Making Architectural Decisions Traceable in the Participatory Design of LLM Agents

Minseo Cho
Dyson School of Design Engineering
Imperial College London
London, United Kingdom
m.cho24@imperial.ac.uk

Pablo A Fonseca
Dyson School of Design Engineering
Imperial College London
London, United Kingdom
p.fonsecaarroyo@imperial.ac.uk

Martina Di Simplicio
Department of Brain Sciences
Imperial College London
London, United Kingdom
m.di-simplicio@imperial.ac.uk

Rafael A Calvo
Dyson School of Design Engineering
Imperial College London
London, United Kingdom
r.calvo@imperial.ac.uk

Participatory design enables HCI researchers to involve stakeholders in shaping interactive systems that reflect their values and contextual needs. In LLM-based conversational agents, however, user experience is shaped not only by visible dialogue or interface features, but also by behind-the-scenes elements such as underlying models and architectures. While participatory design is well-suited to eliciting descriptive user needs, the translation from qualitative design themes to internal agent architectures often remains implicit, distributed across teams, and difficult to trace. This lack of traceability makes it challenging to verify whether the resulting system remains aligned with stakeholder expectations, especially in high-stakes domains like mental health. To address this challenge, this paper explores how to make architectural decisions traceable in the participatory design of LLM-based agents. We present a reflective case situated in the design of an agent for mental health. Through this case, we articulate a translation process shaped by four forms of design judgment: making operations visible, identifying database needs, consolidating overlapping logics, and assigning architectural roles. We show that the same qualitative theme can yield different technical structures, each carrying distinct interactional and ethical consequences. By examining this interpretive translation layer, we illustrate through a reflective case how the movement from participatory insights to agent architecture can be made more traceable.

Conversational Agents, Participatory Design, Agent Architecture, Mental Health

1. INTRODUCTION

Participatory design has long provided Human-Computer Interaction (HCI) practitioners with ways to involve stakeholders in shaping technologies that affect them (Delgado et al., 2023). Through established strategies such as journey mapping, storyboards, personas, and scenario-based activities, participatory design helps surface user needs, expectations, and situated concerns. These methods are especially valuable in sensitive domains such as health and social care, where technologies must respond not only to functional requirements but also to user values and contextual complexity (Viswanathan et al., 2025).

However, the design object of participatory design is changing. While participatory findings have shaped user interface features, interaction flows, and service touchpoints, this mapping becomes more complex when designing Large Language Model (LLM)-based conversational agents (Dray & Symons, 2024). In these agentic systems, user experience is driven not just by interface elements, but also by underlying architectural decisions. For example, the experience of using a chatbot on WhatsApp will change radically by just changing the LLM provider used, even when the interface does not change. Structural choices regarding what an agent remembers, when it retrieves external data, how it reasons across turns, and when it acts autonomously all shapes how users experience the system (*What Are Agentic Workflows?*, 2025). In such systems, architecture is not merely a technical implementation layer; it becomes part of the interaction itself (Ghosh, 2025).

This means that translating participatory findings into LLM-based agents involves more than deciding what the agent should say or do at the interface. For example, a stakeholder need around personalisation may be realised through persistent long-term memory, user-editable profile, or adaptive prompting. Similarly, a desire for companionship may be translated into a dedicated agent role, a system-wide response style, or a safety policy governing relational responses. These are not neutral implementation choices. Each architectural interpretation carries assumptions about users, agents, and the relationships between them.

© Cho et al. Published by BCS Learning and Development Ltd.

Yet this translation from participatory findings to architectural decisions often remains implicit (Kilfoy et al., 2024). While participatory design is effective for eliciting stakeholder values, needs, and expectations, the connection between these qualitative findings and the internal structure of an agent is not self-evident. Participatory materials do not directly specify architectural components such as agent role, memory configuration, or tool access. Rather, in practice, the movement from workshop themes to agent roles and mechanisms may depend on developer interpretation. As a result, important design judgements can become hidden within technical implementation.

This lack of visibility creates a traceability gap. When architectural decisions are not transparent, it becomes difficult to understand why a particular decision was made, whose expectation it reflects, what alternatives were considered, and what assumptions were introduced during translation. If it is not clear how architectural choices link back to stakeholder expectations, participation risks becoming concentrated at the early stage of elicitation, while later architectural decisions remain opaque. This, in turn, may undermine the spirit of participatory design itself.

We therefore suggest a design challenge: how can the translation from participatory design findings to LLM-agent architecture be made traceable?

This paper makes three contributions. First, we identify an emerging HCI problem: while participatory design supports the elicitation of stakeholder values and expectations, there is limited guidance on how these findings can be translated into LLM-based agent architecture. Second, we present a reflective case study that illustrates the translation process from participatory workshop themes to the design of an agent for young people's mental health. Third, we discuss implications for future participatory design workshops, suggesting ways that workshops for LLM-based agents could support examination of architectural concerns such as memory, autonomy, tool use, and human oversight.

2. BACKGROUND

2.1 Participatory Design for Youth Mental Health Technologies

Participatory design is widely used in HCI to integrate stakeholder perspectives and situated knowledge throughout the development of interactive systems (LUMA Institute, 2012). They are important in health and wellbeing spaces, where unintended design flaws can lead to severe consequences (Göttgens & Oertelt-Prigione, 2021). By involving members of the care ecosystem, participatory design can surface contextual concerns, lived experiences, and support needs that may otherwise remain overlooked (Zuber & Moody, 2018).

In the design of digital health technologies for young people, participatory methods are suitable because their care needs and experiences differ from those of adults (Bevan Jones et al., 2020). Designing for this context requires sensitivity to specific developmental, social, emotional, and contextual factors to ensure that the resulting technology is appropriate (Hagen et al., 2012). Furthermore, as young people often exhibit lower acceptance and higher attrition rates with digital interventions, placing their voices at the centre of the design process can increase the acceptability of the technology and improve the likelihood that the resulting system produces its intended benefits.

Participatory design also offers structured pathways to discuss sensitive or abstract topics through the use of design artifacts (Hagen et al., 2012). In mental health contexts, activities such as scenario-building, storyboarding, persona creation, and journey mapping provide a concrete, accessible vocabulary. These artifacts act as communicative proxies, allowing participants to explore complex care mechanisms without requiring full personal self-disclosure or technical expertise.

Within conversational agent design, participatory methods help researchers go beyond identifying user needs to shape the role and interactional qualities of the agent (Sadek et al., 2024). Because participatory design sessions are interactive and creative, they assist in defining not just what support an agent should provide, but how it should engage. This includes determining dialogue pacing, appropriate tone, and boundaries between safe and risky responses. Hence, participatory design is uniquely suited for conversation design, which is difficult to conceptualise with static representations (Zytko et al., 2022).

However, while these outputs are effective for understanding ideal user-facing behaviours, they do not always make visible how such expectations may be realised by the engineering of an LLM system those behaviours in an LLM-based agent. do it in ways that explain how such expectations are affected by engineering of the LLM system. To bridge this gap, HCI needs ways to trace how participatory outcomes are interpreted into agent architecture, making visible the assumptions during this translation.

2.2 Agent Architecture as Design Material

LLM-based conversational agents differ from traditional rule-based chatbots. Rather than adhering to predefined dialogue paths, they can maintain context across multiple interactions, retrieve information from external sources, reason over user input, and generate flexible responses (Gupta et al., 2024). These capabilities mean that modern agent behaviour can no longer be reduced to a fixed script; instead, it is shaped by the architecture through which system operations are organised.

Architectural components, including memory structures, retrieval pipelines, reasoning flows, tool orchestration, and agent coordination, enable systems to support complex user needs (Schmid, 2025). In practice, they dictate what the agent appears to know, what it can access, when it can execute an action, and how responsibility is distributed between human and computational actors. Recent developments in LLM agents have therefore expanded the design scope beyond conversational content and interaction styles.

These architectural arrangements have interactional consequences. For instance, the specific design of a memory structure may influence the continuity and perceived personalisation of an interaction. Similarly, a single prompt may allow a conversation to feel fluid across shifting topics but introduce trade-offs around predictability and safety boundaries (Moore et al., 2023). These cases demonstrate that architecture shapes the user's lived experience of the agent.

In the context youth mental healthcare, architectural choices influence not only what forms of support are technically possible but also how that care is delivered. Even when users do not see the underlying architecture, they experience its systemic effects through the agent's behaviour.

Furthermore, the integration of artificial intelligence into healthcare introduces additional ethical and practical considerations, including transparency, trust, accountability, and data privacy (Auernhammer, 2020; Tahvanainen et al., 2024). These concerns cannot be addressed through interface design alone. They are deeply tied to architectural decisions, such as what information is stored, when the system acts autonomously, and how professional oversight is introduced (Sadek et al., 2024; Trocin et al., 2023). For this reason, agent architecture should be treated as a key behavioural design material that actively constructs the digital healthcare experience.

2.3 Interpretive Gap towards Agent Architecture

While participatory design is increasingly used to shape agent concepts, communication styles, and digital health tools, very little attention has been given to how these qualitative outputs are interpreted into the internal organization of LLM agent systems. Despite growing academic interest in co-designed conversational systems, relatively little work documents how stakeholder data maps onto agent configurations.

We identify this disconnect as an interpretive gap. Participatory workshops may produce rich qualitative materials, that are typically analysed into themes and design requirements. Yet, between these descriptive user needs and the resulting agent architecture lies an interpretive space where consequential design decisions are made. This translation work is not free from values. Each interpretation carries vastly different implications for how user values are structurally preserved. In other words, it involves subjective judgments about how stakeholder needs should be operationalised and, in digital health contexts, how care should be extended.

Still, the gap in the literature is understandable. LLMs are a relatively new technology, and agent architectures and their corresponding terminology are still evolving. As a result, there are few established conventions for documenting how qualitative findings are mapped onto architectural decisions. Without visibility into this translation process, it becomes difficult to understand why a particular architectural decision was made, whose expectation it reflects, what assumptions were introduced, and whether the resulting system remains aligned with the participatory findings.

This paper therefore focuses on making this interpretive translation explicit. Rather than offering a validated framework, we present a reflective case study to show how participatory design insights can be reasoned through as architectural decisions, and to argue that this reasoning should become a more visible concern in participatory HCI work.

3. REFLECTIVE TRANSLATION PROCESS

3.1 Participatory Design Workshop Outputs

Our reflective case study is drawn from the design of an LLM-based conversational agent for youth mental health. We focused on social prescribing, a holistic care pathway that connects individuals to non-clinical, community-based support, such as local activities, interest groups, or social services (Polley et al., 2017). To understand barriers and expectations, we conducted participatory design workshops with two stakeholder groups: young people as intended users, and link workers as professionals driving the care practice.

These workshops produced ten expectation themes. Young people focused on relational and exploratory needs (Personal Profiler, Goal Planner, Social Companion, Check-in Mate, Activity Connector), while link workers emphasised young people's own health accountability and administrative utility (Session Logger, Goal Tracker, Interest Bank, Resource Directory, Activity Recommender).

However, the themes did not by themselves specify how an agent should be organised. Translating these expectations into agent architecture needed a layer of interpretive work.

3.2 Translating Themes into Agent Architecture

In this reflective case, we organised the translation from participatory themes to agent architecture through four kinds of design judgement: making expected operations visible, identifying information and memory needs, consolidating overlapping expectations, and assigning architectural roles.

First, we made expected operations visible. Each expectation theme was examined in terms of how it might unfold in interaction. Rather than treating a theme as a static feature label, we mapped out how each theme would unfold sequentially: what interaction would initiate it, what information the user might provide, what the agent would need to do, what decisions might occur, and what output would be produced. This helped expose assumptions that were not obvious from the theme names alone. For example, *Goal Planner* implied not just a goal-setting feature, but a collaborative process where the agent helps a young person articulate implicit values, break them down into small steps, and revisit progress.

Second, we identified information and memory needs. Once the expected operations were made visible, we examined what information would be required to support them. We categorised whether information was short-term (provided dynamically within the current turn), long-term (persisting across multiple sessions), or external (retrieved from verified repositories). This led to the identification of several data structures, including user profiles, goal tracking logs, activity histories, and community resource indexes. These were not simply technical databases; they represented decisions about what aspects of the young person's mental health journey the system should be able to remember and use.

Third, we consolidated overlapping expectations. The ten themes did not map cleanly onto ten isolated technical components; many overlapped in behavioural logic or data dependencies. For example, *Goal Planner* and *Goal Tracker* both concerned goal-oriented support and relied on persistent progress data, leading us to consolidate them into a singular *Goal Manager* component. Personal Profiler and Interest Bank both involved capturing user attributes, which we merged into a *Contextual Profiler*. Finally, *Activity Connector* and *Activity Recommender* both addressed the discovery of community resources, so they were treated as a unified *Activity Connector* component.

Fourth, we assigned architectural roles. Following consolidation, we revisited each component according to the role it would play in the system. Components requiring sustained dialogue and open-ended reasoning were designated as sub-agents. Bounded tasks with clearer inputs and outputs were treated as tools. Back-end functions that supported the system without direct user interaction were classified as infrastructure or databases, while overarching behavioural qualities were treated as system policies.

Through this reflective translation, the participatory themes were interpreted into a layered architecture. This system is anchored by a reflective primary agent (Check-in Mate), supported by sub-agents (Goal Manager, Activity Connector), alongside contextual profiling tools, external database and shared resource directories.

3.3 Reflecting on Architectural Reasoning

The primary contribution of this case is not the final technical architecture itself, but the design reasoning that justified its construction. The following examples illustrate why this interpretive translation layer is a consequential site of design judgement (Table 1).

Table 1: Interpretive translation from participatory themes to architectural roles

Participatory Theme	Architectural Role	Design Justification and Logic
Check-in Mate	*Primary Agent*	*Interpreted as the primary agent because check-in and reflection are ongoing, open-ended dialogic activities rather than discrete tasks. It serves as the main conversational loop for other features to organically surface.*
Goal Planner & Goal Tracker	*Goal Manager (Sub-Agent)*	*Consolidated and modelled as a sub-agent because both require multi-turn reasoning and shared access to goal records. This role provides adaptive guidance around goal setting, monitoring, and adjustment.*
Activity Connector & Recommender	*Activity Connector (Sub-Agent)*	*Modelled as a sub-agent as it requires synthesising user preferences, contextual constraints, and available resources. This also allows oversight boundaries to be considered around how suggestions are presented.*
Personal Profiler & Interest Bank	*Contextual Profiler (Tool)*	*Designed as a tool rather than an autonomous agent. While personalisation is important, profiling does not need to drive dialogue; it runs in the background to extract and update user attributes.*
Social Companion	*Interaction Policy*	*Rather than creating "Companion Agent" (which would isolate friendliness as a separate feature), companionship was interpreted as a cross-cutting system policy. This embeds an empathetic, non-clinical tone across interactions.*
Session Logger	*Infrastructure*	*While link workers valued session logging for care continuity, and*

		reduced administrative burden, it required no direct conversation with the user. Therefore, it was positioned as background infrastructure.

These examples show that architectural translation involves design judgement. The same qualitative theme can be interpreted through multiple technical structures, each carrying different interactional and care-related implications. Making this reasoning visible enables researchers to inspect how stakeholder expectations were interpreted, transformed, or constrained during the agent design process.

4. DISCUSSION

4.1 Architecture as a Design Material

This reflective case suggests that system architecture should be treated as a participatory design concern in LLM-based agent systems. This does not mean that participants should design software architecture, or that co-design workshops should become technical engineering sessions. Rather, it means that design processes need to attend to the architectural consequences of stakeholder expectations.

In many HCI workflows, architecture is treated as an implementation step that follows established design requirements. In the context of LLM-based agents, however, this rigid separation between interaction design and architecture becomes difficult to maintain as architectural mechanisms shape how the system relates and responds to users, and broader contexts. These choices are therefore not only technical decisions, but socio-technical design decisions.

Consequently, the interpretive work occurring between qualitative themes and agent architectures needs to be made more visible. Without this visibility, stakeholder values may be translated in ways that are difficult to inspect. Making architectural decisions traceable can help HCI researchers and practitioners examine where, why, and how a stakeholder value was structurally embedded into the system.

4.2 Implications for Participatory Workshops

If participatory design is to meaningfully inform LLM-based agent's internal architecture, methodologies and prompts used should also adapt. Conventional activities typically focus on discussing user-facing feature sets, immediate needs, or ideal use cases. While these elements remain necessary, they may be insufficient for capturing unique affordances, dynamics, and risks of agent system architectures.

Future participatory design workshops could scaffold stakeholder engagement around five dimensions.

Memory and Persistence Instead of only exploring what an agent should know, designers should prompt participants to define structural memory boundaries. Workshops should investigate what information must persist across longitudinal sessions, what should remain temporary or session-bound, and what data users should have the right to review and edit.

Data Retrieval and Grounding Design workshops should explicitly map out information boundaries and source credibility. Participants can be prompted to evaluate: What external datasets or personal records is the agent permitted to query? When should the agent restrict itself to a verified directory rather than relying on open-ended language generation?

Autonomy and Operational Boundaries Rather than cataloguing what actions the agent can perform, workshops may investigate the triggers and limits of agentic behaviour. This requires questioning and surfacing boundaries around agent proactivity, when it is appropriate for the agent to take initiative, and when it should maintain a passive posture.

Tool Execution and Component Roles To make efficient use of technical components, designers should determine which design expectations require conversational agent and which can be offloaded to tools or infrastructure. Participants do not need to understand technical terms, but they can inform these architectural roles by mapping out use scenarios.

Human-in-the-Loop Oversight In sensitive care context like mental health, clients and professionals may hold varied expectations regarding human intervention. Workshops should surface these tensions to determine the extent to which a component should be automated, and when and how professional human oversight must be introduced.

In addition, meaningful oversight depends not only on where humans remain in the loop, but also on whether professionals can exercise effective judgement under real care conditions. Participatory design for agent systems should therefore also consider whether proposed forms of oversight are practically sustainable.

4.3 Limitations and Future Directions

This paper presents a single reflective case situated in youth mental health, rather than a validated methodology. While this domain is socially significant and useful for examining architectural concerns such as memory structure, care recommendation, and professional oversight, the translation processes described here may not generalise to other domains. The case also builds on a relatively small set of participatory design workshop sessions involving young people and healthcare providers; other stakeholders may have introduced different expectations and constraints regarding agent role.

The present study did not include a separate validation step in which workshop participants

reviewed the final architectural mappings, nor was the resulting conversational agent architecture evaluated as a deployed system. This paper therefore cannot claim that the resulting architecture improves user experience, care delivery, or health outcomes. Its contribution is instead to make visible the reasoning through which stakeholder expectations were interpreted into architectural decisions.

Finally, the present case offers design considerations rather than a validated participatory method. Future work could develop and evaluate concrete workshop materials that help stakeholders engage with architectural concerns such as memory, autonomy, tool use, and human oversight. Such work should also

distinguish the traceability of design reasoning from the evaluation of system behaviour: making an architectural decision traceable does not by itself demonstrate that an implemented LLM will consistently enact the stakeholder expectation it was intended to represent.

5. CONCLUSION

This paper highlighted an emerging HCI design challenge for participatory design approach: in LLM-based agents, user experience is shaped not only by visible interaction, but also by architectural decisions. While participatory design is well suited to eliciting stakeholder values, and requirements engineering can help specify system functions and constraints, the translation between qualitative expectations and agentic architectural choices often remains implicit.

Through a reflective case study in mental health, we illustrated how stakeholder themes can be interpreted into architectural decisions. The case showed that themes do not map directly onto components. Instead, they require design judgement. In our case, self-reflection concept became the primary agent, goal-related expectations became a sub-agent, companionship became an interaction policy, and resource information became a curated repository. These decisions were not purely technical; they actively shaped the user experience and safe and accountable technology use within care practices.

We therefore propose a design space for HCI: making architectural decisions traceable in the participatory design of LLM agents. As agentic systems become more common in sensitive domains, HCI researchers and practitioners will need ways to show how stakeholder values are translated into the structures that shape system behaviour. Making this translation visible can support more accountable design practice by enabling researchers and practitioners to inspect how stakeholder expectations become embedded in architectural decisions.

6. REFERENCES

Auernhammer, J. (2020, August 11). *Human-centered AI: The role of Human-centered Design Research in the development of AI*. DRS2020: Synergy. https://doi.org/10.21606/drs.2020.282

Bevan Jones, R., Stallard, P., Agha, S. S., Rice, S., Werner-Seidler, A., Stasiak, K., Kahn, J., Simpson, S. A., Alvarez-Jimenez, M., Rice, F., Evans, R., & Merry, S. (2020). Practitioner review: Co-design of digital mental health technologies with children and young people. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, *61*(8), 928–940. https://doi.org/10.1111/jcpp.13258

Delgado, F., Yang, S., Madaio, M., & Yang, Q. (2023). The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. *Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–23. https://doi.org/10.1145/3617694.3623261

Dray, J., & Symons, D. (2024). Review of Innovative Mental Health Support for Children and Young People: Generative AI Co-design Applications and Challenges. *Current Developmental Disorders Reports*, *12*(1). Scopus. https://doi.org/10.1007/s40474-025-00328-z

Ghosh, B. (2025, October 26). Agentic Design Patterns. *Medium*. https://medium.com/@bijit211987/agentic-design-patterns-cbd0aae2962f

Göttgens, I., & Oertelt-Prigione, S. (2021). The Application of Human-Centered Design Approaches in Health Research and Innovation: A Narrative Review of Current Practices. *JMIR mHealth and uHealth*, *9*(12), e28102. https://doi.org/10.2196/28102

Gupta, D., Swami, V., Shukla, D., & Nimala, K. (2024). Design and Implementation of an AI-Driven Mental Health Chatbot: A Generative AI Model. *2024 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES)*, 1–7. https://doi.org/10.1109/ICSES63760.2024.10910363

Hagen, P., Collin, P., Metcalf, A., Nicholas, M., Rahilly, K., & Swainston, N. (2012). *Participatory design of evidence-based online youth mental health promotion, intervention and treatment*. Young and Well Cooperative Research Centre. http://handle.uws.edu.au:8081/1959.7/534977

Kilfoy, A., Hsu, T.-C. C., Stockton-Powdrell, C., Whelan, P., Chu, C. H., & Jibb, L. (2024). An umbrella review on how digital health intervention co-design is conducted and described. *Npj Digital Medicine*, *7*(1), 374. https://doi.org/10.1038/s41746-024-01385-1

LUMA Institute (Ed.). (2012). *Innovating for people: Handbook of human-centered design methods* (First edition). LUMA Institute, LLC.

Moore, R. J., An, S., & Ren, G.-J. (2023). The IBM natural conversation framework: A new paradigm for conversational UX design. *Human–Computer Interaction*, *38*(3–4), 168–193. https://doi.org/10.1080/07370024.2022.2081571

Polley, M. J., Fleming, J., Anfilogoff, T., & Carpenter, A. (2017, August 18). *Making Sense of Social Prescribing* [Technical report]. University of Westminster. https://www.westminster.ac.uk/patient-outcomes-in-health-research-group/projects/social-prescribing-network

Sadek, M., Calvo, R. A., & Mougenot, C. (2024). The Value-Sensitive Conversational Agent Co-Design Framework. *International Journal of Human–Computer Interaction*, 1–32. https://doi.org/10.1080/10447318.2024.2426737

Schmid, P. (2025, May 5). *Zero to One: Learning Agentic Patterns*. https://www.philschmid.de/agentic-pattern

Tahvanainen, L., Tetri, B., & Ahonen, O. (2024). Exploring and Extending Human-Centered Design to Develop AI-Enabled Wellbeing Technology in Healthcare. In M. Särestöniemi, P. Keikhosrokiani, D. Singh, E. Harjula, A. Tiulpin, M. Jansson, M. Isomursu, M. van Gils, S. Saarakkala, & J. Reponen (Eds), *Digital Health and Wireless Solutions* (pp. 288–306). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-59091-7_19

Trocin, C., Mikalef, P., Papamitsiou, Z., & Conboy, K. (2023). Responsible AI for Digital Health: A Synthesis and a Research Agenda. *Information Systems Frontiers*, *25*(6), 2139–2157. https://doi.org/10.1007/s10796-021-10146-4

Viswanathan, S., Ibrahim, S., Shankar, R., Binns, R., Van Kleek, M., & Slovak, P. (2025). The Interaction Layer: An Exploration for Co-Designing User-LLM Interactions in Parental Wellbeing Support Systems. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, CHI '25*. https://doi.org/10.1145/3706598.3714088

What Are Agentic Workflows? Patterns, Memory, Use Cases, and Examples | Weaviate. (2025, March 6). https://weaviate.io/blog/what-are-agentic-workflows

Zuber, C. D., & Moody, L. (2018). Creativity and Innovation in Health Care: Tapping Into Organizational Enablers Through Human-Centered Design. *Nursing Administration Quarterly*, *42*(1), 62–75. https://doi.org/10.1097/NAQ.0000000000000267

Zytko, D., J. Wisniewski, P., Guha, S., P. S. Baumer, E., & Lee, M. K. (2022). Participatory Design of AI Systems: Opportunities and Challenges Across Diverse Users, Relationships, and Application Domains. *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–4. https://doi.org/10.1145/3491101.3516506