

Agentic AI for Context-Aware Adaptive UI Design in Wearable Healthcare

Dr. MK Jayanthi Kannan¹, Mariam Safar Mohammed Alshahrani²

¹*Professor, School of Computing Science and Engineering, VIT Bhopal University, Bhopal-Indore Highway, Kothrikalan, Sehore, Madhya Pradesh – 466114*

²*Digital Government Authority (DGA), Digital Government Authority of KSA, Riyadh Province, Riyadh, KSA Postal Code: 11564, Kingdom of Saudi Arabia.*

Abstract—Wearable healthcare devices sense physiological and behavioural state continuously, and the mechanisms for acting on that state are individually mature. Physiological computing has closed the sensing-to-adaptation loop since the 1990s; just-in-time adaptive interventions have formalised context-triggered decision-making; adaptive user interface research has produced runtime interface generation; reinforcement learning and combinatorial optimisation supply policy and candidate selection; and control theory supplies provable stability. These traditions have developed largely independently and are rarely combined. This chapter is a conceptual and integrative contribution. It does not propose a new algorithm, and it does not argue that existing algorithms are inadequate. It argues instead that the mechanisms required for context-aware adaptation of wearable healthcare interfaces already exist across separate literatures, and that their integration is obstructed by a conflation of distinct speeds. Sensing, inference, decision-making and interface change are commonly treated as a single adaptation rate, whereas adaptive interface research has

repeatedly shown that fast structural change degrades user performance even when the underlying prediction is accurate (Findlater & McGrenere, 2004; Gajos et al., 2008).

We present an integrated architecture organised as Sense → Infer → Confidence/Context Validation → Persistence/Change Detection → Decide Whether to Adapt → Select Adaptation → Optimise Candidates → Apply Change Safely → Monitor Response → Learn/Recover, assigning an established mechanism to each stage and labelling every element by provenance. The central proposition is that fast sensing and fast decision-making do not entail fast structural interface change, and that separating these speeds is what makes responsive adaptation compatible with predictable, glanceable and safe wearable interaction.

Index Terms—Adaptive user interfaces; wearable healthcare; physiological computing; context-aware computing; just-in-time adaptive interventions; reinforcement learning; control theory; mixed-initiative interaction; agentic AI; human factors.

Provenance labelling used throughout

Label	Meaning
[ESTABLISHED]	A mechanism or finding directly demonstrated in the cited literature.
[TRANSFERRED]	An established mechanism conceptually carried from another domain into wearable healthcare UI adaptation, where it has not been demonstrated.
[PROPOSED]	Our own synthesis, integration, architectural relation or conceptual extension.
[UNKNOWN]	Insufficient evidence located. Not converted into a claim anywhere in this chapter.

These labels are load-bearing. A mechanism marked [TRANSFERRED] or [PROPOSED] is not to be read as an established fact.

I. INTRODUCTION

A wearable health device on a user's wrist can estimate cognitive load from photoplethysmography in real time, classify physical activity from accelerometry, and infer attention state from electroencephalography with reported classification accuracies near eighty per cent. A reinforcement learning policy can select an adaptation in milliseconds. An integer programme can generate an optimised grid layout in seconds. Every technical precondition for a continuously self-adapting wearable health interface is satisfied. Yet the adaptive interface literature has spent two decades documenting that interfaces which change frequently perform worse than interfaces that do not change at all, even when their predictions are correct (Findlater & McGrenere, 2004; Lavie & Meyer, 2010). The obstacle is not capability. It is that adaptation imposes costs on the user — loss of spatial memory, loss of predictability, loss of awareness of unused features — and those costs scale with the rate and the structural depth of change. This chapter treats that tension as the organising problem for wearable healthcare interface adaptation, and argues that it is addressable by decomposition rather than by invention. The contribution is integrative: identifying which established mechanism addresses which stage of an adaptation lifecycle, and specifying how those stages relate to one another. Section 2 states the conceptual position and its limits. Section 3 surveys the existing mechanisms and their provenance. Section 4 presents the integrated architecture. Section 5 gives the mechanism matrix. Section 6 sets out the adaptation-level gradient and the human-factor constraints that govern it. Section 7 develops the multi-speed argument. Section 8 examines where agentic AI can and cannot be justified. Section 9 states the proposed conceptual contribution explicitly. Section 10 records limitations and the empirical work that would be required to test any of this.

II. RESEARCH SCOPE AND CONCEPTUAL POSITION

This chapter is conceptual and integrative research. It is important to state plainly what it is not.

- It is not a systematic review. No PRISMA protocol was followed and no exhaustive search was

conducted; sources were selected for mechanism relevance.

- It is not an empirical study. No data were collected and no hypothesis is tested here.
- It does not introduce a new reinforcement learning, control or optimisation algorithm.
- It does not claim that existing algorithms are inadequate. Each mechanism surveyed performs the task it was designed for.
- It does not claim that agentic AI is superior to reinforcement learning or control theory. Section 8 argues the opposite for most of the architecture.

What the chapter does claim is narrower and, we think, defensible: that the mechanisms needed for responsive yet stable adaptation of wearable healthcare interfaces exist across several research traditions that rarely cite one another, and that assembling them requires separating speeds which are usually conflated.

III. EXISTING MECHANISMS AND THEIR PROVENANCE

3.1 Physiological computing and biocybernetics adaptation

Fairclough (2009) coined physiological computing for a class of systems built on closed-loop control, in which implicit measures from the brain and body inform adaptation at the user interface [ESTABLISHED]. The underlying biocybernetics loop predates the term: Pope, Bogart and Bartolome (1995) built a system evaluating operator engagement in an automated task, and Prinz et al. (2000) a closed-loop system for adaptive task allocation. The mechanism is a negative feedback loop. Psychophysiological data is captured, analysed and converted into a control input in real time, and the loop initiates an adaptive response intended to move a target psychological variable — engagement, workload, attention — toward a desired range. Inference validity is the load-bearing element. Every biocybernetic system rests on a psychophysiological inference, such as inferring arousal from skin conductance. Ewing, Fairclough and Gilleade (2016) structured their adaptive game development explicitly around validating that inference before building the loop, and state that where the link between measure, psychological construct and adaptive logic is weak, the closed-loop system is compromised [ESTABLISHED]. What these systems adapt is task

demand or task allocation, not interface structure. Across the physiological computing corpus surveyed here we located no treatment of adaptation cost, predictability, spatial memory or user override. More recent systems extend the loop to structural visual adaptation. SYMBIOTIK drives dashboard adaptation from EEG-estimated mental workload using reinforcement learning, evaluated with 120 participants (López-Cardona, Attygalle et al., 2025) [ESTABLISHED]. EEG-driven adjustment of dynamic visual complexity in virtual reality improved task performance and reduced perceived workload against a reverse-adaptation control. Neither system is a wearable.

3.2 Just-in-time adaptive interventions

JITAI is the mature framework for data-driven adaptive intervention in health, specified around six elements: distal outcomes, proximal outcomes, tailoring variables, decision points, decision rules and intervention options. HeartSteps is the canonical instantiation (Klasnja et al., 2019) [ESTABLISHED]. Over forty-two days, at up to five user-selected decision points per day, the system chose whether to deliver a contextually tailored activity suggestion and, if so, whether it should be a walking or an antisedentary suggestion. The reward was the log step count over the following thirty minutes, analysed by centred and weighted least squares. Availability gating is the mechanism most directly relevant to wearable healthcare. Participants were treated as unavailable, and therefore ineligible for intervention, if currently walking or running, if they had completed an activity bout within the previous ninety seconds, if driving, if disconnected, or if they had manually disabled the intervention; availability was determined from on-device activity recognition [ESTABLISHED]. This is a burden-and-safety gate operating upstream of the decision policy, and it is the closest existing analogue to a dedicated “should we act at all” stage. The decision mechanism developed across three generations: micro-randomisation for causal estimation of proximal effects; then action-centred contextual bandits, whose contribution is a linear contextual bandit variant permitting a complex baseline reward model alongside a simple treatment-effect model, decoupled through action centring (Greenewald et al., 2017); then online reinforcement learning personalised per user (Liao et al., 2020)

[ESTABLISHED]. What JITAI does not adapt is equally important. Notification content and timing constitute the entire action space. There is no layout, density, modality, navigation or visual hierarchy adaptation, adaptation occurs at a handful of discrete points per day rather than continuously, and there is no model of interface change cost because the interface does not change.

3.3 Adaptive user interfaces

SUPPLE formalised interface generation as a constrained optimisation problem: given a functional specification, device characteristics, available widgets and a user- and device-specific cost function, search a space of up to 10^{17} possible interfaces for the rendition minimising estimated user effort. Exact solutions were obtained in under a second in most cases, making runtime generation feasible (Gajos, Weld & Wobbrock, 2010) [ESTABLISHED]. SUPPLE++ extended this to individual motor capabilities with adjustment for vision. Generated interfaces significantly improved speed, accuracy and satisfaction for people with motor impairments relative to manufacturers' defaults [ESTABLISHED]. The authors were explicit that the system was not intended to replace designers but to serve users whose devices, abilities and tasks were not addressed by hand-crafted designs. Personalisation across this tradition operates on device parameters, abilities, task, preferences and usage traces. It does not use physiological signals. Todi, Bailly, Leiva and Oulasvirta (2021) introduced model-based reinforcement learning for adaptive UIs, simulating candidate sequences of adaptations and evaluating them with predictive HCI models, using Monte Carlo tree search with a pretrained value network to hold computation time approximately constant as search depth grows. The objective explicitly incorporates the negative effects of change: user surprise and relearning effort [ESTABLISHED]. This is the only mechanism in the surveyed corpus that models the cost of changing an interface over time, and it was demonstrated on menus.

3.4 The evidence that adaptation itself carries cost

This body of findings constrains everything that follows, and it is well evidenced. Adaptation rate can invert the benefit. Slow-paced adaptation of split menus produced benefits over a non-adaptive baseline,

whereas fast-paced adaptation produced negative results (Sears & Shneiderman, 1994; Findlater & McGrenere, 2004; synthesised in Lavie & Meyer, 2010) [ESTABLISHED]. Adaptation frequency determines the relative weight users place on its costs and benefits. Spatial memory is the mechanism. When frequently used items are moved rather than copied, frequent adaptation prevents users from forming spatial memory of the layout, degrading both performance and satisfaction relative to a static interface [ESTABLISHED]. Elective adaptation behaves differently from mandatory adaptation. Sears and Shneiderman's split menus copied predicted items rather than moving them, so users could scan the prediction or navigate to the known location; that structural difference separates the positive results from the negative ones [ESTABLISHED]. Accuracy is necessary but not sufficient: predictability contributes independently of accuracy (Gajos et al., 2008) [ESTABLISHED]. And even accurate adaptation costs something — highly accurate adaptive menus were efficient but reduced users' awareness of the full feature set, a measurable efficiency-versus-awareness trade-off (Findlater & McGrenere, 2008) [ESTABLISHED]. Small screens are not exempt. Low-accuracy adaptation performed poorly regardless of screen size, and small-screen participants relied more heavily on adaptive predictions than large-screen participants (Findlater & McGrenere, 2008) [ESTABLISHED]. Whether these findings transfer to watch-sized glanceable displays is [UNKNOWN]; no study located tested structural adaptation at that scale. A related caution: integrating human feedback into a reinforcement learning reward model for UI adaptation improved objective engagement but did not improve subjective engagement or user satisfaction (Gaspar-Figueiredo et al., 2025) [ESTABLISHED]. Metrics moving is not equivalent to users being better served.

3.5 Optimisation and multi-objective methods

Computational design treats interface generation as optimisation over predictive models of human perception, behaviour and experience (Oulasvirta, 2017). GRIDS demonstrated integer programming for grid layout, solving to high quality within seconds on commodity hardware, while noting that black-box multi-objective methods often return an unbalanced sample of the Pareto front (Dayama et al., 2020)

[ESTABLISHED]. Work on Pareto-optimal layouts for adaptive mixed reality is directly relevant to the temptation to build a single composite quality score. Its authors observe that prevailing approaches use global criterion methods such as weighted sums, which are inflexible, difficult to use and may fail to find preferred solutions; they propose generating a Pareto set instead, which supports fully automated, semi-automated or largely manual adaptation depending on how the set is resolved [ESTABLISHED]. The demonstration adapted the position of a single UI element. The implication for the architecture below [PROPOSED] is that optimisation is not a rival to policy learning. It defines the admissible candidate set within which a policy chooses, and a Pareto formulation preserves the option of handing the final choice to the user rather than scalarising it away.

3.6 Mixed-initiative interaction: deciding whether to act

Horvitz's (1999) principles of mixed-initiative user interfaces supply the mechanism absent from every physiological system above [ESTABLISHED]. The expected utility of autonomous action given evidence combines the utilities of four outcomes — acting or not, with the inferred goal correct or incorrect — weighted by their probabilities, producing a threshold probability for action. Context-dependent shifts in any of those utilities move the threshold: where erroneous action is costly, the system requires greater confidence before intervening. The accompanying principles govern override and recovery: introduce automation only where it adds value over direct manipulation; provide clear means to invoke or terminate it; minimise the cost of errors through natural rejection gestures; allow the degree of automation to be adjusted to match uncertainty; continue learning from observation; and let users refine suggestions rather than only accept or reject them. A 2025 scoping review of mixed-initiative visual analytics reports that only a narrow subset of these principles has been adopted in practice, and argues that safeguards minimising poor action and poor timing become more necessary as automation levels rise [ESTABLISHED].

3.7 Ambiguity and mediation

Where confidence is insufficient for autonomous action, mediation is the established alternative to

guessing. Dey and Mankoff (2005) define mediation as the dialogue between user and system that resolves ambiguous context, arguing that context-aware services routinely assume the context they consume is accurate whereas both sensing and interpretation are frequently ambiguous, and that correct handling will often require involving the user [ESTABLISHED]. The OOPS toolkit surveyed existing correction techniques for recognition-based interfaces and found they cluster into two strategies, repetition and choice (Mankoff, Abowd & Hudson, 2000) [ESTABLISHED]. Mediation has not been applied to physiological state inference; whether it transfers is [UNKNOWN]. Applying it there is [TRANSFERRED].

3.8 Change-point detection: transient state versus regime shift

Bayesian online changepoint detection maintains a posterior over run length — the time elapsed since the last changepoint — permitting real-time detection of genuine regime shifts as distinct from noise (Adams & MacKay, 2007), with classical antecedents in CUSUM and generalised likelihood ratio procedures and a modern review in Truong, Oudre and Vayatis (2020) [ESTABLISHED]. The method is established on physiological signals in clinical real time: detection of pathological state transitions in preterm infants with apnea and bradycardia; online segmentation of epileptic activity with a linear-complexity approximation developed specifically for real-time monitoring; change detection in heart-rate variability during sleep; and adaptive change detection in heart-rate trend monitoring in anaesthetised children (Yang, Dumont & Ansermino, 2006). It has not been used to gate an interface adaptation [UNKNOWN]. Doing so is [TRANSFERRED], and it is the cleanest transfer available in this chapter: the mechanism is mature, real-time capable, and validated on precisely the signal class a wearable produces.

3.9 Control theory and self-adaptive systems

Control-based self-adaptation provides guarantees that no interface mechanism above possesses. Filieri, Maggio and colleagues (2015) set out four properties analytically guaranteed from the mathematical definition of the controller: stability, meaning the setpoint is reached where possible; absence of overshoot; bounded settling time; and robustness to

model imprecision and to delayed or inaccurate measurement. Transient behaviour can be constrained so that the response does not oscillate around the setpoint [ESTABLISHED]. MAPE-K — Monitor, Analyze, Plan, Execute over shared Knowledge — is the reference control model for self-adaptive systems, with formal modelling and verification frameworks (Arcaini, Riccobene & Scandurra, 2017) [ESTABLISHED]. A known open problem is directly relevant here. Reconciling control-theoretic properties such as stability and settling time with the non-functional properties self-adaptive systems care about is explicitly identified as difficult and unresolved (Cámara et al., 2020) [ESTABLISHED]. That reconciliation has been posed for software properties. It has not been posed for human properties — spatial memory, predictability, awareness, surprise and relearning cost — anywhere we located [UNKNOWN].

3.10 Agentic AI over wearable and healthcare context

Large language model agents reasoning over wearable data are established at the highest level of publication. PHIA combines multistep reasoning, code generation and retrieval over behavioural health data, evaluated across benchmark datasets of over four thousand health-insight questions with a 650-hour expert evaluation, reporting 84% accuracy on objective numerical questions (Merrill et al., 2026) [ESTABLISHED]. Its action space is tool invocation and its output is language; it does not decide anything about an interface. Separately, a deep Q-learning approach combined with a metaheuristic optimiser has been proposed for adaptive user interfaces on wearable medical devices, personalising navigation flow, button placement and notification timing from real-time physiological and behavioural data. This is the nearest existing work to the present chapter's scope. Meta-control over policies is itself an active area, and it is dominated by reinforcement learning rather than language agents: deep RL meta-controllers routing between models on multi-objective criteria, and learned meta-policies over high-level deliberative actions [ESTABLISHED]. We located no instance of a language agent occupying an orchestration layer above deterministic adaptation controllers in any interface system [UNKNOWN].

IV. INTEGRATED ADAPTIVE ARCHITECTURE

4.1 Relation to MAPE-K

The proposed decomposition refines MAPE-K rather than replacing it [PROPOSED]. Sense and Infer correspond to Monitor and Analyze; Decide-whether, Select and Optimise to Plan; Apply to Execute; and Monitor-response and Learn feed Knowledge. Two refinements are made, each with independent mechanism backing:

- Decide Whether to Adapt is separated from planning, justified by Horvitz's expected-utility threshold, which is a decision distinct from choosing what to do. [TRANSFERRED]
- Apply Change Safely is separated from selection, justified by control-theoretic rate and settling-time bounds, which govern how a change is delivered independently of what was chosen. [TRANSFERRED]

Figure 1. Layered architecture for context-aware adaptation of wearable healthcare interfaces

Each stage carries an established mechanism; the composition is the proposed contribution.




Figure 1. Layered architecture for context-aware adaptation of wearable healthcare interfaces. Colour indicates provenance; the stage sequence and inter-stage relations are the proposed contribution.

4.2 The stages and their mechanisms

Each stage is populated by a mechanism drawn from an established tradition. The composition, and the assignment of timescales to stages, is what this chapter proposes.

V. MECHANISM MATRIX

Table 1 maps each architectural stage to its strongest established mechanism, its origin tradition, the objective it serves, its native timescale and its provenance status.

Table 1. Mechanism assignment across the adaptation lifecycle.

Stage	Mechanism	Origin	Objective	Timescale	Provenance
1. Sense	PPG, EEG, HRV, accelerometry	Physiological computing	Signal fidelity	ms–s	[ESTABLISHED]
2. Infer	Validated psychophysiological inference	Fairclough; Ewing et al.	Construct validity	s	[ESTABLISHED]
3. Confidence / Validation	Ambiguity modelling; mediation	Dey & Mankoff; OOPS	Resolve ambiguity, involve user	per event	[ESTABLISHED]; [TRANSFERRED] to physiology
4. Persistence / Change detection	Bayesian online changepoint (run length)	Adams & MacKay	Distinguish regime shift from noise	s–min	[ESTABLISHED] on physiology; [TRANSFERRED] to UI
5. Decide whether	Expected-utility threshold; availability gating	Horvitz; HeartSteps	Act only above threshold	per event	[ESTABLISHED]
6. Select	Contextual bandit, DQN, MCTS	HeartSteps; Todi et al.	Reward, minus cost of change	ms	[ESTABLISHED]
7. Optimise candidates	NSGA-III, integer programming, MOBO	GRIDS; Pareto MR layouts	Non-scalarised trade-offs	ms–s	[ESTABLISHED]
8. Apply safely	Rate limiting, settling time, overshoot bounds	Control theory	Stability, bounded transient	continuous	[ESTABLISHED]; never applied to UI
9. Monitor response	Micro-randomisation, proximal outcome	Klasnja et al.	Causal effect estimate	per decision point	[ESTABLISHED]
10. Learn	Online per-user RL	Liao et al.	Long-run reward	days–weeks	[ESTABLISHED]
10b. Recover	No mechanism located	—	—	—	[UNKNOWN]

5.1 Complementarity rather than competition

Read by timescale rather than by function, these mechanisms cease to be rivals. Physiological sensing supplies the state on which a policy acts rather than competing with it. Pareto optimisation defines the candidate space within which a policy chooses rather than replacing the policy. Control theory regulates how and when a selected change is applied rather than


selecting it. Todi's planning accounts for the cost of changing the interface over time, which none of the others do. The literature does not assemble this stack anywhere we located, but each layer's mechanism is independently established. The assembly is [PROPOSED].

VI. ADAPTATION LEVELS AND HUMAN-FACTOR CONSTRAINTS

Adaptation is not one thing. The evidence supporting it thins sharply as structural depth increases, and the defensible degree of automation falls with it.

Figure 3. The adaptation-level gradient

As adaptation becomes more structural, cost rises, evidence thins, and control shifts to the user.



Adaptation level	Evidence	Cost of change	Control mode
Content	Strong	Low	Automatic
Timing	Strong	Low	Automatic
Notification	Strong	Interruption	Automatic, gated
Information density	Moderate (VR)	Moderate	Assisted
Layout	Moderate, desktop	Spatial memory	Assisted
Component	Moderate	Awareness loss	Assisted
Navigation	Harm at speed	High	User-controlled
Modality	Insufficient	Mode errors	Not determined
Full structure	None located	Compounded	Not supported

[ESTABLISHED]
 [TRANSFERRED]
 [UNKNOWN]

Figure 3. The adaptation-level gradient. As adaptation becomes more structural, cost rises, evidence thins, and control shifts from the system to the user.

Table 2. Adaptation levels, their evidential support, and the defensible degree of automation at each.

Level	Mechanism precedent	Evidence	Cost of change	Timescale	Control mode
Content	JITAI tailoring variables	Strong	Low	Seconds	Automatic
Timing	Decision points, availability gating	Strong	Low	Minutes–hours	Automatic
Notification	Bandit arms; interruption management	Strong	Interruption	Event-driven	Automatic, gated
Information density	EEG-driven visual complexity; SYMBIOTIK	Moderate, not wearable	Moderate	Seconds–minutes	Assisted
Layout	SUPPLE, GRIDS	Moderate, desktop	Spatial memory	Slow or on request	Assisted
Component	SUPPLE++	Moderate	Awareness loss	Slow	Assisted

Level	Mechanism precedent	Evidence	Cost of change	Timescale	Control mode
Navigation	Todi et al.; Sci. Rep. 2025	Harm documented at speed	High	Slow, elective	User-controlled
Modality	None located	[UNKNOWN]	Mode errors	—	Not determined
Full structure	None located	[UNKNOWN]	Compounded	—	Not supported

The gradient is the practical design guidance this chapter offers [PROPOSED]. Sears and Shneiderman's copy-rather-than-move principle is the concrete implementation of the assisted mode [ESTABLISHED].

6.1 Constraints that must be treated as hard

- Adaptation rate, which can invert the sign of the benefit. [ESTABLISHED]
- Predictability, which contributes independently of accuracy. [ESTABLISHED]
- Stability and consistency of interface structure across glances. [ESTABLISHED]
- A minimum accuracy threshold, below which adaptation degrades performance at any screen size. [ESTABLISHED]
- Feature awareness, which even accurate adaptation erodes. [ESTABLISHED]

- Spatial memory preservation, the mechanism underlying glance-length interaction. [ESTABLISHED] on desktop; transfer to wearables [UNKNOWN]
- Elective rather than mandatory adaptation wherever structural change is involved. [ESTABLISHED]
- Validity of the psychophysiological inference, a precondition rather than a constraint. [ESTABLISHED]

VII. MULTI-SPEED ADAPTATION AND FAST CHANGE

The literature commonly treats adaptation speed as a single variable. We argue it is at least four, and that conflating them is what makes responsive adaptation appear incompatible with stable interaction.

Figure 2. The multi-speed adaptation model




Figure 2. The multi-speed adaptation model. The four upper stages may run at their native speed; the lower stages are governed by human-factor constraints.

Table 3. The four speeds, their mechanisms and their provenance.

Speed	Established mechanism	Status
Fast sensing	PPG, EEG at second-level resolution	[ESTABLISHED]
Fast inference	Classifiers; ~79% attention-state accuracy reported	[ESTABLISHED]
Fast decision	Bandits, DQN, pretrained value networks	[ESTABLISHED]
Fast candidate generation	Integer programming; layouts in seconds	[ESTABLISHED]
Deliberately slow structural change	Settling-time bounds, rate limiting, elective presentation	[ESTABLISHED] in control; [TRANSFERRED] to UI
Fast recovery	None located	[UNKNOWN]

The practical consequence [PROPOSED] is that fast sensing and fast decision are compatible with deliberately slow interface change, and the mechanisms for holding them apart already exist: control theory bounds rate of change and settling time,

and decision theory raises the threshold for action rather than accelerating it when the cost of error is high. A system may detect a state in milliseconds and legitimately decline to restructure the interface for minutes, or at all.

7.1 Stability mechanisms and their provenance

Mechanism	Original literature	Problem solved there	Status in adaptive UI
Hysteresis / dead-band	Control engineering	Prevent chatter around a switching threshold	[UNKNOWN] — not located
Cooldown / minimum dwell	Control, event systems	Bound actuation frequency	[UNKNOWN] — not located
Debounce / smoothing	Signal processing	Reject transient noise	[UNKNOWN] — not located
Confidence threshold	Decision theory; Horvitz (1999)	Act only above expected-utility crossover	[ESTABLISHED]
Change-point / run length	Adams & MacKay (2007)	Distinguish regime shift from noise	[ESTABLISHED] on physiology; [TRANSFERRED] to UI
Settling time, overshoot bounds	Filieri et al. (2015)	Bound transient behaviour with proof	[TRANSFERRED]
Model predictive control	Control theory	Optimise over receding horizon under constraints	Nearest analogue: Todi et al. MCTS planning
Mediation	Dey & Mankoff (2005)	Resolve ambiguous context with the user	[ESTABLISHED]; [TRANSFERRED] to physiological state

Several of these have never crossed into adaptive user interface research. That table is offered as a map of what remains available to be imported, not as a claim that importing them would succeed.

VIII. THE ROLE OF AGENTIC AI

The evidence does not support placing an agent throughout this architecture, and this chapter does not do so.

8.1 Where an agent is not justified

Not in the Apply stage. Stability guarantees are analytic and derive from the mathematical definition of the controller; substituting a language model forfeits the guarantee without an evidenced compensating benefit. Not in the Decide-whether gate. Horvitz's expected-utility threshold performs this function with better properties and lower computational cost. Not in Select. Contextual bandits and deep Q-networks demonstrably perform adaptation selection, including on wearable medical devices. We located no evidence that agentic reasoning improves adaptation decisions over these methods [UNKNOWN].

8.2 Where an agent may be justified

Three placements are supportable on present evidence: Heterogeneous context interpretation at the Infer stage. PHIA evidences reasoning over multi-source, irregular health data that a fixed feature vector represents poorly, as a capability to this role. Explanation of why an adaptation occurred. This is the only function uniquely demonstrated by language agents and absent from every adaptive system surveyed here, and Horvitz's principles arguably presuppose it. Mediation dialogue at the Confidence stage. Mediation requires a dialogue with the user to resolve ambiguous context, and the strategies available when it was formalised were repetition and choice. A language agent is a materially richer mediator. This is the single point where agentic capability addresses a documented limitation of an established mechanism rather than duplicating a function reinforcement learning already performs. This grading is deliberately conservative. The title of this chapter places agentic AI at the centre; the evidence places it at the periphery of the decision path and at the centre of interpretation and explanation. We have reported the evidence rather than the framing.

IX. THE PROPOSED CONCEPTUAL CONTRIBUTION

Stated precisely, so that it cannot be read as a stronger claim than it is: What is not claimed: that physiological adaptive interfaces are novel; that agentic AI outperforms reinforcement learning or control theory; that any component of the architecture is new; or that the integration has been empirically validated.

9.1 The recovery stage

The Learn/Recover stage is the architecture's weakest element and is presented as such. We located no work on safely reversing an interface adaptation after the inferred state proves transient or wrong, no work on returning to a prior interface state, and no work distinguishing a temporary physiological excursion from a durable change in user preference [UNKNOWN]. Mediation supplies the user-facing half of a possible solution; the system-side half has no precedent we located. This element is [PROPOSED] rather than synthesised, and it is the part of the architecture most in need of empirical work.

X. EMPIRICAL RESEARCH OPPORTUNITIES

Every mechanism in this chapter is imported from a non-wearable context. Physiological adaptation evidence comes from games, virtual reality, air traffic control and desktop dashboards. Adaptation-cost evidence comes from desktop and laptop menus. Stability guarantees come from server and infrastructure software. No element of the architecture has been demonstrated on a watch-sized glanceable display. The transfer of adaptation-cost findings to wearables is untested. Our reasoning that small screens and glance-length interaction increase the importance of spatial memory is an inference from adjacent findings, not a result. No validated glanceability metric was located, which limits the operationalisation of the construct most specific to this domain. Modality adaptation and accessibility in adaptive wearables were not adequately covered here. The accessibility omission is the more serious. AI-generated interfaces have been shown to achieve low WCAG compliance, and one study reports that explicitly specifying accessibility requirements in generation prompts decreased rather than improved compliance — a finding that contradicts an earlier

study in the same venue. This is **conflicting evidence** and is recorded as such. Adaptive behaviour and accessibility can conflict directly, and healthcare users are disproportionately exposed to the consequences. This chapter proposes no empirical validation and offers none. Its components are individually evidenced; their integration is not.

10.1 Four questions for empirical work

- Do the adaptation-cost findings from desktop menus hold at wearable scale, where interaction is glance-length and spatial memory carries more weight?
- Can control-theoretic stability properties be expressed in terms of human-factor constraints, extending the reconciliation problem already posed for software properties?
- What mechanism supports safe recovery when a physiological inference proves transient or incorrect?
- Does a language agent outperform existing mediation strategies at resolving ambiguous physiological context?

XI. CONCLUSION

The mechanisms required for context-aware adaptation of wearable healthcare interfaces exist. They are distributed across physiological computing, just-in-time adaptive interventions, adaptive user interface research, reinforcement learning, computational design, mixed-initiative interaction, context-aware computing and control theory — traditions that cite one another rarely and have not been assembled into a single account. This chapter has proposed one such assembly, organised around the proposition that the speed of sensing and decision-making should be separated from the speed and structural depth of interface change, and that established mechanisms exist to enforce that separation at every stage but one. Where a stage has no established mechanism, we have said so rather than filling the gap with assertion. The contribution is integrative. Its value, if it has any, lies in making visible a set of relations between mechanisms that already work, and in stating clearly which of those relations remain untested.

REFERENCES

Sources marked with a dagger (†) were assessed from abstract, indexed metadata or a secondary reference rather than full text, and require verification before submission. No bibliographic field has been inferred where it was not observed; fields that could not be confirmed are omitted rather than estimated.

- [1] Adams, R. P., & MacKay, D. J. C. (2007). Bayesian online changepoint detection. arXiv preprint [arXiv:0710.3742](https://arxiv.org/abs/0710.3742). <https://arxiv.org/abs/0710.3742>
- [2] Arcaini, P., Riccobene, E., & Scandurra, P. (2017). Formal design and verification of self-adaptive systems with decentralized control. *ACM Transactions on Autonomous and Adaptive Systems*. † <https://doi.org/10.1145/3019598>
- [3] Cámara, J., et al. (2020). Towards bridging the gap between control and self-adaptive system properties. arXiv preprint [arXiv:2004.11846](https://arxiv.org/abs/2004.11846). † Author list requires confirmation. <https://arxiv.org/abs/2004.11846>
- [4] Dayama, N. R., Todi, K., Saarelainen, T., & Oulasvirta, A. (2020). GRIDS: Interactive layout design with integer programming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3313831.3376553>
- [5] Dey, A. K., & Mankoff, J. (2005). Designing mediation for context-aware applications. *ACM Transactions on Computer-Human Interaction*, 12(1). <https://doi.org/10.1145/1057237.1057241>
- [6] Dey, A. K., Mankoff, J., Abowd, G. D., & Carter, S. (2002). Distributed mediation of ambiguous context in aware environments. In *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology (UIST '02)* (pp. 121–130). <https://doi.org/10.1145/571985.572003>
- [7] Ewing, K. C., Fairclough, S. H., & Gilleade, K. (2016). Evaluation of an adaptive game that uses EEG measures validated during the design process as inputs to a biocybernetic loop. *Frontiers in Human Neuroscience*, 10, 223. <https://doi.org/10.3389/fnhum.2016.00223>
- [8] Fairclough, S. H. (2009). Fundamentals of physiological computing. *Interacting with Computers*, 21, 133–145. † Issue number not confirmed.

- [9] Fairclough, S. H. (n.d.). Physiological computing [Author blog]. † Non-peer-reviewed source; cited only for the author's own characterisation of the field. <https://physiologicalcomputing.org/>
- [10] Filieri, A., Maggio, M., et al. (2015). Software engineering meets control theory. In Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS). † Full author list requires confirmation. <https://misailo.web.engr.illinois.edu/papers/2015-seams.pdf>
- [11] Findlater, L., & McGrenere, J. (2004). A comparison of static, adaptive, and adaptable menus. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04) (pp. 89–96). <https://doi.org/10.1145/985692.985704>
- [12] Findlater, L., & McGrenere, J. (2008). Impact of screen size on performance, awareness, and user satisfaction with adaptive graphical user interfaces. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08). https://www.cs.ubc.ca/~joanna/papers/CHI2008_Findlater.pdf
- [13] Gajos, K. Z., Everitt, K., Tan, D. S., Czerwinski, M., & Weld, D. S. (2008). Predictability and accuracy in adaptive user interfaces. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08). http://aiweb.cs.washington.edu/ai/puirg/papers/k_gajos-chi08-predictability.pdf
- [14] Gajos, K. Z., Weld, D. S., & Wobbrock, J. O. (2010). Automatically generating personalized user interfaces with SUPPLE. Artificial Intelligence, 174(12–13), 910–950. <https://www.sciencedirect.com/science/article/pii/S0004370210000822>
- [15] Gaspar-Figueiredo, D., et al. (2025). A comparative study on reward models for user interface adaptation with reinforcement learning. Empirical Software Engineering. † Full author list requires confirmation. <https://doi.org/10.1007/s10664-025-10659-5>
- [16] Greenewald, K., Tewari, A., Klasnja, P., & Murphy, S. (2017). Action centered contextual bandits. In Advances in Neural Information Processing Systems 30. https://proceedings.neurips.cc/paper_files/paper/2017/file/4fa177df22864518b2d7818d4db5db2d-Paper.pdf
- [17] Horvitz, E. (1999). Principles of mixed-initiative user interfaces. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99) (pp. 159–166). <https://doi.org/10.1145/302979.303030>
- [18] Klasnja, P., Smith, S., Seewald, N. J., Lee, A., Hall, K., Luers, B., Hekler, E. B., & Murphy, S. A. (2019). Efficacy of contextually tailored suggestions for physical activity: A micro-randomized optimization trial of HeartSteps. Annals of Behavioral Medicine, 53(6), 573–582. <https://doi.org/10.1093/abm/kay067>
- [19] Lavie, T., & Meyer, J. (2010). Benefits and costs of adaptive user interfaces. International Journal of Human-Computer Studies. † Volume and pages require confirmation. http://www.cs.tufts.edu/~jacob/250aui/Adaptive_Benefits_Lavie_IJoHCI10.pdf
- [20] Liao, P., Greenewald, K., Klasnja, P., & Murphy, S. (2020). Personalized HeartSteps: A reinforcement learning algorithm for optimizing physical activity. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 4(1), 1–22. <https://doi.org/10.1145/3381007>
- [21] López-Cardona, A., Attygalle, N., et al. (2025). Context-aware adaptive visualizations for critical decision making (SYMBIOTIK). In Proceedings of the European Conference on Artificial Intelligence (ECAI 2025), Demo Track. † Demonstration-track paper; full author list requires confirmation. <https://doi.org/10.3233/FAIA251433>
- [22] Mankoff, J., Abowd, G. D., & Hudson, S. E. (2000). OOPS: A toolkit supporting mediation techniques for resolving ambiguity in recognition-based interfaces. Computers & Graphics, 24(6), 819–834.
- [23] Merrill, M. A., Paruchuri, A., Rezaei, N., et al. (2026). Transforming wearable data into personal health insights using large language model agents. Nature Communications, 17, 1143. <https://doi.org/10.1038/s41467-025-67922-y>
- [24] Oulasvirta, A. (2017). User interface design with combinatorial optimization. Computer, 50, 40–47. † Issue number not confirmed.

- [25] Oulasvirta, A., De Pascale, S., Koch, J., Langerak, T., Jokinen, J., Todi, K., Laine, M., Kristhombuge, M., Zhu, Y., Miniukovich, A., Palmas, G., & Weinkauff, T. (2018). Aalto Interface Metrics (AIM): A service and codebase for computational GUI evaluation. In UIST '18 Adjunct Proceedings (pp. 16–19). <https://doi.org/10.1145/3266037.3266087>
- [26] Pope, A. T., Bogart, E. H., & Bartolome, D. S. (1995). Biocybernetic system evaluates indices of operator engagement in automated task. *Biological Psychology*, 40, 187–195.
- [27] Prinzel, L. J., Freeman, F. G., Scerbo, M. W., Mikulka, P. J., & Pope, A. T. (2000). A closed-loop system for examining psychophysiological measures for adaptive task allocation. *The International Journal of Aviation Psychology*, 10(4), 393–410.
- [28] Schaule, F., Johanssen, J. O., Bruegge, B., & Loftness, V. (2018). Employing consumer wearables to detect office workers' cognitive load for interruption management. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(1), 32.
- [29] Sears, A., & Shneiderman, B. (1994). Split menus: Effectively using selection frequency to organize menus. *ACM Transactions on Computer-Human Interaction*. † Volume, issue and pages require confirmation.
- [30] Todi, K., Bailly, G., Leiva, L. A., & Oulasvirta, A. (2021). Adapting user interfaces with model-based reinforcement learning. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3411764.3445497>
- [31] Truong, C., Oudre, L., & Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing*, 167, 107299.
- [32] Yang, P., Dumont, G., & Ansermino, J. M. (2006). Adaptive change detection in heart rate trend monitoring in anesthetized children. *IEEE Transactions on Biomedical Engineering*, 53(11), 2211–2219.
- [33] Adaptive user interfaces for wearable medical devices using deep Q-learning and Golden Jackal Optimization. (2025). *Scientific Reports*. † Assessed from abstract and indexed metadata only. Author list, volume, article number and validation method all require confirmation from the full text; see verification note in Section 3.10. <https://www.nature.com/articles/s41598-025-28937-z>
- [34] Bayesian online changepoint detection of physiological transitions. (2018). † Author list and venue require confirmation. PubMed ID 30440337. <https://pubmed.ncbi.nlm.nih.gov/30440337/>
- [35] Generated inaccessible: Measuring WCAG violations in AI UI design tools. (2026). In *Proceedings of the 23rd International Web for All Conference (W4A)*. † Author list requires confirmation. <https://doi.org/10.1145/3800424.3800430>
- [36] When LLM-generated code perpetuates user interface accessibility barriers, how can we break the cycle? (2025). In *Proceedings of the 22nd International Web for All Conference (W4A)*. † Author list requires confirmation. <https://doi.org/10.1145/3744257.3744266>
- [37] Pareto optimal layouts for adaptive mixed reality. (2023). In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. † Author list requires confirmation. <https://doi.org/10.1145/3544549.3585732>
- [38] A scoping review of mixed-initiative visual analytics in the automation renaissance. (2025). arXiv preprint arXiv:2509.19152. † Author list requires confirmation. <https://arxiv.org/abs/2509.19152>
- [39] Real-time cognitive load measurement using wearable photoplethysmogram sensor. *ACM Transactions on Applied Perception*. † Authors, year, volume and pages require confirmation. <https://doi.org/10.1145/3779293>
- [40] Nahum-Shani, I., et al. Just-in-time adaptive interventions (JITAIs): Six core elements. † Cited here from secondary sources only; the primary reference must be obtained and cited directly before submission.