

Viewpoint

From Prediction to Interaction: A Behavioral Informatics Framework for AI in Preventive Health

Anand Gavai, Dr IR; Jaap Heringa, Prof Dr IR

Bioinformatics Section, Department of Computer Science, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Corresponding Author:

Anand Gavai, Dr IR
Bioinformatics Section, Department of Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081A
Amsterdam 1081 HV
The Netherlands
Phone: 0031 6 34 64 99 52
Email: anand.gavai@gmail.com

Abstract

In this viewpoint, we present our perspective—informed by published evidence and by our own work on AI-driven digital health systems—on why AI has had limited real-world impact in preventive health and how it should be redesigned. Most AI systems in health care are built as episodic prediction tools, whereas preventive care depends on sustained behavioral change; the consequences are visible in consistently high attrition and low long-term adherence across digital health interventions. The aim of this viewpoint is to argue that this failure is structural rather than incremental and to propose an integrated, interaction-centric alternative. We advance 3 key messages. First, the limited impact of AI in preventive health reflects a structural misalignment between episodic, prediction-centric system design and the longitudinal dynamics of health behavior. Second, addressing this misalignment requires treating continuous interaction as the primary unit of design, which we articulate through a behavioral informatics framework integrating 6 interdependent layers: data, knowledge, reasoning, interaction, behavior, and governance. The contribution of this framework is explicitly integrative: it synthesizes previously separate strands, including just-in-time adaptive interventions, behavioral intervention technologies, and human-AI interaction research, into a single system-level account. Third, evaluation must shift from predictive accuracy and short-term engagement toward measurable behavioral trajectories and longitudinal outcomes, with explicit attention to feasibility and to conditions under which simpler, non-AI interventions remain preferable. This viewpoint is addressed to digital health researchers and AI developers, behavioral scientists, clinicians involved in preventive care, and policymakers responsible for governing continuously interacting AI systems.

J Med Internet Res 2026;28:e98961; doi: [10.2196/98961](https://doi.org/10.2196/98961)

Keywords: artificial intelligence; AI; agentic AI; large language models; digital health; human-AI interaction; preventive health; behavioral informatics; digital behavior change

Introduction

Episodic AI in a Longitudinal Domain

AI in health care has traditionally been structured around prediction, where systems transform input data into outputs such as risk scores, classifications, or recommendations. These model-centric approaches have demonstrated strong performance in domains such as diagnostics and decision support [1,2]. However, their impact in preventive health remains limited, because prevention depends not on isolated decisions but on sustained behavioral adaptation, where

outcomes emerge from sequences of small, context-dependent actions repeated in everyday life [3,4].

The empirical record quantifies this gap. Eysenbach's "law of attrition" identified steep nonusage attrition as a defining feature of eHealth interventions [5]. Subsequent studies found that only about 50% of participants fully adhered to web-based interventions [6], while app-based chronic disease interventions had a pooled dropout rate of 43% (95% CI 29%-57%), rising to 49% in observational studies [7]. Across popular mental health apps, median retention was 3.9% at 15 days and 3.3% at 30 days [8]. Chatbot interventions have improved physical activity, diet, and sleep [9], but long-term

maintenance remains uncertain. These findings, rather than predictive accuracy alone, motivate this viewpoint.

Generative AI and large language models (LLMs) enable personalized recommendations through conversational interfaces [10], but most implementations remain episodic: they respond to user input and terminate. Even conversational or agentic systems may not maintain longitudinal context or support sustained behavioral change [11,12]. Their implicit assumption—that accurate information is sufficient—is difficult to reconcile with persistent attrition. Information that is not engaged with cannot change behavior, and engagement is a central determinant of intervention effectiveness [13,14]. Episodic AI is therefore misaligned with the temporal dynamics of preventive health behavior.

Definitions and Scope

Because several terms are used throughout this viewpoint, we define them at the outset. Interaction-centric AI denotes systems in which the primary unit of design and evaluation is the continuous loop of exchanges between system and user—insight, action, feedback, and adaptation—rather than the accuracy of individual outputs. A behavioral informatics framework is an integrated account that couples the informatics components of such systems (data, knowledge, and reasoning) with their behavioral components (interaction design and behavior change mechanisms) under explicit governance constraints. Agentic coordination refers to reasoning organized as the orchestration of multiple specialized components across user state, context, and knowledge, rather than autonomous decision-making by a single model [15].

We use preventive health in a deliberately broad sense that includes both primary prevention (interventions before disease onset) and secondary prevention or early risk management, such as lifestyle modification for individuals with elevated but subclinical risk factors, as in the prevention of hypertension, type 2 diabetes, and obesity. What unites these settings, and distinguishes them from acute or episodic care, is that outcomes depend on behavior sustained over months to years rather than on discrete clinical decisions.

Aim, Contribution, and Roadmap

The aim of this viewpoint is to explain why episodic, prediction-centric AI is structurally unsuited to preventive health and to propose an integrated framework for designing and evaluating interaction-centric alternatives. Rather than restate the 3 key messages given in the abstract, we note where each is developed: the first in this Introduction and the following section, the second in the framework section and the illustrative scenario, and the third in the section on evaluation and feasibility.

Two additional points clarify the contribution and readership. Interaction-centric AI has been proposed previously, including in our work on agentic coordination [15] and human-AI interaction research [11,16]. This viewpoint contributes an integrative synthesis showing how behavioral, computational, and governance concerns couple. The framework sections primarily address AI

developers and behavioral scientists, the feasibility subsection addresses clinicians and preventive care organizations, and the governance subsection addresses policymakers and regulators.

From Episodic Prediction to Continuous Behavioral Support

The Interaction Loop

Traditional health care AI uses linear pipelines that generate predictions or recommendations for users or clinicians to interpret and act upon [1,2]. This suits discrete decisions but not behaviors such as activity, diet, and sleep, which vary with routines, context, and motivation. A recommendation has little impact if it is not acted upon and reinforced; repeated context-specific actions can instead accumulate into meaningful outcomes.

Interaction-centric systems embed outputs in iterative cycles of insight, action, feedback, and adaptation. Insights are coupled with actionable guidance, behavioral responses are observed, and subsequent recommendations adapt. This requires longitudinal context, adaptation to changing user states, and calibrated interaction frequency, while introducing risks including fatigue and reduced autonomy.

Positioning Relative to Existing Frameworks

Three research traditions have addressed parts of this problem, and our framework builds directly on them. Just-in-time adaptive interventions (JITAs) formalize when to intervene: decision points, tailoring variables, and decision rules that deliver support at moments of receptivity [17, 18]. The behavioral intervention technology (BIT) model formalizes what to deploy, organizing design around 5 questions: why the intervention exists (the clinical aim: eg, blood pressure reduction), what strategy serves it (eg, self-monitoring), how that strategy is instantiated conceptually (eg, a reminder), how it is delivered technically (eg, a push notification), and when delivery occurs (the workflow rule: eg, every evening at 8 PM) [19]. Human-AI interaction research formalizes how systems should communicate: guidelines for transparency, feedback, and user control at the interface [11,16].

The comparison is complementary. JITAs use predefined decision points and rules, whereas interaction-centric systems may repeatedly derive decision points, strategies, and framing from evolving user state. BIT can also specify adaptive workflows in which delivery changes with predefined user or environmental variables. Here, adaptation is a recurring cross-layer operation that may revise decision points, behavioral strategy, framing, and delivery. For example, after observing that an 8 PM weekday reminder is ignored but weekend meal planning is engaged with, the system may shift to Saturday-morning action planning. The distinction is not static versus adaptive intervention, but whether runtime reconfiguration and its dependencies across data, knowledge, reasoning, behavior, interaction, and

governance are central architectural processes. Governance also constrains each operating cycle rather than remaining solely external compliance. Existing traditions isolate timing, intervention content, or interface, yet failures often arise at their couplings: timing cannot rescue poor framing, framing cannot rescue context-insensitive advice, and failures remain invisible without behavioral feedback. A JITAI can therefore be one framework configuration with a fixed-rule reasoning layer, while BIT is a design specification that can be instantiated within the framework; the present architecture makes the cross-layer processes governing runtime adaptation explicit.

A Behavioral Informatics Framework for Interaction-Centric AI


Overview and Operating Cycle

We present a conceptual architecture of 6 interdependent layers—data, knowledge, reasoning, interaction, behavior, and governance—

and governance (Figure 1)—that clarifies dependencies and design constraints rather than prescribing implementation. Continuous feedback couples the layers, and performance emerges from their interaction. Membership is determined by what each layer owns; explicit rules distinguish data from knowledge and interaction from behavior. All layers depend on a time-varying representation of user state, context, and behavioral trajectories.

A loop traversal proceeds as follows: (1) the data layer updates user state; (2) reasoning detects a decision point; (3) knowledge grounds candidate recommendations; (4) reasoning acts, defers, or remains silent and selects a strategy within governance constraints; (5) interaction chooses timing, framing, and modality; (6) behavior records the response and updates strategy estimates; and (7) the response re-enters the data layer. The following subsections describe each layer and its integration with care ecosystems.

Figure 1. Conceptual architecture of an interaction-centric AI system for preventive health. Six interdependent layers connect continuous sensing, knowledge representation, reasoning, adaptive communication, behavioral feedback, and governance.



Data Layer: Continuous and Contextual Inputs

The data layer owns what is true of this person—time-indexed observations and derived state—whereas the knowledge layer owns person-independent evidence. The distinction is not raw versus structured or sensed versus retrieved: a graph of the user's history remains data, while a guideline threshold remains knowledge even when presented as retrieved text. The layers differ in update trigger (user action versus evidence change), validity condition (measurement reliability versus evidential currency), governance obligation (data minimization versus provenance), and failure mode (missensed steps versus outdated guidance). The reasoning layer joins them during decision formation. Data streams may include self-reports, sensors, and context. Although wearables and apps already collect many such data, they often lack integration with adaptive reasoning and behavioral feedback. Because longitudinal data are incomplete and noisy, this layer must construct a temporally coherent user state representation through signal integration, missing data handling, temporal smoothing, imputation, and adaptive sampling [20]. Collection must also limit burden and intrusion, which undermine engagement [17].

Knowledge Layer: Hybrid Representation

Reliable recommendations require both retrieved sources such as clinical guidelines and structured representations such as knowledge graphs [12,21]. Retrieval accommodates evolving evidence; structure supports relational reasoning. Hybrid systems must resolve incomplete, outdated, or conflicting information [22] and maintain coherence with current evidence and prior behavior [23], escalating unresolved high-stakes conflicts to human oversight. Because sources are continuously updated and partly inconsistent, and LLM consistency over time remains contested [24], knowledge integration is an epistemic limitation as well as an engineering task.

Reasoning Layer: Agentic Coordination

The reasoning layer integrates signals from the data and knowledge layers to generate context-aware, actionable outputs. Rather than relying on a single monolithic model, reasoning can be organized as coordination across specialized components for behavioral interpretation, contextual analysis, and knowledge grounding. In this framing, agentic AI means coordination across knowledge, context, and behavior, not autonomous decision-making, aligning with perspectives on AI as an orchestrator of distributed reasoning [15,25], and raising deliberate questions about how much autonomy is achievable or desirable in safety-critical health contexts.

Three functions define this layer. Personalization requires combining rule-based constraints, which guarantee safety and alignment with established evidence, with adaptive mechanisms responsive to individual trajectories, avoiding both rigidity and uncontrolled variability [26]. Temporal orchestration requires deciding not only what to recommend but when to intervene and when to remain passive [27]: excessive

interaction produces fatigue, and insufficient interaction forfeits effectiveness [28]. Uncertainty handling requires representing and communicating confidence honestly because overconfidence undermines trust while excessive caution reduces usability [23]; situations exceeding the system's competence must be identified and escalated to clinical care.

Interaction Layer: Adaptive Communication

The interaction layer governs how decisions are communicated. Its scope is deliberately bounded: this layer owns the communication channel, timing, framing, modality, and dose of exchanges, whereas the selection of behavioral strategies and the monitoring of behavioral response belong to the behavior layer described next. The primary end user throughout this framework is the individual whose behavior is being supported, not the clinician; clinicians interact with these systems as secondary users through oversight and escalation pathways, where design requirements differ substantially (aggregated trends and provenance rather than progressive disclosure).

Communication design follows the principle of cognitive economy: each exchange delivers minimal, actionable information to limit cognitive load [29], with progressive disclosure offering concise guidance first and detail on request. Interactions must be nonintrusive; deferred under high cognitive load; prioritized when action is feasible [18]; and adapted to health literacy, preferences, and communication styles [10]. Such adaptivity carries a documented cost: simplified, tailored outputs can obscure the full basis of a recommendation, and LLM-generated guidance has been shown to deviate from clinical practice guidelines when operating on real-world, unstructured inputs [30], which is why explanation obligations and evidence grounding are assigned to this layer by the governance layer below.

Behavior Layer: Behavior Change Mechanisms and Adaptation

Where the interaction layer determines how the system speaks, the behavior layer determines what the system is trying to change and whether change is occurring. It selects and adapts behavior change techniques—goal setting, self-monitoring, prompts and cues, feedback on performance, and reinforcement—in response to user state and context, consistent with established models of capability, opportunity, and motivation [3]. Goal setting and self-monitoring belong to this layer because they define and track behavioral targets; the interaction layer only renders them. The BIT model conceptualizes digital interventions as a mapping between clinical aims, behavioral strategies, and technological instantiation [19]; interaction-centric systems differ in continuously reconfiguring these mappings in response to observed behavior, which brings flexibility but also potential instability.

A taxonomy of techniques does not by itself address the dependency and autonomy concerns this viewpoint raises. Self-determination theory supplies the distinction that does: behavior may be regulated autonomously, through goals the

person endorses, or through controlled regulation driven by external contingency, and autonomous regulation is generally more strongly associated with persistence once external support is withdrawn than controlled regulation [31]; meta-analytic evidence in health contexts links autonomy-supportive practice to psychological need satisfaction, autonomous motivation, and better health outcomes [32]. The architectural consequence is direct: the behavior layer's objective is not adherence but the internalization of regulation, so a system that raises adherence while shifting regulation from autonomous to controlled has failed on its own terms. Persistence during and after tapering may therefore provide one observable indication that regulation has become more internalized.

Motivational interviewing supplies the corresponding communicative technology and strategies that are autonomy-supportive by construction: eliciting the person's own reasons for change rather than supplying them, responding to ambivalence without confrontation, and keeping the decision to act the user's [33]. Technology-delivered adaptations have been applied across chronic disease prevention and management, though reviews report wide variation in implementation fidelity [34], and systems delivering motivational interviewing through conversational AI are only now being characterized, with rigorous and standardized fidelity assessment remaining a major open problem [35]. The behavior layer's strategy set should therefore be drawn from techniques compatible with autonomous regulation, and fidelity to them should be monitored rather than assumed.

This layer must detect disengagement, identify barriers to action, and adapt accordingly, modifying recommendations, adjusting interaction frequency, or reframing guidance. A central tension is maintaining engagement without fostering dependency: frequent prompts may raise short-term adherence while creating reliance on the system. Tapering strategies include prompt modulation, gradual reduction of intervention intensity, and transition to maintenance phases as behavioral patterns stabilize; these strategies support the shift from guided interaction to autonomous behavior. Interventions optimized for adherence can shade into subtle behavioral influence, a risk taken up under governance.

This locates a boundary the viewpoint has so far only named: when does behavioral support become manipulation? Persuasive technology research defines persuasion as influence that changes attitudes or behavior without coercion or deception, and catalogs the system features through which it operates [36,37], but it assumes a strategy fixed at design time and in principle inspectable; an interaction-centric system reconfigures its strategy at run time, so what the user is subject to is neither fixed nor visible. Accounts of online manipulation supply the criterion that survives this shift: influence is manipulative when it exploits cognitive or affective vulnerabilities and depends on the target not

recognizing it [38]. A substantial reduction in an adaptation's effectiveness when its purpose and mechanism are disclosed may therefore serve as a warning sign of manipulative influence, although disclosure sensitivity alone is not sufficient to establish manipulation. The behavior layer should therefore record which strategy is active and why, while the interaction and governance layers should make that strategy inspectable and contestable by the user. Users should be able to request an explanation of the strategy's purpose and operation.

Governance Layer: Constraints and Accountability

Because these systems continuously influence behavior, governance is an architectural layer and design driver. The General Data Protection Regulation and the European Union Artificial Intelligence Act principles—data minimization, purpose limitation, transparency, and human oversight—must be operationalized [39,40]. Adaptive sampling and on-device processing can minimize collection; behavioral data must not silently become training data for other purposes. The knowledge layer requires provenance and conflict resolution audit trails; the interaction layer requires explanations at recommendation time [41]; reasoning requires escalation thresholds; and the behavior layer requires user-controlled prompt limits, inspectable and contestable goals, and penalty-free disengagement.

Equity is a governance concern of the same rank: sustained engagement is achieved more readily among digitally literate and motivated users, so without explicit safeguards—inclusive communication design, multimodal interaction, attention to access constraints—interaction-centric systems risk amplifying existing health inequities. Finally, governance must address system-level risks of prolonged interaction, including bias, drift, and behavioral influence that is difficult to monitor when personalization is opaque; continuous auditing and monitoring are required to keep systems aligned with user needs and societal expectations [15]. Choices about data collection, personalization, and interaction strategy are normative: they embed assumptions about acceptable influence, autonomy, and risk.

Ecosystem Integration

Interaction-centric systems must integrate with electronic health records, wearables, and other applications while addressing interoperability and cross-system governance [42]. Hybrid human-AI pathways should enable clinicians to review outputs, monitor trends, and intervene beyond automated scope [11,16]. Deployment must cover onboarding, active engagement, maintenance, scalability, robustness, and behavioral or performance drift [43]. [Table 1](#) summarizes the layers, challenges, and evaluation metrics.

Table 1. Core layers of interaction-centric AI systems and associated design considerations and evaluation metrics^a.

Layer	Primary function	Example capabilities	Key design challenges	Evaluation metrics
Data	Continuous and contextual input	Behavior tracking, sensor integration, and user-reported data	Missing data, noise, temporal drift, and user burden	Data completeness, signal reliability, sampling frequency adherence, and user input compliance
Knowledge	Context grounding and reasoning support	Hybrid representations (retrieval and structured knowledge)	Inconsistency, outdated information, and conflict resolution	Retrieval relevance, factual consistency, knowledge coverage, and update latency
Reasoning	Coordination across inputs and knowledge	Personalized recommendations and adaptive decision logic	Coherence, uncertainty handling, and control vs flexibility	Recommendation consistency, uncertainty calibration ^b , decision accuracy (contextual), and adaptation effectiveness
Interaction	Communication with the user	Adaptive prompts, progressive disclosure, and multimodal interaction	Cognitive load, engagement vs fatigue, and usability variability	Interaction quality (user ratings), response time, drop-off rate, and cognitive load indicators ^b
Behavior	Feedback and behavioral adaptation	Behavioral loops, reinforcement strategies, and habit formation support	Sustained adherence, behavioral variability, and disengagement	Behavioral adherence, action completion rate, habit persistence, and longitudinal engagement
Governance	Ethical, regulatory, and safety constraints	Transparency, explainability, user control, and compliance mechanisms	Trust, bias, regulatory compliance, and accountability	Explainability scores ^b , user trust measures, audit outcomes, bias metrics, and compliance adherence

^aThe evaluation metrics presented in this table, including behavioral adherence, action completion rate, and drop-off rate, can be measured using existing instruments, except for those identified by superscript b.
^bUncertainty calibration, cognitive load indicators, and explainability scores require further operationalization and context-specific validation for longitudinal interaction-centric AI systems.

The structural coherence of the framework does not guarantee practical viability. Interaction-centric systems presuppose sustained engagement, reliable behavioral signals, and consistent adaptation—conditions rarely met in real-world settings, as the attrition evidence cited above makes clear. There is a real risk that such systems reproduce the failures of earlier digital health interventions under a new architectural label, particularly if behavioral loops degrade into intermittent or superficial interaction. We treat this as a testable failure mode rather than a footnote, and return to it under feasibility below.

The Framework in Operation: An Illustrative Scenario

To demonstrate how the layers interact in practice, consider an individual identified through routine screening as having elevated blood pressure, managed initially through lifestyle modification without pharmacological treatment. This is a secondary prevention setting—early risk management before established disease—and thus falls within the scope of preventive health defined in the Introduction. The scenario is a speculative design walkthrough, not a validated result; we annotate each step with the layer it exercises.

The system identifies dietary, activity, and routine patterns from wearable signals and meal logging (data). Elevated evening sodium intake combined with prolonged sedentary periods creates a decision point (reasoning), grounded in guidelines and user context (knowledge). At meal planning, it suggests a low-sodium alternative and prompts a brief walk (interaction). Logging and sensor data show whether the user acts (behavior); the system reinforces effective patterns or adjusts timing and framing, such as moving suggestions to weekend meal preparation. Governance caps interaction

frequency and data sampling, while sustained blood pressure elevation triggers clinical escalation.

Over successive loop traversals, the system tapers from frequent intervention to an adaptive monitoring role, providing targeted input only when deviations from stable patterns are detected. Each element of this walkthrough corresponds to one step of the operating cycle described above; the scenario is the cycle instantiated in a concrete preventive context.

Evaluation, Feasibility, and Implications

From Predictive Accuracy to Measurable Behavioral Trajectories

The shift to interaction-centric AI requires rethinking evaluation. We propose 3 dimensions, each with measurable indicators. Interaction quality covers the relevance, timing, and usability of system outputs: measurable now through prompt response rates, time-to-response, conversation or session drop-off rates, and validated usability and engagement instruments [14,43]. Behavioral adherence covers whether interactions translate into action: measurable as the proportion of recommended actions completed within defined time windows (sensor-verified where possible), consistency of targeted routines, and time-to-disengagement, analyzed with survival methods. Longitudinal impact covers health-relevant change: trajectories of risk factors (eg, blood pressure, weight, and glycemic markers) and validated behavioral instruments, analyzed as slopes and inflection points in mixed-effects or growth-curve models rather than as single end point differences. Additional longitudinal indicators

include time to habit formation, persistence of change after intervention tapering, and relapse rates.

Evaluation data arise largely from interaction traces, sensors, and periodic validated self-report. Individual adaptation complicates between-arm comparisons and supports within-person randomization. Microrandomized trials randomize intervention options at decision points, estimating proximal effects where the framework locates its mechanism [44]. Unlike standard person-level designs, their estimand is the effect of interaction policies rather than a fixed package. Conventional randomized controlled trials remain necessary for distal health outcomes, and AI-supported trial design is emerging [45].

Evaluation dimensions interact, and not always benignly. Increased adherence may come at the expense of autonomy, and optimized interaction frequency may accelerate fatigue. Adherence itself is not neutral: it presupposes goals that may not reflect user priorities, and systems optimized for adherence risk shifting from support toward subtle control, particularly when adaptation mechanisms are opaque or difficult to contest. Evaluation of interaction-centric systems is therefore normative as well as technical, and the boundary between support and manipulation must be an explicit evaluation criterion rather than an afterthought. Because evaluation data are continuously produced, they should feed back into system operation—a closed-loop process in which evaluation contributes directly to system evolution.

Feasibility, Resources, and When Simpler Interventions Suffice

Feasibility remains uncertain. Persistent attrition [5,7,8] means no architecture guarantees continued engagement. Continuous sensing burdens devices and users; LLM reasoning creates recurring costs; escalation requires clinical capacity; and health-record integration demands maintenance. These constraints define where interaction-centric systems can justify their complexity.

Simpler interventions may be more cost-effective for motivated users with short-term goals, low-risk populations,

or single barriers such as forgetting. A stepped-care approach should begin with the simplest adequate intervention and escalate only when relapse, multifactor risk, or failed alternatives justify added complexity. Head-to-head and noninferiority comparisons with simple alternatives should therefore be standard; superiority over inexpensive baselines is an empirical obligation, not an assumption.

Conclusions

This viewpoint has argued that the limited real-world impact of AI in preventive health is structural: episodic, prediction-centric systems are misaligned with a domain whose outcomes emerge from behavior sustained over time. The behavioral informatics framework presented here reorients design around continuous interaction, integrating data, knowledge, reasoning, interaction, behavior, and governance into a single system-level account whose effectiveness depends on the couplings between layers rather than on any component in isolation. Interaction-centric systems must balance engagement against autonomy, personalization against comparability, and adaptability against stability, while operating within ethical and regulatory constraints.

Evaluating such systems requires metrics centered on behavioral trajectories and longitudinal outcomes rather than predictive accuracy alone, supported by designs such as microrandomized trials that match the unit of randomization to the unit of intervention. It remains an open question whether this level of system complexity is warranted in all preventive contexts; simpler interventions will sometimes achieve comparable outcomes at lower cost and risk, and continuous interaction carries its own hazards of overmonitoring, dependency, and cognitive burden. Future work should translate these principles into real-world implementations, validate them across diverse populations, and develop shared evaluation frameworks. Achieving this will require collaboration across AI, behavioral science, health care, and human-computer interaction. AI in preventive health is best understood not as a predictive tool but as an ongoing sociotechnical process embedded in everyday behavior.

Acknowledgments

During the preparation of this manuscript, the authors used a generative AI-based language tool solely to improve the clarity and readability of author-written text. No scientific content, analyses, interpretations, or references were generated by AI. The authors reviewed and edited all text and take full responsibility for the content of this manuscript.

Funding

The authors declared no financial support was received for this work.

Data Availability

No datasets were generated or analyzed during this study.

Authors' Contributions

AG conceived the study, developed the conceptual framework, and drafted the manuscript. JH contributed to conceptual refinement and critical revision of the manuscript. All authors approved the final version.

Conflicts of Interest

None declared.

References

1. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. Jan 2019;25(1):44-56. [doi: [10.1038/s41591-018-0300-7](https://doi.org/10.1038/s41591-018-0300-7)] [Medline: [30617339](https://pubmed.ncbi.nlm.nih.gov/30617339/)]
2. Rajkomar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med*. Apr 4, 2019;380(14):1347-1358. [doi: [10.1056/NEJMr1814259](https://doi.org/10.1056/NEJMr1814259)] [Medline: [30943338](https://pubmed.ncbi.nlm.nih.gov/30943338/)]
3. Michie S, van Stralen MM, West R. The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implement Sci*. 2011;6:42. [doi: [10.1186/1748-5908-6-42](https://doi.org/10.1186/1748-5908-6-42)] [Medline: [21513547](https://pubmed.ncbi.nlm.nih.gov/21513547/)]
4. Neal DT, Wood W, Quinn JM. Habits—a repeat performance. *Curr Dir Psychol Sci*. Aug 2006;15(4):198-202. [doi: [10.1111/j.1467-8721.2006.00435.x](https://doi.org/10.1111/j.1467-8721.2006.00435.x)]
5. Eysenbach G. The law of attrition. *J Med Internet Res*. Mar 31, 2005;7(1):e11. [doi: [10.2196/jmir.7.1.e11](https://doi.org/10.2196/jmir.7.1.e11)] [Medline: [15829473](https://pubmed.ncbi.nlm.nih.gov/15829473/)]
6. Kelders SM, Kok RN, Ossebaard HC, Van Gemert-Pijnen J. Persuasive system design does matter: a systematic review of adherence to web-based interventions. *J Med Internet Res*. Nov 14, 2012;14(6):e152. [doi: [10.2196/jmir.2104](https://doi.org/10.2196/jmir.2104)] [Medline: [23151820](https://pubmed.ncbi.nlm.nih.gov/23151820/)]
7. Meyerowitz-Katz G, Ravi S, Arnolda L, Feng X, Maberly G, Astell-Burt T. Rates of attrition and dropout in app-based interventions for chronic disease: systematic review and meta-analysis. *J Med Internet Res*. Sep 29, 2020;22(9):e20283. [doi: [10.2196/20283](https://doi.org/10.2196/20283)] [Medline: [32990635](https://pubmed.ncbi.nlm.nih.gov/32990635/)]
8. Baumel A, Muench F, Edan S, Kane JM. Objective user engagement with mental health apps: systematic search and panel-based usage analysis. *J Med Internet Res*. Sep 25, 2019;21(9):e14567. [doi: [10.2196/14567](https://doi.org/10.2196/14567)] [Medline: [31573916](https://pubmed.ncbi.nlm.nih.gov/31573916/)]
9. Singh B, Olds T, Brinsley J, et al. Systematic review and meta-analysis of the effectiveness of chatbots on lifestyle behaviours. *NPJ Digit Med*. Jun 23, 2023;6(1):118. [doi: [10.1038/s41746-023-00856-1](https://doi.org/10.1038/s41746-023-00856-1)] [Medline: [37353578](https://pubmed.ncbi.nlm.nih.gov/37353578/)]
10. Gavai AK, van Hillegersberg J. AI-driven personalized nutrition: RAG-based digital health solution for obesity and type 2 diabetes. *PLOS Digit Health*. May 2025;4(5):e0000758. [doi: [10.1371/journal.pdig.0000758](https://doi.org/10.1371/journal.pdig.0000758)] [Medline: [40327647](https://pubmed.ncbi.nlm.nih.gov/40327647/)]
11. Pasch S, Ha SY. Human–AI interaction and user satisfaction: empirical evidence from online reviews of AI products. *Int J Hum Comput Interact*. 2026;42(7):5570-5591. [doi: [10.1080/10447318.2025.2545463](https://doi.org/10.1080/10447318.2025.2545463)]
12. Arslan M, Ghanem H, Munawar S, Cruz C. A survey on RAG with LLMs. *Procedia Comput Sci*. 2024;246:3781-3790. [doi: [10.1016/j.procs.2024.09.178](https://doi.org/10.1016/j.procs.2024.09.178)]
13. Huang X, Zhang S. Engagement in LLM chatbot-supported learning: the pivotal roles of GenAI competency and emotion. *Comput Educ Artif Intell*. Jun 2026;10:100559. [doi: [10.1016/j.caeai.2026.100559](https://doi.org/10.1016/j.caeai.2026.100559)]
14. Yardley L, Spring BJ, Riper H, et al. Understanding and promoting effective engagement with digital behavior change interventions. *Am J Prev Med*. Nov 2016;51(5):833-842. [doi: [10.1016/j.amepre.2016.06.015](https://doi.org/10.1016/j.amepre.2016.06.015)] [Medline: [27745683](https://pubmed.ncbi.nlm.nih.gov/27745683/)]
15. Gavai AK, Meuwissen MP. Agentic AI as a coordination paradigm in digital health and agri-food systems. *Patterns (N Y)*. Mar 13, 2026;7(3):101496. [doi: [10.1016/j.patter.2026.101496](https://doi.org/10.1016/j.patter.2026.101496)] [Medline: [42028409](https://pubmed.ncbi.nlm.nih.gov/42028409/)]
16. Amershi S, Weld D, Vorvoreanu M, et al. Guidelines for human-AI interaction. In: CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery; 2019:1-13. [doi: [10.1145/3290605.3300233](https://doi.org/10.1145/3290605.3300233)]
17. Goldstein SP, Evans BC, Flack D, et al. Return of the JITAI: applying a just-in-time adaptive intervention framework to the development of m-health solutions for addictive behaviors. *Int J Behav Med*. Oct 2017;24(5):673-682. [doi: [10.1007/s12529-016-9627-y](https://doi.org/10.1007/s12529-016-9627-y)] [Medline: [28083725](https://pubmed.ncbi.nlm.nih.gov/28083725/)]
18. Nahum-Shani I, Smith SN, Spring BJ, et al. Just-in-time adaptive interventions (JITAI) in mobile health: key components and design principles for ongoing health behavior support. *Ann Behav Med*. May 18, 2018;52(6):446-462. [doi: [10.1007/s12160-016-9830-8](https://doi.org/10.1007/s12160-016-9830-8)] [Medline: [27663578](https://pubmed.ncbi.nlm.nih.gov/27663578/)]
19. Mohr DC, Schueller SM, Montague E, Burns MN, Rashidi P. The behavioral intervention technology model: an integrated conceptual and technological framework for eHealth and mHealth interventions. *J Med Internet Res*. Jun 5, 2014;16(6):e146. [doi: [10.2196/jmir.3077](https://doi.org/10.2196/jmir.3077)] [Medline: [24905070](https://pubmed.ncbi.nlm.nih.gov/24905070/)]
20. Coffey JD, Christopherson LA, Williams RD, et al. Development and implementation of a nurse-based remote patient monitoring program for ambulatory disease management. *Front Digit Health*. 2022;4:1052408. [doi: [10.3389/fdgh.2022.1052408](https://doi.org/10.3389/fdgh.2022.1052408)] [Medline: [36588748](https://pubmed.ncbi.nlm.nih.gov/36588748/)]
21. Wang F, Shi D, Aguilar J, et al. LLM-KGMA: large language model-augmented multi-hop question-answering system based on knowledge graph in medical field. *Knowl Inf Syst*. 2025;67(8):6461-6503. [doi: [10.1007/s10115-025-02399-1](https://doi.org/10.1007/s10115-025-02399-1)]
22. Nan S, Qiao Y, Xie Y, et al. Automatic conflict detection and resolution in digital service network requirement models using large language models. *SOCA*. 2026;20:187-198. [doi: [10.1007/s11761-025-00480-9](https://doi.org/10.1007/s11761-025-00480-9)]
23. Lai X, Chen J, Lai Y, et al. Using large language models to enhance exercise recommendations and physical activity in clinical and healthy populations: scoping review. *JMIR Med Inform*. May 27, 2025;13:e59309. [doi: [10.2196/59309](https://doi.org/10.2196/59309)] [Medline: [40424584](https://pubmed.ncbi.nlm.nih.gov/40424584/)]

24. Hila A. The epistemological consequences of large language models: rethinking collective intelligence and institutional knowledge. *AI Soc.* 2025;41(1):79-97. [doi: [10.1007/s00146-025-02426-3](https://doi.org/10.1007/s00146-025-02426-3)]
25. Saup TO, Asghar J, Kanbach DK, Kraus S. From pilots to decision systems: embedding generative AI into strategic decision-making through a socio-technical and governance lens. *J Decis Syst.* 2026;35(1):2597835. [doi: [10.1080/12460125.2025.2597835](https://doi.org/10.1080/12460125.2025.2597835)]
26. Aghae F, Shaker HR. RB-LLM control: an intelligent control framework with rule-based large language model decision-making. *Aerosp Sci Technol.* Jan 2026;168:111259. [doi: [10.1016/j.ast.2025.111259](https://doi.org/10.1016/j.ast.2025.111259)]
27. Hori K, Suzuki K, Ogata T. Enhancement of long-horizon task planning via active and passive modification in large language models. *Sci Rep.* Feb 28, 2025;15(1):7113. [doi: [10.1038/s41598-025-91448-4](https://doi.org/10.1038/s41598-025-91448-4)] [Medline: [40016395](https://pubmed.ncbi.nlm.nih.gov/40016395/)]
28. Perez JE, Ong EC. Prompting for engagement: using the ICAP framework to guide prompt design in LLM-powered dialogue-based tutoring system for novice programmers. *Procedia Comput Sci.* 2025;272:113-121. [doi: [10.1016/j.procs.2025.10.186](https://doi.org/10.1016/j.procs.2025.10.186)]
29. Norman DA. *The Design of Everyday Things*. MIT Press; 2013. ISBN: 9780262525671
30. Roeschl T, Hoffmann M, Hashemi D, et al. Assessing the limitations of large language models in clinical practice guideline-concordant treatment decision-making on real-world data: retrospective study. *JMIRx Med.* Nov 3, 2025;6:e74899. [doi: [10.2196/74899](https://doi.org/10.2196/74899)] [Medline: [41190890](https://pubmed.ncbi.nlm.nih.gov/41190890/)]
31. Ryan RM, Deci EL. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Am Psychol.* 2000;55(1):68-78. [doi: [10.1037/0003-066X.55.1.68](https://doi.org/10.1037/0003-066X.55.1.68)] [Medline: [11392867](https://pubmed.ncbi.nlm.nih.gov/11392867/)]
32. Ng JY, Ntoumanis N, Thøgersen-Ntoumani C, et al. Self-determination theory applied to health contexts: a meta-analysis. *Perspect Psychol Sci.* Jul 2012;7(4):325-340. [doi: [10.1177/1745691612447309](https://doi.org/10.1177/1745691612447309)] [Medline: [26168470](https://pubmed.ncbi.nlm.nih.gov/26168470/)]
33. Miller WR, Rollnick S. *Motivational Interviewing: Helping People Change*. 3rd ed. Guilford Press; 2013.
34. Pedamallu H, Ehrhardt MJ, Maki J, Carcone AI, Hudson MM, Waters EA. Technology-delivered adaptations of motivational interviewing for the prevention and management of chronic diseases: scoping review. *J Med Internet Res.* Aug 9, 2022;24(8):e35283. [doi: [10.2196/35283](https://doi.org/10.2196/35283)] [Medline: [35943775](https://pubmed.ncbi.nlm.nih.gov/35943775/)]
35. Karve Z, Calpey J, Machado C, Knecht M, Mejia MC. New doc on the block: scoping review of AI systems delivering motivational interviewing for health behavior change. *J Med Internet Res.* Sep 16, 2025;27:e78417. [doi: [10.2196/78417](https://doi.org/10.2196/78417)] [Medline: [40957014](https://pubmed.ncbi.nlm.nih.gov/40957014/)]
36. Fogg BJ. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann Publishers; 2003. ISBN: 9780080479941
37. Oinas-Kukkonen H, Harjumaa M. Persuasive systems design: key issues, process model, and system features. *Commun Assoc Inf Syst.* 2009;24. [doi: [10.17705/1CAIS.02428](https://doi.org/10.17705/1CAIS.02428)]
38. Susser D, Roessler B, Nissenbaum HF. Online manipulation: hidden influences in a digital world. *Geo Law Technol Rev.* 2019;4(1):1-45. URL: <https://georgetownlawtechreview.org/online-manipulation-hidden-influences-in-a-digital-world/GLTR-01-2020/> [Accessed 2026-10-14]
39. Feretzakis G, Vagena E, Kalodanis K, Peristera P, Kalles D, Anastasiou A. GDPR and large language models: technical and legal obstacles. *Future Internet.* 2025;17(4):151. [doi: [10.3390/fi17040151](https://doi.org/10.3390/fi17040151)]
40. Nolte H, Rateike M, Finck M. Robustness and cybersecurity in the EU Artificial Intelligence Act. In: *FACt '25: Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery; 2025:283-295. [doi: [10.1145/3715275.3732020](https://doi.org/10.1145/3715275.3732020)]
41. Chittimalla SK, Potluri LK. Explainable AI frameworks for large language models in high-stakes decision-making. In: *2025 International Conference on Advanced Computing Technologies (ICoACT)*. IEEE; 2025:1-6. [doi: [10.1109/ICoACT63339.2025.11005089](https://doi.org/10.1109/ICoACT63339.2025.11005089)]
42. Idrissi-Yaghir A, Arzideh K, Schäfer H, et al. Using a diverse test suite to assess large language models on fast health care interoperability resources knowledge: comparative analysis. *J Med Internet Res.* Aug 12, 2025;27:e73540. [doi: [10.2196/73540](https://doi.org/10.2196/73540)] [Medline: [40795315](https://pubmed.ncbi.nlm.nih.gov/40795315/)]
43. Perski O, Blandford A, West R, Michie S. Conceptualising engagement with digital behaviour change interventions: a systematic review using principles from critical interpretive synthesis. *Transl Behav Med.* Jun 2017;7(2):254-267. [doi: [10.1007/s13142-016-0453-1](https://doi.org/10.1007/s13142-016-0453-1)] [Medline: [27966189](https://pubmed.ncbi.nlm.nih.gov/27966189/)]
44. Klasnja P, Hekler EB, Shiffman S, et al. Microrandomized trials: an experimental design for developing just-in-time adaptive interventions. *Health Psychol.* Dec 2015;34S:1220-1228. [doi: [10.1037/hea0000305](https://doi.org/10.1037/hea0000305)] [Medline: [26651463](https://pubmed.ncbi.nlm.nih.gov/26651463/)]
45. Jin L, Ong JC, Elangovan K, et al. Large language models in randomized controlled trials design: observational study. *J Med Internet Res.* Sep 3, 2025;27:e67469. [doi: [10.2196/67469](https://doi.org/10.2196/67469)] [Medline: [41055064](https://pubmed.ncbi.nlm.nih.gov/41055064/)]

Abbreviations

BIT: behavioral intervention technology

JITAI: just-in-time adaptive intervention

LLM: large language model

Edited by Ivan Steenstra; peer-reviewed by Spandana Poddutoori, Stanislav Dadelo, Suraj Kath; submitted 21.Apr.2026; final revised version received 06.Aug.2026; accepted 06.Aug.2026; published 28.Oct.2026

Please cite as:

Gavai A, Heringa J

From Prediction to Interaction: A Behavioral Informatics Framework for AI in Preventive Health

J Med Internet Res 2026;28:e98961

URL: <https://www.jmir.org/2026/1/e98961>

doi: [10.2196/98961](https://doi.org/10.2196/98961)

© Anand Gavai, Jaap Heringa. Originally published in the Journal of Medical Internet Research (<https://www.jmir.org>), 28.Oct.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in the Journal of Medical Internet Research (ISSN 1438-8871), is properly cited. The complete bibliographic information, a link to the original publication on <https://www.jmir.org/>, as well as this copyright and license information must be included.