

Advancing Personalised Healthcare with OpenCHA



Large Language Models (LLMs) have become a cornerstone in digital transformation, especially in healthcare. However, their standalone deployment in conversational health agents (CHAs) has limitations. Conventional CHAs excel in natural dialogue but lack personalised data integration, multimodal input processing and complex problem-solving. Addressing these shortcomings, the openCHA framework introduces a modular and flexible architecture that enhances CHAs through data orchestration, intelligent planning and integration with diverse external sources. This approach not only boosts the performance of CHAs but also ensures they remain relevant, explainable and adaptive to individual health contexts.

Must Read: [Transforming Care with AI Voice Agents](#)

From Conversations to Intelligent Health Agents

Traditional LLM-based systems like ChatGPT or BioGPT have demonstrated proficiency in handling health-related conversations. Yet, their capabilities remain restricted to pre-trained textual outputs, lacking the means to access real-time patient data or coordinate complex health queries. openCHA transforms this paradigm by enabling CHAs to function not just as chatbots but as intelligent agents. By integrating with electronic health records, biosignals, nutritional databases and multilingual interfaces, openCHA-based CHAs can reason through queries, identify the necessary information and sequence appropriate actions.

Each query is processed through an Orchestrator comprising a Task Planner and Task Executor. These modules break down requests into steps, consult appropriate data sources or AI tools and refine results for final output. For example, when asked to assess a patient's sleep pattern, openCHA retrieves relevant data, performs statistical analysis and offers a tailored response. This process shifts CHAs from static dialogue generators to interactive decision support tools capable of adapting to user needs and contexts.

Demonstrated Capabilities Across Domains

The flexibility of openCHA has been proven through demonstrations and real-world use cases. One demo showed a CHA generating detailed reports from patient health records, calculating sleep trends and activity levels upon request. Another focused on stress estimation, processing heart rate variability from wearable-collected PPG signals using AI models to determine stress levels, even responding in multiple languages.

Beyond prototypes, openCHA has underpinned several high-performance agents. ChatDiet, for instance, merges personal health data with broader population models to provide dynamic and personalised food recommendations, outperforming baseline models like GPT4. In diabetes management, another CHA used domain-specific guidelines and data analysis tools to deliver accurate dietary advice, achieving a 92% accuracy rate. Mental health support was enhanced by a CHA that evaluated other chatbots' empathy and reliability with a lower error margin than leading LLMs. In physiological monitoring, an openCHA agent estimated heart rates with superior precision compared to GPT-4o, thanks to its ability to process real-time biosignals and conduct advanced analysis.

Benefits and Challenges of a Modular Framework

The core advantage of openCHA lies in its modular design. Developers can integrate healthcare-specific LLMs, select preferred planning techniques and link to custom external tools. This adaptability makes it suitable for varied clinical and research applications. Moreover, its transparency features enable users to query how conclusions were reached, improving trust in AI-generated recommendations.

Personalisation is another strength. CHAs developed with openCHA can tailor responses using longitudinal health data and behavioural patterns, ensuring greater relevance than generic LLM outputs. Reliability is also enhanced by delegating analytical tasks to validated external models while reserving LLMs for reasoning and communication.

However, these benefits come with trade-offs. Latency increases as tasks and data sources multiply, potentially hindering real-time interactions. Solutions like parallel execution and task fusion are being explored to reduce delays. Token limits in LLMs constrain the complexity of queries, although newer models with higher capacities are alleviating this issue. Security and compliance remain critical concerns, as openCHA does not govern deployment. Responsibility falls on implementers to ensure privacy standards are upheld, including data de-identification and access controls. Scalability, too, depends on user-side infrastructure and application-specific configurations.

openCHA offers a structured yet highly flexible approach to designing conversational agents that go beyond generic dialogue. By orchestrating multiple components—from data pipelines to AI models—openCHA enables CHAs to deliver personalised, multimodal and trustworthy health support. The framework's effectiveness across nutrition, diabetes care, mental health and physiological monitoring demonstrates its potential to reshape digital health services. As further enhancements address latency, robustness and privacy, openCHA is set to support scalable, adaptive and impactful health agents.

Source: [JAMIA Open](#)

Image Credit: [iStock](#)

Published on : Fri, 1 Aug 2025