Agentic Wearables: Sensor-Informed Intelligent Systems for Autonomous Health Support

I envision a future where intelligent systems provide us with just-in-time, personalized health support during everyday life. Advances in wearables, smartphones, and other mobile health (mHealth) technologies provide an unprecedented opportunity to complement clinical care and fill a critical gap in health care delivery. Sudden bouts of symptoms (e.g., panic attack) and the need for continual behavior change are hallmarks of mental and chronic illnesses, requiring timelier support than can be clinically initiated.[1] Yet, much of the latest work on AI for wearables focuses on interpreting sensor data for routine clinical support,[2] leaving the gap in autonomous everyday support unaddressed.

My research addresses this gap by enabling intelligent systems that *autonomously* provide health support, leveraging *sensor data* and *online learning* for personalization. A difficult yet remarkable aspect of health care delivery is that the same exact intervention can produce vastly different outcomes when delivered in different contexts or to different people. My research addresses this need for precision and personalization by **focusing on three components of AI systems for autonomous mHealth support (Figure 1)**. Sensor data can be used to infer an individual's biobehavioral context or "state." Digital twin models can accordingly predict the effects of intervention decisions on outcomes. Decision-making algorithms can then optimally control interventions to improve outcomes. For each component, online learning enables individualized training on incoming data to optimize future inference and decision making.

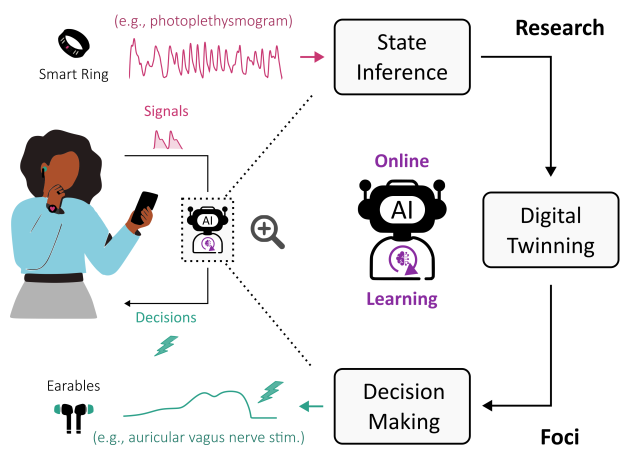

Figure 1: Research program on "agentic wearables." My program's three foci enable the next generation of sensor-informed closed-loop systems for mHealth support.

Three intertwined **c**hallenges motivate my past contributions and lab's future research directions:

C1) Modeling and inference errors can cause personalized support to backfire. If a system suggests a person de-stress when they are in fact relaxed, confidence in the system can reduce and lead to disengagement. Wearable sensor data is often corrupted, and predictions of states like stress are fallible. Autonomous systems must account for these uncertainties to personalize support intelligently.

C2) Quantitative models to inform decision making are lacking. Mechanistic models in behavioral science are often qualitative, while data-driven models often ignore person-specific differences and struggle notoriously with group-to-individual generalizability. Models are needed that account for between-person heterogeneity and leverage domain knowledge to overcome blindspots in data.

C3) Current control systems focus either on behavioral or biological interventions, each with distinct modes of operation. Behavioral interventions (e.g., mindfulness reminders) require conscious engagement, which is impractical if suggested too often or in moments like panic. Biological interventions (e.g., neuromodulation) bypass consciousness but are riskier and do not promote self-efficacy.

To address C1, C2, and C3, my lab will build upon the foundation laid by my prior work and pursue three specific research **d**irections: **D1) uncertainty-informed, online decision making** for precision health support, **D2) hybrid modeling for digital twin design** informed by domain science, and **D3) optimal control of biobehavioral intervention systems** to enable the next generation of autonomous systems for personalized health support – accessible in everyday life.

[1] M Ghahramanlou-Holloway, *Psychiatry*, 2022 [2] D A Adler et al., *ACM IMWUT*, 2024

Past Contributions

My past contributions lay the groundwork to address challenges C1, C2, and C3. I designed online biosignal processing and state inference pipelines that automatically assess signal quality (i.e., measurement uncertainty); modeled person-specific stress dynamics in response to non-invasive vagus nerve stimulation (nVNS), a nascent biological intervention; and developed a digital twin framework to optimize reinforcement learning (RL) algorithms for just-in-time adaptive interventions (JITAIs), nascent closed-loop mHealth systems. Through team science, my use-inspired research has also led to real-world impact, including **Food and Drug Administration Breakthrough Designation of nVNS** for posttraumatic stress disorder (PTSD) based on my team's findings.

Dynamic Modeling and Digital Twin Simulations of Closed-Loop nVNS

I produced the first-ever simulations of closed-loop nVNS's effects on stress (Figure 2), modeling the dynamics of physiological responses to nVNS at the temporal resolution of seconds [J34] (lettered citations refer to my CV). nVNS electrically stimulates a "rest and digest nerve" to counteract the "fight or flight" response. Prior work investigated nVNS effects over minutes or hours, but understanding the time course of shorter-term effects is critical to delivering nVNS in closed-loop.

My key insight was abstracting nVNS as the input and physiological markers as the outputs of a dynamical system. The system's latent states then capture changes shared across markers. My person-specific models uncovered nVNS effects within 10-15 s and can predict physiological changes 20 s into the future [J1], [J34].

For this work, I published two journal articles and four conference papers as lead author [J1], [J34], [C1], [C10], [C12], [C16]. My work was awarded an IEEE BHI 2019 Best Poster Award, a NYC Neuromodulation 2020 Outstanding Presentation Award, and the InterfaceRice 2023 Best Lightning Talk Award.

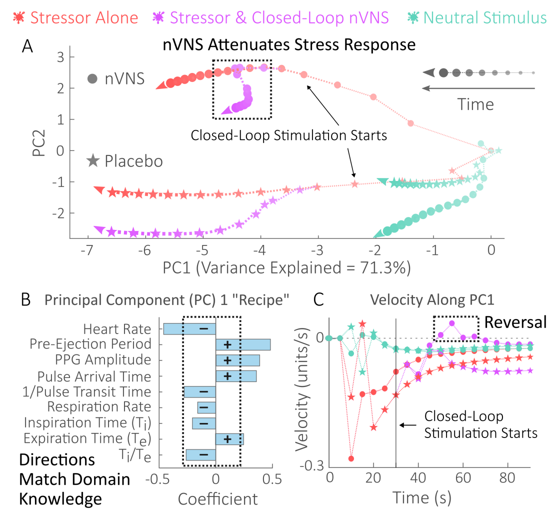

Figure 2: First-ever simulations of closed-loop nVNS. Latent state dynamics are projected into a principal component (PC) space. Two sets of responses are shown for patients with PTSD: a nVNS group and a placebo group (double-blind clinical trial). **A)** Closed-loop nVNS attenuates the stress response. **B)** PC1 lies in the orthant stress psychophysiology suggests. **C)** Closed-loop nVNS reverses the stress response briefly.

Digital Twins to Optimize RL Algorithms and Continually Improve JITAIs

I created the JITAI-Twin framework, the first instantiation of the digital twins of a subpopulation concept in health.[3] JITAI-Twins synthesize the need to deploy a JITAI multiple times for continual improvement with the need to select between and optimize RL algorithms prior to each deployment.

My key insight was recognizing that while a RL algorithm may yield person-specific policies, a meta-layer of optimization is needed to tune the algorithm itself for the entire target subpopulation. Tuning is guided by JITAI-Twin simulations, and data is used to continually improve the twin.

For this work, I am the lead author of a journal article and a manuscript in revision [J39], [J43].

Biology-Informed Methods for Measurement Uncertainty and State Inference

I developed state inference methods for cardiovascular and respiratory signals, using physiology-informed signal quality assessment (SQA) to quantify measurement uncertainty (e.g., signal corrup-

[3] NASEM, *Opportunities and Challenges for Digital Twins in Biomedical Sciences: Proceedings of a Workshop*, 2023

tion). Standard SQA uses outlier detection or accelerometry, only detecting artifacts due to motion.

My key insight was leveraging the quasi-periodicity of heartbeats and breaths for SQA. Assuming stationarity for 10-20 s, quality respiratory signals have concentrated power spectra [J5], [C7]. Heartbeat segments can be compared against templates formed using past high-quality segments [J17], [C4]. My SQA methods improve inference, especially in real-world settings with low signal-to-noise ratio (Figure 3).

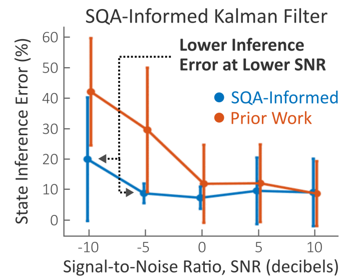

Figure 3: Improved state inference using SQA to inform measurement noise covariance.

For this work, I published two journal articles and four conference papers as lead author and won an IEEE BHI 2021 **Best Paper Award** [J5], [J24], [C2], [C3], [C4], [C7]. I also produced an open-source toolbox, used for example by collaborators at the Emory School of Medicine.

Future Directions

My prior work demonstrates expertise in (1) signal processing and state inference for wearable sensor data; (2) dynamic modeling and digital twin design for biological and behavioral interventions; and (3) closed-loop decision making with model-based control and RL for mHealth systems. Building on this, my lab will address challenges C1, C2, and C3 by pursuing directions D1, D2, and D3.

D1) Uncertainty-Informed State Inference and RL Algorithms for JITAIs

An immediate direction of my lab will be to develop uncertainty-informed state inference and RL algorithms for JITAIs (Figure 4). JITAIs thus far have ignored uncertainty quantification (UQ) when inferring state. RL algorithms that account for state uncertainty (e.g., POMDPs) are impractical for JITAIs.

Our first aim will be to develop UQ methods tailored to supervised learning with wearable sensor data. We will enhance state-of-the-art UQ methods by using SQA to inform data UQ (i.e., aleatoric). Our second aim will be to develop contextual bandit algorithms that account for uncertainty in state. We will investigate interventions that improve data quality to target aleatoric uncertainty. An extension we will research is UQ for JITAIs with large language model (LLM) agents. LLMs implicitly infer state when personalizing interactions, but without principled UQ for LLM inference, challenge C1 returns.

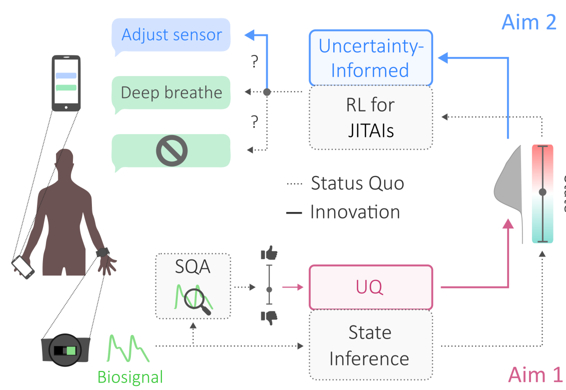

Figure 4: Uncertainty-informed JITAIs. Aim 1 is to design UQ methods leveraging SQA, and Aim 2 is to develop uncertainty-informed RL for JITAIs.

To fund this research, I have secured over $700k for my lab's first three years from the National Institute of Biomedical Imaging and Bioengineering. My lab will continue my collaborations with behavioral scientists at the Utah Center for Health Outcomes and Population Equity to deploy our UQ and RL algorithms for a smoking cessation JITAI.

D2) Hybrid Modeling for Psychophysiology-Informed Digital Twin Design

Inspired by physics-informed machine learning,[4] another immediate direction of my lab will be to develop models for digital twins that incorporate psychophysiological insight. The challenge is that differential equations and other conservation principles are often unknown in psychophysiology.

Our first aim will be to create what I call "systems physiology-informed neural networks

[4] G Karniadakis et al., *Nature Reviews Physics*, 2021

(SPINNs)" for applications such as autonomic nervous system (ANS) modeling for stress. While differential equations for ANS changes are unknown, causal directed acyclic graphs are known with "+" or "−" symbols for known monotonicity (Figure 2B). This can inform architecture and loss functions (e.g., penalize disparate gradients). Our second aim will be to train models such as SPINNs via decision-focused learning (DFL) for use in digital twins. Digital twins are often trained to predict future states, then used for decision making. However, better predictions do not imply better prediction-informed decisions. DFL focuses on decision quality, rather than intermediary prediction.

To fund this research, I will first target the National Science Foundation SCH program, which supports computational research informed by biomedical sciences for health innovation. My lab will continue my collaborations with Ziping Xu at UNC-Chapel Hill and Emre Ertin at Ohio State University on pretraining JITAI-Twin models with big data (e.g., All of Us dataset).

D3) Optimal Online Decision Making for Biobehavioral Intervention Systems

A longer-term direction of my lab will be to develop modeling and control methods for next generation mHealth systems that use behavioral *and* biological interventions, analogous to treatment regimens combining pharmacological and behavioral interventions that improve outcomes more than either alone.[5] How to best intervene biologically and/or behaviorally is an optimal control problem.

Our first aim will be to unify input-output modeling methods for improved optimization experiments for biobehavioral intervention systems. We will investigate the connection between persistent excitation for system identification and statistical power for micro-randomized trials in RL for JITAIs. Our second aim will be to devise reward functions for adaptive optimal controllers (note that RL can be viewed as adaptive optimal control) of combined mindfulness and nVNS intervention systems. Reward functions will include standard user burden and nVNS energy consumption terms, as well as user preference terms. Optimality must factor in how users prefer each form of support.

To fund this research, I will first target the Defense Advanced Research Projects Agency's (DARPA's) programs similar to STRENGTHEN, which supports the design of multimodal, multidimensional interventions for PTSD. I have interacted with DARPA program managers in the past and helped secure funding for a DARPA NEAT project. My lab will continue my collaborations with Conor Walsh at Harvard on wearable and mobile intervention systems that combine behavioral interventions, neuromodulation, and robotic support for stroke rehabilitation.

Conclusion

I am an engineer with interdisciplinary expertise in designing autonomous mHealth systems that personalize biobehavioral interventions to biobehavioral sensor feedback. As one of 32 Schmidt Science Fellows selected from around the world in 2023, I pivoted for my postdoc to focus on online decision making for JITAIs, complementing my PhD on modeling and state inference with wearables. I will leverage this unique expertise to lead a lab on "agentic wearables," enabling just-in-time health support in our everyday lives – *when and how we need it.*

[5] P Cuijpers et al., *World Psychiatry*, 2020