
- We offer consultation services to help customers identify areas where better alignment can be achieved.

Customers have full observability into the decision process of the AI solutions and can analyze them at any time to determine if any adjustments are needed to improve alignment. Our platform is inherently flexible, accommodating diverse business needs and brand standards to allow for customization to meet various customer demands across different industries and use cases.

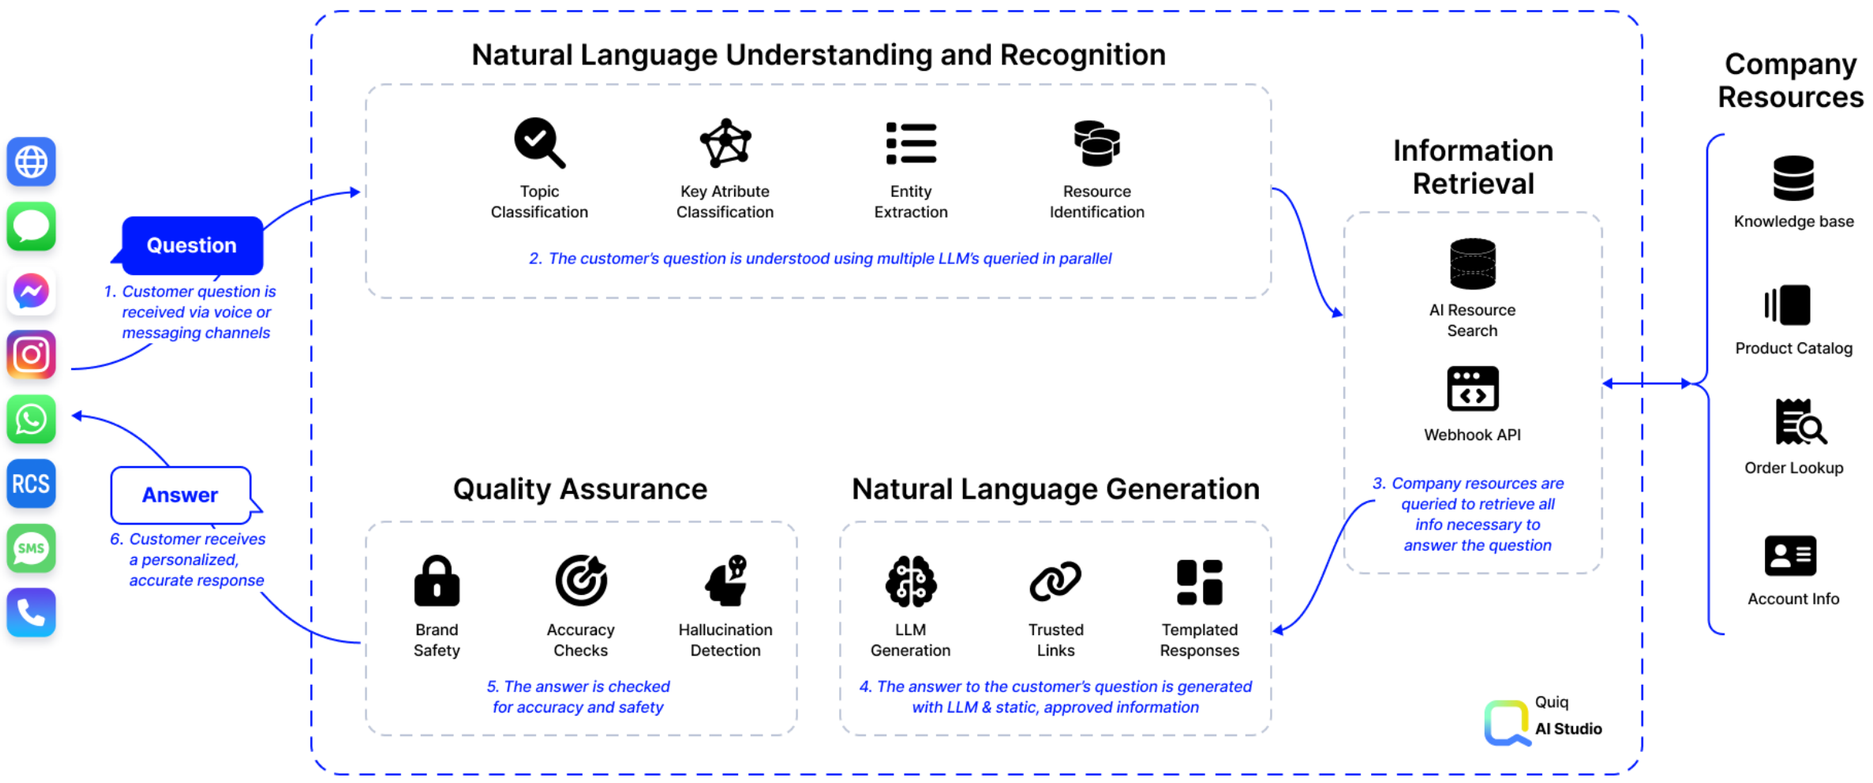

Transparency.

At Quiq, we believe in the importance of transparency in our AI systems. We are committed to clearly documenting the purpose, capabilities, and limitations of our AI solutions. Resource integration is a key part of this transparency. Once an inbound message is determined to be in scope, company-approved resources—such as knowledge articles, brand facts, product catalogs, and user-specific data—are retrieved for the Agentic AI to consult. This ensures that the AI's responses are not only accurate but also transparent in their sourcing, providing customers with a clear understanding of how and why certain responses are generated. Our AI Studio platform is designed to provide visibility into the workings of our Agentic AI, allowing our customers to understand and customize their behavior.

Privacy.

Quiq is committed to stringent ethical standards and the safeguarding of user privacy. Our platform collects only the essential data necessary to facilitate conversations, such as SMS numbers. Any additional required data can be customized within the product by our customers. We prioritize the protection of data and strictly adhere to relevant privacy regulations.

SOC 2 TYPE II
Quiq has undertaken the AICPA SOC 2 Type II certification to attest to Security, Confidentiality, and Availability controls in place in accordance with the AICPA Trust Service Criteria.

Encryption
All data in Quiq Messaging is strongly encrypted, both during transmission and at rest.

Data center security
Quiq uses Amazon Web Services. AWS is the industry standard for cloud computing infrastructure, featuring ISO, SOC, and CSA compliance programs.

Security testing
Vulnerability scans are performed regularly to ensure the Quiq system is not open to hackers. We use third-party Pen testing in addition to manual and automated internal vulnerability scans.

Fault-tolerant
Quiq is built using a microservices architecture with multiple nodes for every service so that there are no single points of failure. Quiq is fully redundant across all of our data centers.

Disaster recovery
Quiq tests and maintains Disaster Recovery and Business Continuity plans to help ensure a high level of continuous service, even in a geographical disaster. Backups are continuously mirrored to other AWS data centers in geographic locations.

Defense-in-depth
A good security strategy should never rely on just one thing, which is why Quiq uses redundant layers of protection throughout our stack.

Data stewardship
Your data in Quiq belongs to you. Quiq will never use your confidential information for any other purpose than to serve your business.

Ensuring Safety and Precision.

Quiq's AI Studio provides a toolkit to produce supervised Agentic AI that operate within the rules of your business domain. This toolkit includes conversational classification, a rules engine, prompt instructions, claim verification, regression testing, full observability and exception handling. This approach allows us to deliver AI solutions that are precise, powerful and safe, while adhering to standards of responsibility. Additionally, we actively monitor and endorse the OWASP Top 10 for Large Language Model Applications, further supporting our commitment to safety and precision.

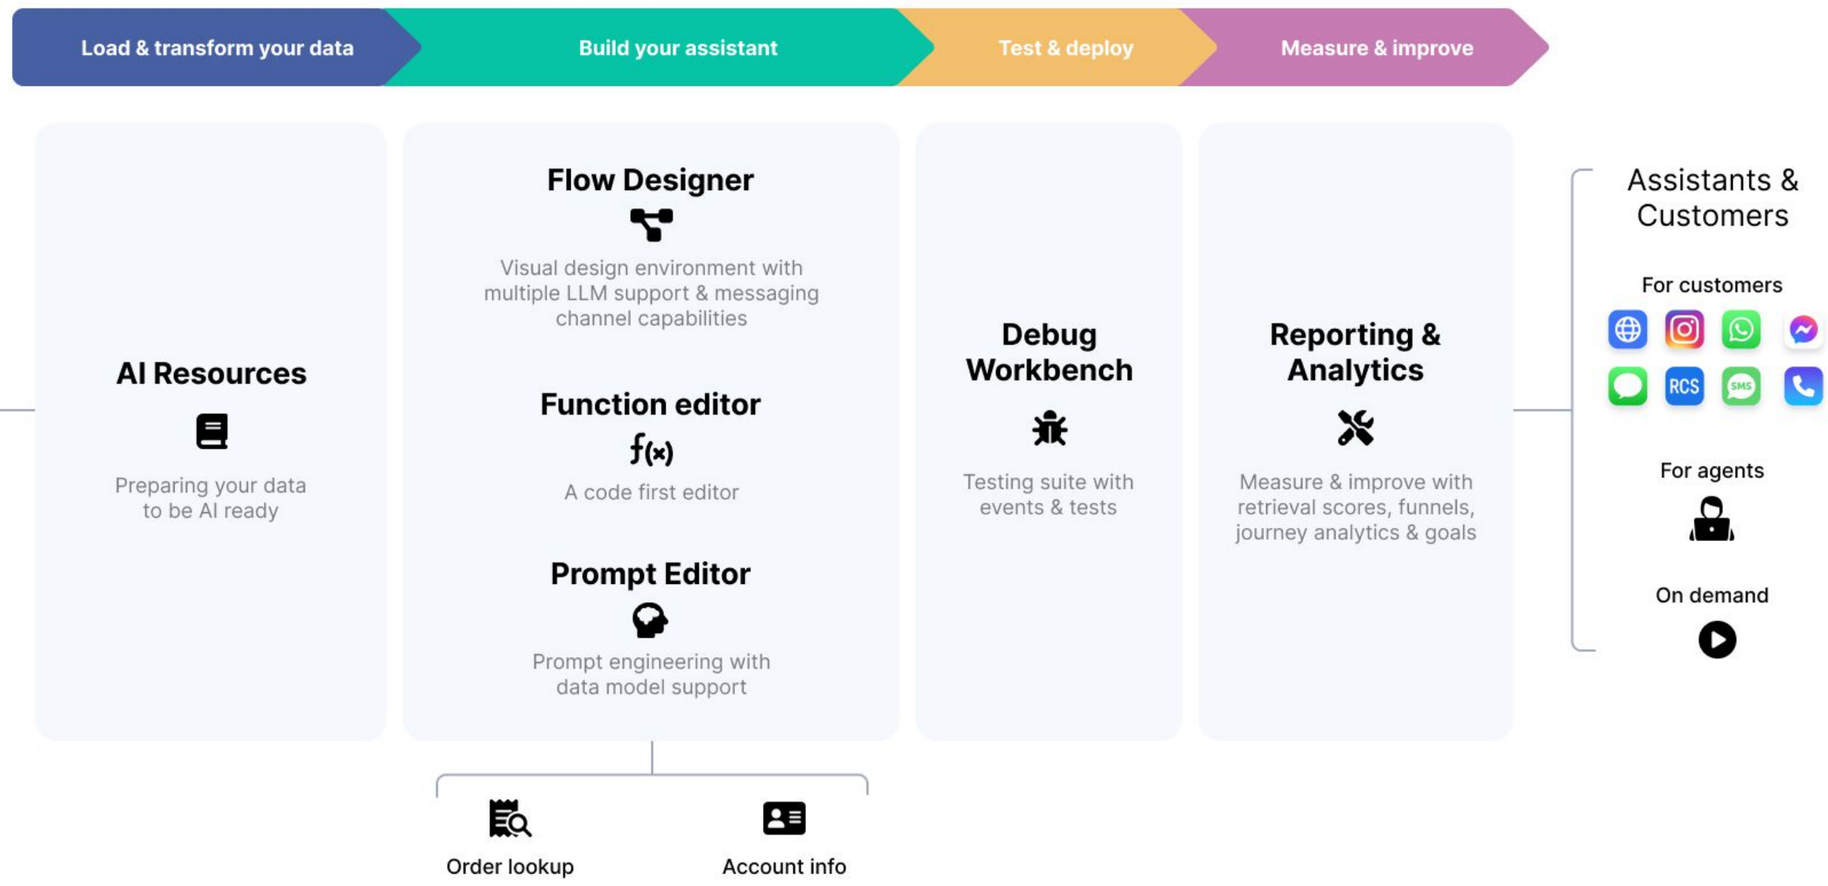

Risk Mitigation.

Quiq takes a comprehensive approach to risk mitigation by implementing pre- and post-answer generation guardrails, ensuring that our Agentic AI provide answers only from approved sources and that these answers are corroborated by the information or datasets they reference. This process involves filtering inbound and outbound messages to prevent harmful content and ensure appropriate responses. Inbound filtering uses semantic similarity searches to identify and reject malicious or out-of-scope messages. Outbound filtering checks for prompt leakage, semantic similarity to source data, and ensures all links and contact information are relevant. These measures help brands confidently adopt AI solutions while managing risks.

Response verification is also crucial. Involving thorough checks to ensure AI-generated responses align with brand values and are factually accurate.

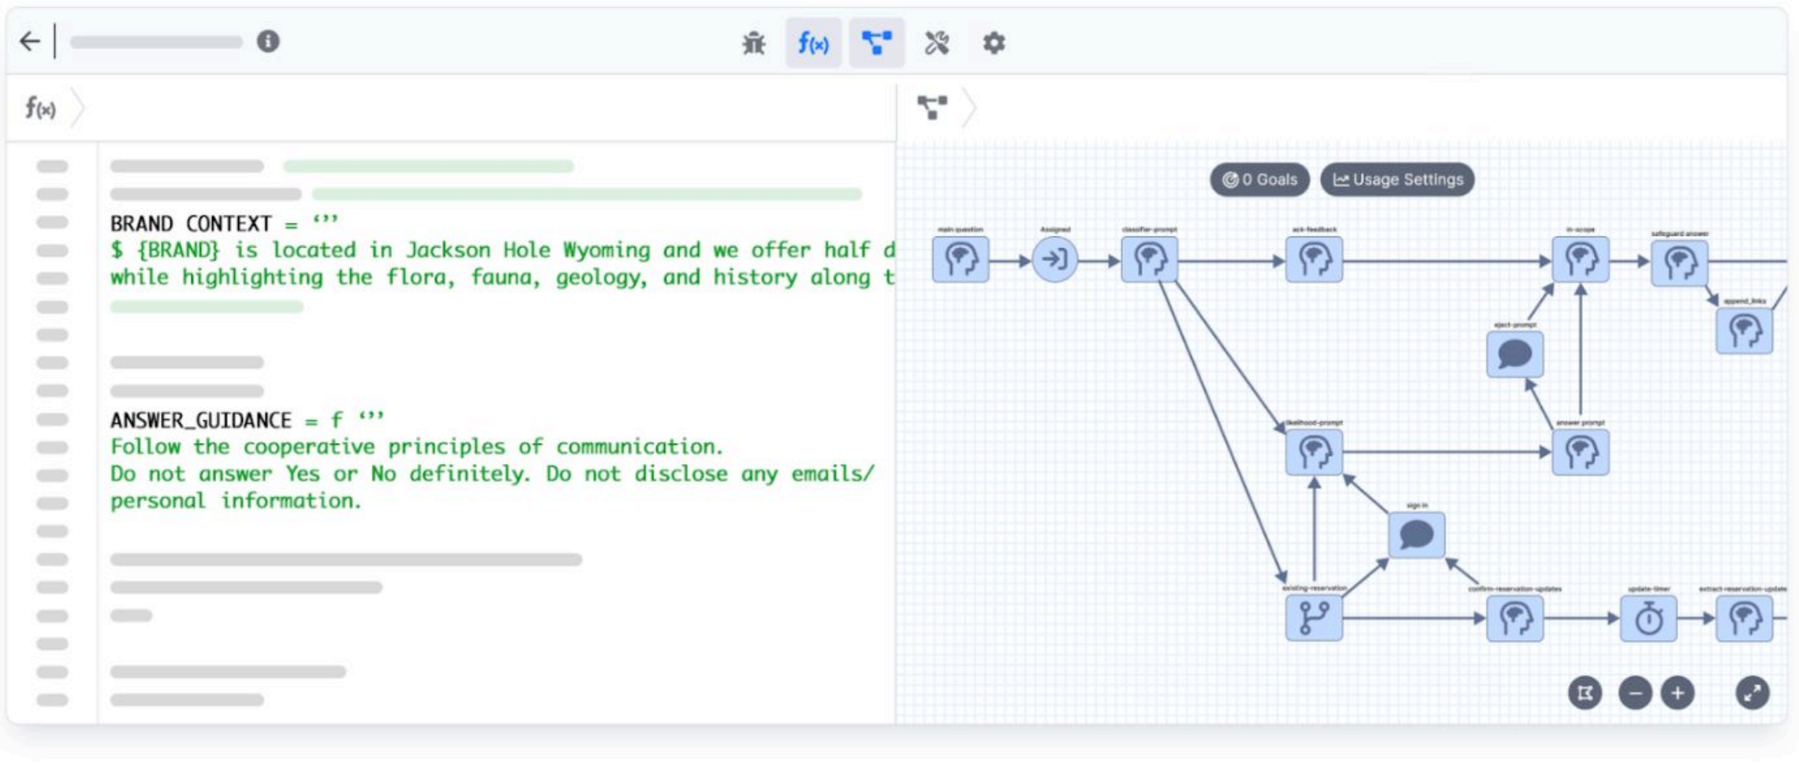

AI Model Flexibility.

Quiq takes a model-agnostic approach, allowing for integration with various AI models. AI Studio supports both frontier models supports both frontier models (such as OpenAI's GPT-3.5 and GPT-4, and Anthropic's Claude offerings) and our customers' AI models. This flexibility enables us to adapt to the specific needs of our customers and the evolving AI landscape.

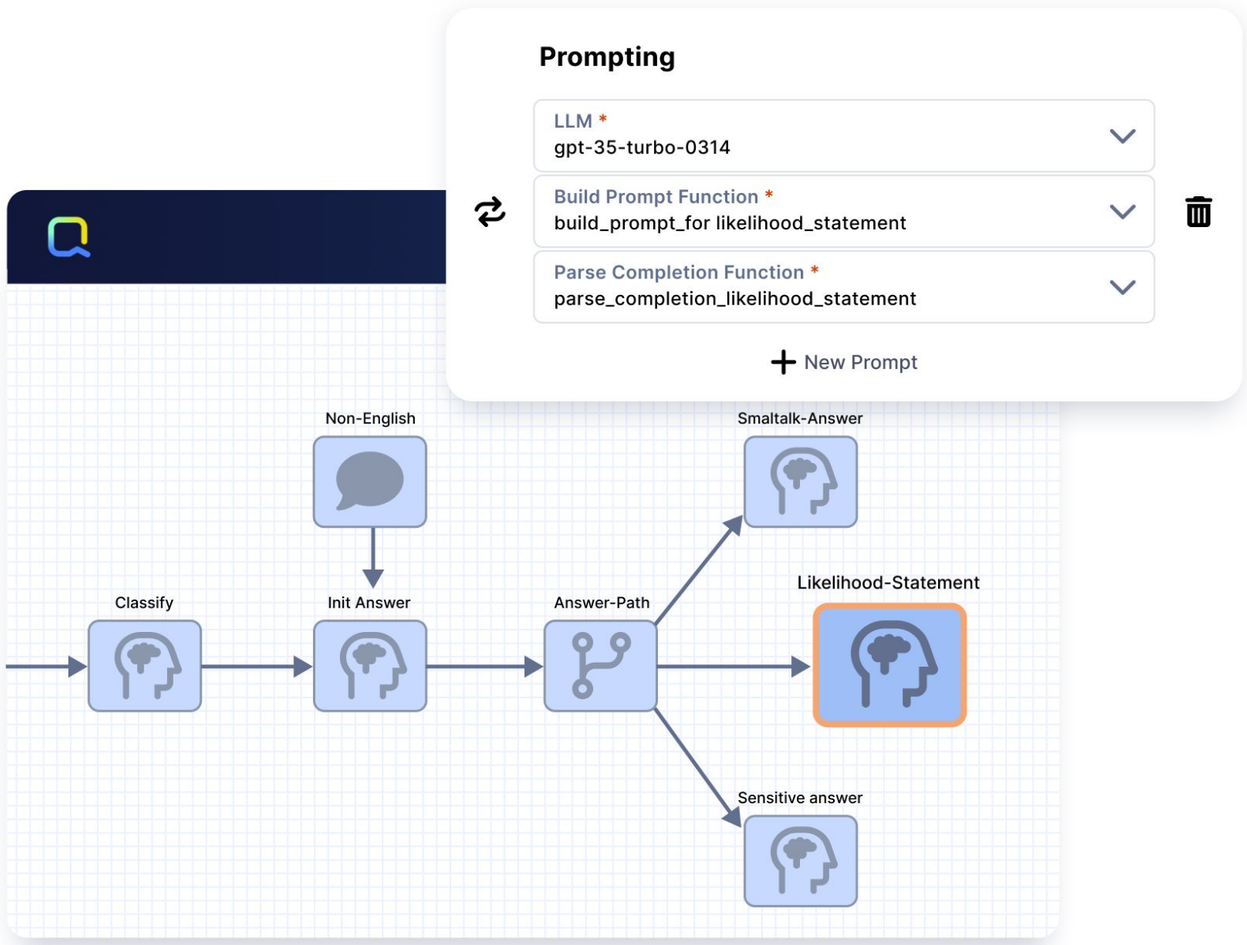

Conclusion.

At Quiq, we view the responsible use of artificial intelligence as crucial to our mission of enhancing customer experience. We are committed to accurately aligning to our customer's brand standards and maintaining transparency, privacy, and quality in our AI solutions. As AI technology continues to advance, we will continuously evaluate and evolve our approach to ensure we remain a trusted partner for our customers in their customer engagement efforts. By adhering to these responsible AI principles, we enable our clients to enhance their customer journeys, reduce operational costs, and significantly improve customer satisfaction scores. Our customer-centric AI solutions, built on these principles, empower businesses to seamlessly integrate AI and human roles in customer interactions, delivering optimal CX outcomes.

Want to see how Quiq works in real time?

We'd love to show you how our platform has successfully settled build vs. buy arguments for some of today's top brands. Let's talk!

*The information provided is for informational purposes only. Any usage of the AI solution and the performance of the AI solution will be subject to the terms of the agreement, including software warranty and SLA set forth in quiq.com/TOS or the agreement between Quiq and the customer. In addition, experiences that Quiq builds in the AI solution are built to the customer's specifications and can be modified by the customer. Quiq strives to minimize and remediate any vulnerabilities inherent to an AI solution as described in this document. Quiq does not guarantee any results from the AI solution and reference should be made by the customer to the terms and conditions of each third party technology LLM provider used in the AI solution. Terms and conditions applicable to third party AI models and LLMs are set forth on their websites.

Quiq

Customer-centric AI for CX

quiq.com