

Kinesis: A Multi-Agent Wearable System for Contextualized Posture Correction

Jiayue Ni*
Harvard University
Cambridge, MA, USA
chloe_ni@gsd.harvard.edu

Jianing Yu*
Harvard University
Graduate School of Design
Cambridge, MA, USA
nomy_yu@gsd.harvard.edu

Lilith Yu
Harvard University
Cambridge, MA, USA
lilith_yu@gsd.harvard.edu




Figure 1: Kinesis is a multi-agent wearable system that coordinates specialized sensing and decision agents to deliver adaptive bodily interventions. The example deployment scenarios illustrate how interventions adapt during outdoor running (Scenario 1) and a formal interview (Scenario 2).

Abstract

Most current posture correction solutions are limited to a rigid "sensing-to-feedback" loop, failing to integrate user context and positive reinforcement for long-term behavioral adaptation. This paper introduces Kinesis, a multi-agent wearable system for adaptive bodily intervention. Kinesis coordinates three specialized agents: a *Context Agent* that interprets ambient social and environmental cues through the camera-equipped wearable devices, a *Body Agent* that monitors electromyographic signals for posture assessment, and a *Coach Agent* that integrates bodily, contextual, and external health data to determine when and how to intervene. The agents coordinate through a shared-state blackboard, while Model Context Protocol (MCP)-compatible connectors expose data and capabilities from external wearable devices. Together, these components form

a closed loop that integrates posture, context, and health data to determine when and how to intervene, deliver feedback, and verify the bodily response. We present the design and implementation of a working prototype and discuss how multi-agent architectures can support context-aware, somatically grounded behavior change in everyday personal health computing.

CCS Concepts

• **Human-centered computing** → Ubiquitous and mobile computing systems and tools.

Keywords

Wearable Computing; Multi-agent Systems; Context-aware Computing; Just-In-Time Adaptive Intervention

*These two authors contributed equally to this research.



ACM Reference Format:

Jiayue Ni, Jianing Yu, and Lilith Yu. 2026. Kinesis: A Multi-Agent Wearable System for Contextualized Posture Correction. In *Companion of the 2026 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '26)*, October 11–15, 2026, Shanghai, China. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3798063.3839204>

1 Introduction

Poor posture is a common consequence of prolonged sedentary work and the widespread use of digital devices. To address this, wearable and ubiquitous computing systems have explored a variety of techniques for recognizing postural deviations and delivering corrective feedback [10]. Despite advances in inertial, pressure, vision-based, and electromyographic sensors for real-time assessment [9], posture correction typically relies on a direct sensing-to-feedback pipeline. This paradigm treats posture correction predominantly as a sensing challenge, largely overlooking user context and the opportunity to foster long-term behavioral adaptation through positive feedback loops.

Research on Just-In-Time Adaptive Interventions (JITAs) has similarly emphasized that support should be delivered not only when it is needed, but also when users are receptive to it [12]. Contextualized posture correction therefore requires coordinating three dimensions of knowledge: bodily state, environmental context, and the optimal intervention strategy given both.

We present **Kinesis** (Figure 1), a multi-agent wearable system that treats posture correction as a contextual decision-making process. Kinesis decouples sensing, reasoning, and actuation through a shared-state architecture based on Model Context Protocols (MCPs). Within this ecosystem, three specialized agents coordinate across distinct timescales: A *Body Agent* monitors spinal alignment and muscular activity through inertial and surface electromyographic sensing; a *Context Agent* interprets ambient social and environmental conditions through camera-equipped wearable devices; and a *Coach Agent* integrates various signals with external health data and intervention history to decide whether, when, and how corrective feedback should be delivered. This architecture facilitates a closed-loop intervention process, leveraging cumulative interaction history to ultimately drive the longitudinal adaptation of multidimensional feedback.

Through the design and implementation of a working prototype, this paper explores a potential path toward more situated personal health interventions by offering: (1) a tri-agent architecture that separates bodily assessment, contextual interpretation, and intervention planning; (2) a protocol-driven coordination system designed to help heterogeneous wearable and health devices collaborate in a shared intervention loop. We use posture correction as a preliminary case study to investigate how agentic architectures might support more situated, adaptive, and interpretable forms of personal health intervention.

2 Related Work

2.1 Context-Aware Wearable Systems for Posture Correction

Wearable and ubiquitous systems have explored a range of techniques for recognizing posture and musculoskeletal activity, including inertial, pressure, ear-worn, camera-based, and electromyographic sensing [9, 10]. Recent systems increasingly combine heterogeneous signals to estimate bodily posture and biomechanical states during everyday movement. For example, KneeGuard integrates inertial and surface electromyographic sensing for gait retraining [15], while PressInPose combines insole pressure sensing

with wearable inertial measurements for full-body pose estimation [3]. Posture-correction systems commonly translate detected deviations into visual, auditory, or vibrotactile feedback [2, 13]. However, this interaction is often operationalized as a direct mapping between a detected posture and a predefined alert.

Research on JITAs argues that behavioral support should account for both users' momentary needs and their receptivity to intervention [12]. Prior work has therefore examined how behavioral, temporal, and environmental context can indicate whether a user is likely to notice, engage with, or act upon a mobile health intervention [1, 6, 11]. More recently, LLM-based approaches have explored replacing fixed intervention rules with adaptive reasoning over intervention timing and content [5]. Building on these directions, Kinesis coordinates bodily sensing with ambient context interpretation to determine not only whether posture correction is needed, but also when and how feedback should be delivered.

2.2 Agentic Architectures for Embodied Health Intervention

Recent health-oriented AI systems have incorporated wearable data into LLM-based physical-activity, sleep, and well-being coaching [7, 8, 14]. These systems demonstrate how personal sensing data, behavioral theories, and contextual information can ground personalized health recommendations. HealthGuru, for example, adopts a multi-agent framework that coordinates wearable-data interpretation, contextual factors, behavior-change support, and adaptive recommendation selection [14]. Other work investigates how coaching agents can dynamically allocate reasoning effort while maintaining interaction quality and responsiveness [4].

However, these systems primarily support users through conversational guidance or retrospective recommendations. Kinesis instead applies a multi-agent architecture to continuous embodied intervention. It decomposes posture correction into a *Body Agent* that assesses muscular and postural signals, a *Context Agent* that interprets the surrounding social and environmental situation, and a *Coach Agent* that plans intervention timing and modality. This architecture forms a closed loop in which Kinesis senses posture, determines an appropriate intervention opportunity, delivers feedback, and observes the resulting bodily response to inform subsequent interventions.

To bridge the gap between continuous kinesiological sensing and context-aware behavior modification, our system utilizes a decoupled hardware-software ecosystem consisting of several biometric wearables and AI agents to execute a continuous behavioral adaptation loop from sense to interpret, decide, intervene, measure, and adapt.

3 System Design

Kinesis implements posture correction as an event-driven, multi-agent decision process rather than a direct mapping from sensor thresholds to feedback. As illustrated in Figure 2, the system separates bodily assessment, situational interpretation, and intervention planning across three specialized agents. These agents operate asynchronously at different update rates and coordinate through a shared-state blackboard. External wearable devices and health-data sources are incorporated through protocol-compatible connectors,




Figure 2: Event-driven coordination and closed-loop intervention workflow in Kinesis.

enabling the *Coach Agent* to reason over both immediate sensor observations and broader health context.

In Kinesis, an agent is a stateful computational component with a distinct observation space, functional objective, and set of available actions. Each agent independently interprets its inputs, publishes an assessment to the shared coordination layer, and reacts to relevant state changes. This decomposition allows the sensing and decision components to be developed and updated independently while preserving an inspectable intervention process.

3.1 Specialized Agents and Data Sources

3.1.1 Body Agent. The *Body Agent* maintains a real-time estimate of the user’s postural and muscular state. It is implemented through a back-worn sensor-actuator assembly controlled by an ESP32 microcontroller. Two MPU6050 inertial measurement units positioned on the upper and lower back estimate relative spinal deviation, while a three-channel surface electromyography array monitors muscle engagement in the rhomboid and lumbar regions. Operating at approximately 500 ms intervals, the *Body Agent* classifies postural states, including upright, slouching, hunched, and lateral-leaning postures, and estimates signs of muscular fatigue. It publishes these assessments to the shared state but does not independently determine whether an intervention should occur. When the *Coach Agent* authorizes an intervention, the *Body Agent* activates four eccentric rotating mass motors distributed across the shoulder and lumbar regions. It subsequently samples the sensor streams again to determine whether the user corrected the posture and whether the targeted muscles were engaged. This post-intervention assessment provides the somatic outcome used to close the feedback loop.

3.1.2 Context Agent. The *Context Agent* estimates whether the user’s current situation is appropriate for an intervention. Camera- and microphone-equipped AI glasses, implemented using a Xiao ESP32 S3 Sense, provide egocentric visual and ambient signals. At a slower update rate of approximately two to three seconds, the

agent interprets the user’s activity and surrounding social environment, such as focused individual work, a conversation, or a formal presentation. Rather than issuing intervention commands, the *Context Agent* publishes a structured assessment of activity, social setting, and estimated interruptibility. This assessment functions as an opportunity signal for the *Coach Agent*. For example, the same postural deviation may be treated as immediately actionable during private desk work but deferred during an interview or presentation.

3.1.3 MCP-Enabled Device Integration. The ecosystem of Kinesis is device-agnostic; individual nodes are wrapped as modular MCP servers that expose data directly to the agent. As a proof of concept, a third-party Whoop wristband is integrated to supply systemic physiological metrics like sleep scores and HRV, enabling the system’s actuators to deliver haptic prompts calibrated to the user’s broader somatic condition. This protocol-driven decoupling allows new environmental sensors or commercial wearables to seamlessly join the system without modifying the core agent architecture.

3.1.4 Coach Agent. While the Body and Context Agents handle reactive, situational sensing, the *Coach Agent* serves as the central orchestration and decision engine, maintaining the system’s long-term memory to manage behavioral strategies across hours, days, and weeks. As shown in Figure 2, its decision process comprises three stages: 1) need assessment determines whether the user’s postural deviation and muscular state warrant intervention; 2) the opportunity gate evaluates whether the user is currently receptive to interruption, incorporating activity, social setting, recent intervention frequency, and an attention budget; 3) policy selection determines the appropriate intervention action and parameters, including modality, location, intensity, and timing.

The *Coach Agent* also maintains a rolling history of recent observations, decisions, and outcomes, supporting reasoning over temporal patterns rather than isolated sensor readings.




Figure 3: Kinesis interface for monitoring personal agent states, intervention decisions, and longitudinal posture history (left), and agent-to-agent communication (right).

3.2 Event-Driven Multi-Agent Coordination

Rather than coupling agents through direct function calls, Kinesis mediates all coordination through a shared-state blackboard implemented as a central MCP server. Each agent reads and writes to namespaced state resources — such as `state://body/posture`, `state://glasses/context`, and `state://coach/mode` — establishing a canonical, inspectable representation of the system’s beliefs at any moment. When an agent updates its state, the server pushes notifications to all subscribers via MCP resource subscriptions, enabling downstream agents to react without polling or tight coupling.

This blackboard pattern produces a natural information flow: the *Body Agent* publishes posture classification and EMG readings; the *Context Agent* publishes scene type and interruptibility; the *Coach Agent* consumes both, along with external biometrics from MCP-connected devices, and publishes an authorized intervention action that the *Body Agent* reads and executes. No agent imports or invokes another — coordination emerges from shared state. For cases requiring richer exchange than publish-subscribe, the system supports structured inter-agent discussions routed through the blackboard. Before triggering any haptic cue, the *Body Agent* issues a query to the *Context Agent* — e.g., describing the current deviation and asking whether the moment is interruptible. The *Context Agent* responds with a situational assessment, and only then does the *Body Agent* proceed or defer. These exchanges are logged as first-class events, providing a full audit trail visible on the system dashboard.

Each agent’s reasoning layer is powered by LLMs invoked with MCP tool-use capabilities, allowing the model to read state, update state, send haptic commands, or initiate inter-agent queries within a single multi-step reasoning turn. This design means that adding a new agent — or a new MCP-connected device — requires only registering new state resources on the blackboard; existing agents can discover newly registered resources through a common interface, although incorporating new signals may still require schema definitions, authorization policies, and updates to agent instructions.

3.3 User-Facing Interface

Figure 3 (left) shows the Kinesis dashboard, which provides two complementary views of system behavior. The left panel presents the current state of the personal agent ecosystem. Users can inspect whether the Coach, Body, Context, and external-device connectors are active, view the latest posture and context assessments, and review the external health signals available to the system. The interface also displays the *Coach Agent*’s current interpretation and selected intervention policy, making visible why the system intervenes, defers, or suppresses a cue. The right panel exposes communication and history at two timescales. The agent-communication view presents recent exchanges and coordination events, including detected posture patterns, contextual assessments, deferred interventions, and verified outcomes. The longitudinal insight view aggregates these events into higher-level observations, such as recurring posture triggers, contexts associated with successful correction, and changes in posture patterns over time.

The interface is not intended only as a monitoring console. It provides an interpretability layer between the user and the multi-agent system. Internally, the event history supports the *Coach Agent*’s future adaptation; externally, it allows users to examine which data sources contributed to an intervention, how the decision was reached, and whether the intervention produced a measurable bodily response. In future deployments, this interface could also support user correction of agent inferences, adjustment of intervention preferences, and control over which external data sources participate in decision-making.

4 Discussion, Limitations, and Future Work

Kinesis demonstrates the technical feasibility of coordinating bodily sensing, contextual interpretation, and intervention planning through a shared multi-agent architecture. However, the current prototype does not establish whether agent-based contextual coordination improves posture-correction outcomes, reduces inappropriate interruptions, or supports sustained behavior change relative to simpler intervention pipelines. We therefore discuss the

implications and limitations of this architectural decomposition and outline priorities for evaluation, data governance, safety, and longitudinal personalization.

4.1 Architectural Implications of Multi-Agent Intervention

Separating bodily assessment, contextual interpretation, and intervention planning makes the assumptions behind each intervention more inspectable than in a monolithic sensing-to-feedback pipeline. It also allows agents and sensors operating at different temporal resolutions to update independently: the *Body Agent* can continuously monitor posture, while the *Context Agent* interprets environmental conditions less frequently and the *Coach Agent* reasons over their combined histories. The shared-state blackboard provides a common representation through which these assessments and decisions can be inspected without tightly coupling individual components.

This decomposition also introduces new coordination challenges. Delayed observations, stale state, conflicting assessments, and accumulated latency may affect when an intervention is delivered. Moreover, an inspectable event log does not by itself guarantee that an intervention is reliable or interpretable to the user. Future implementations should associate observations with timestamps and confidence estimates, define explicit conflict-resolution and fallback policies, and constrain the actions available to each agent. These mechanisms are necessary for determining when the system should intervene, request additional evidence, or defer action under uncertainty.

4.2 Preliminary Observations and Evaluation Plan

During development, we conducted informal engineering tests to verify the operation of the sensing, communication, and actuation pipeline. These tests included checking whether the *Body Agent*'s posture classifications aligned with visually observed posture states, whether the *Context Agent* distinguished coarse individual and social settings, and whether intervention commands produced the intended haptic output. These observations supported iterative development but were not collected through a controlled protocol and should not be interpreted as evidence of intervention effectiveness.

A systematic evaluation should distinguish the contribution of contextual sensing from that of multi-agent reasoning. One possible study could compare Kinesis with two progressively stronger baselines: a body-only condition that directly maps detected posture deviations to haptic feedback, and a deterministic context-aware condition that applies predefined rules to gate interventions. Comparing these conditions would help determine whether contextual gating reduces inappropriate interruptions and whether agent-based coordination provides additional value beyond fixed context rules. System-level measures should include posture- and context-recognition accuracy, intervention latency, state consistency, action-execution reliability, correction success rate, and false or deferred alert rates. User-facing measures should examine perceived appropriateness, comfort, trust, interruption burden, and perceived control. Longer deployments would additionally be required to assess sustained posture change, habituation, and notification fatigue.

4.3 Privacy, Safety, and User Control

Kinesis coordinates camera and microphone observations, IMU and electromyographic signals, intervention histories, and health information from external wearable devices. These data differ in sensitivity and create risks when combined into a shared representation. The current prototype invokes LLM-based agents over derived system states but does not yet provide a complete account of where all reasoning executes, which data leave individual devices, or how long observations and interaction histories are retained. These limitations must be addressed before deployment beyond controlled prototyping. Future implementations should minimize the information exposed to each agent, process raw audiovisual and physiological signals locally where feasible, and publish only the derived states required for intervention planning. Access to blackboard resources should follow explicit authorization and least-privilege policies, with auditable retention and deletion controls. Safety mechanisms should also bound haptic intensity and intervention frequency, default to deferral when agent assessments conflict or remain uncertain, and allow users to pause, override, or disable interventions. The dashboard could support these controls by showing which signals informed a decision and allowing users to correct inferred context, revise intervention preferences, and remove accumulated history. Kinesis should therefore be understood as a behavior-support prototype rather than a diagnostic or medical system.

4.4 Longitudinal Personalization and Inter-Agent Extensions

The rolling event history maintained by the *Coach Agent* provides a foundation for longitudinal personalization. By recording when postural deviations occur, which intervention strategies are selected, and whether users subsequently correct and maintain their posture, future versions could adapt intervention timing, modality, intensity, and frequency to individual patterns. Such personalization should remain inspectable and revisable: users should be able to review the patterns inferred by the system, correct inaccurate assumptions, and reset or limit the use of historical data. This is important for preventing temporary routines or erroneous inferences from becoming persistent intervention policies.

A more speculative extension concerns interoperability across independently governed personal agents. As an initial feasibility exploration, we used mock user data to visualize agent-to-agent communication in the Kinesis interface (Figure 3, right). Rather than exchanging raw biometric data, personal agents could selectively share abstracted intervention strategies or aggregated outcomes under explicit user consent. This functionality is not part of the currently evaluated prototype and would require clear mechanisms for authorization, provenance, revocation, and privacy-preserving aggregation. We therefore position cross-user coordination as a longer-term research direction, while keeping the present contribution focused on contextual intervention within a single user's personal agent ecosystem.

5 Conclusion

Kinesis presents a somatically grounded closed loop for adaptive posture intervention through a tri-agent architecture coordinated

by an MCP blackboard. By integrating posture signals, situational context, and health history, Kinesis determines not only if correction is needed, but also when and how feedback should be delivered. The prototype demonstrates situated, interpretable posture interventions via multi-agent coordination, laying the groundwork for future longitudinal studies on long-term effectiveness and user control.

References

- [1] Woohyeok Choi, Sangkeun Park, Duyeon Kim, Youn-kyung Lim, and Uichin Lee. 2019. Multi-stage receptivity model for mobile just-in-time health intervention. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 3, 2 (2019), 1–26.
- [2] Minh Duc Dang, Duy Phuoc Luong, Christopher Napier, and Lawrence H Kim. 2025. Co-Design & Evaluation of Visual Interventions for Head Posture Correction in Virtual Reality Games. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [3] Yang Gao, Wenbo Zhang, Junbin Ren, Ruihao Zheng, Yingcheng Jin, Di Wu, Lin Shu, Xiangmin Xu, and Zhanpeng Jin. 2024. Pressinpose: Integrating pressure and inertial sensors for full-body pose estimation in activities. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 4 (2024), 1–28.
- [4] Andreas Göldi, Roman Rietsche, and Lyle Ungar. 2025. Efficient Management of LLM-Based Coaching Agents' Reasoning While Maintaining Interaction Quality and Speed. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [5] David Haag, Devender Kumar, Sebastian Gruber, Dominik P Hofer, Mahdi Sareban, Gunnar Treff, Josef Niebauer, Christopher N Bull, Albrecht Schmidt, and Jan David Smeddinck. 2025. The last JITAI? Exploring large Language models for issuing Just-in-Time adaptive interventions: fostering physical activity in a prospective cardiac rehabilitation setting. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [6] Brennan Jones, Yan Xu, Qisheng Li, and Stefan Scherer. 2024. Designing a Proactive Context-Aware AI Chatbot for People's Long-Term Goals. In *Extended abstracts of the CHI conference on human factors in computing systems*. 1–7.
- [7] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A Landay. 2025. GPTCoach: towards LLM-based physical activity coaching. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–46.
- [8] Justin Khasentino, Anastasiya Belyaeva, Xin Liu, Zhun Yang, Nicholas A Furlotte, Chace Lee, Erik Schenck, Yojan Patel, Jian Cui, Logan Douglas Schneider, et al. 2025. A personal health large language model for sleep and fitness coaching. *Nature Medicine* 31, 10 (2025), 3394–3403.
- [9] Christian Krauter, Katrin Angerbauer, Aimée Sousa Calepso, Alexander Achberger, Sven Mayer, and Michael Sedlmair. 2024. Sitting posture recognition and feedback: A literature review. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [10] Hosub Lee, Young Sang Choi, Sunjae Lee, and Eunsoo Shim. 2013. Smart pose: mobile posture-aware system for lowering physical health risk of smartphone users. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2257–2266.
- [11] Varun Mishra, Florian Künzler, Jan-Niklas Kramer, Elgar Fleisch, Tobias Kowatsch, and David Kotz. 2021. Detecting receptivity for mHealth interventions in the natural environment. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 5, 2 (2021), 1–24.
- [12] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. 2016. Just-in-time adaptive interventions (JITAI) in mobile health: key components and design principles for ongoing health behavior support. *Annals of behavioral medicine* (2016), 1–17.
- [13] Yushi Takayama, Shun Ishii, Anna Yokokubo, and Guillaume Lopez. 2021. Detecting forward leaning posture using eSense and developing a posture improvement promoting system. In *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers*. 178–179.
- [14] Xingbo Wang, Janessa Griffith, Daniel A Adler, Joey Castillo, Tanzeem Choudhury, and Fei Wang. 2025. Exploring personalized health support through data-driven, theory-guided LLMs: A case study in sleep health. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [15] Baichen Yang, Xinyi Zhang, Jiayi Zhang, Zirui Huang, Qiqi Lu, Jin Zhang, Hai Hu, and Qian Zhang. 2024. KneeGuard: A Calibration-free Wearable Monitoring System for Knee Osteoarthritis Gait Re-training via Effortless Wearing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 4 (2024), 1–29.