Modeling Latent Dynamics of the Autonomic Nervous System in Response to Trauma Recall and Non-Invasive Vagus Nerve Stimulation

Asim H. Gazi, *Member, IEEE*, Sungtae An, Jesus Antonio Sanchez-Perez, Michael Chan, Mohammad Nikbakht, David J. Lin, Shlok Natarajan, Kyle A. Johnsen, J. Douglas Bremner, Jin-Oh Hahn, *Senior Member, IEEE*, Omer T. Inan, *Fellow, IEEE*, Christopher J. Rozell, *Senior Member, IEEE*

***Abstract*— Objective: To develop a person-specific dynamic modeling approach to quantify the time course of multiple peripheral markers of the autonomic nervous system (ANS) in response to acute stressors and potential stress-reducing interventions.**

Methods: We curated data (N=50 participants) from a double-blind, randomized, sham-controlled trial of non-invasive vagus nerve stimulation (nVNS) for posttraumatic stress disorder (PTSD). For each participant, a multi-input, multi-output, linear state space model (SSM) was trained on ∼3500 s of cardiovascular and respiratory marker time series data and tested for predictive validity on ∼2000 s of held-out data. The inputs to each SSM indicated the presence or absence of potential stressors (e.g., traumatic memories) and interventions (e.g., nVNS or sham stimulation). We analyzed the SSMs' step responses to each input and compared the responses to real data. We then simulated the effects of just-in-time nVNS delivered during a traumatic memory.

Results: The SSMs outperformed baseline forecasting methods on held-out data (P<.05). Responses to nVNS were in the opposite direction of responses to traumatic memories, with neutral conditions (included for comparison) remaining in between. For participants with PTSD, just-in-time nVNS attenuated–and briefly reversed–the response to traumatic memories along the principal axis of variance, which explained ∼50% or more variance and mirrored expected ANS changes. Just-in-time sham stimulation produced no attenuation.

Conclusion: Our methods capture latent ANS dynamics during increasing and decreasing stress levels associated with traumatic memories and nVNS, respectively.

Significance: Traumatic memories can cause pathological stress responses during daily life that just-in-time non-invasive neuromodulation may potentially help mitigate.

***Index Terms*— dynamic modeling, multimodal fusion, personalized, stress, unsupervised learning**

Manuscript submitted on February 9, 2025. A. H. Gazi was supported by a National Science Foundation Graduate Research Fellowship (DGE-2039655) and is currently supported by Schmidt Science Fellows, in partnership with the Rhodes Trust. This research was supported by the Defense Advanced Research Projects Agency Biological Technologies Office Targeted Neuroplasticity Training program and ElectRx program through the Naval Information Warfare Center (N66001-16-2-4054 and N66001-19-2-4002) and the Office of Naval Research (N00014-21-1-2031). The active and sham neuromodulation devices used in this study were provided free of charge by electroCore.

A. H. Gazi (agazi@schmidtsciencefellows.org) is with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. He was previously with the School of Electrical and Computer Engineering (ECE), Georgia Institute of Technology (GT), Atlanta, GA, USA.

S. An and S. Natarajan were with the School of Interactive Computing, GT, Atlanta, GA, USA.

J. A. Sanchez-Perez, M. Nikbakht, D. J. Lin, O. T. Inan, and C. J. Rozell are with the School of ECE, GT, Atlanta, GA, USA.

M. Chan, K. A. Johnsen, and O. T. Inan are with the Coulter Department of Biomedical Engineering, GT, Atlanta, GA, USA.

J. D. Bremner is with the Department of Psychiatry and Behavioral Sciences and the Department of Radiology and Imaging Sciences, Emory School of Medicine, Atlanta, GA, USA. J. D. Bremner is also with the Atlanta Veterans Affairs Health Care System, Decatur, GA, USA.

J.-O. Hahn is with the Department of Mechanical Engineering, University of Maryland, College Park, MD, USA.

I. Introduction

THE autonomic nervous system (ANS) is the primary driver of the body's short-term (i.e., acute) stress response. Acute stress involves latent changes in ANS activity that – within seconds – cause observable physiological changes throughout the body [1]. These observable changes include variations in cardiovascular, respiratory, and perspiratory function, measurable using physiological sensing [2]. Physiological sensors and other technologies that interface with the ANS have the potential to improve quality of life for many. For example, individuals with posttraumatic stress disorder (PTSD) face episodes of pathological stress that current healthcare systems struggle to manage [3]. These stress episodes can be triggered by traumatic memories and have been shown to induce changes in several peripheral physiological markers including heart rate, sweat production, cardiac contractility, peripheral vasoconstriction, respiratory timings, and certain markers' variabilities [4]–[7]. Many of these stress-induced changes are exacerbated or blunted in PTSD [4], [8], [9]. To better manage episodes of acute stress, stress-induced changes in ANS activity could be monitored and, when necessary, counteracted during daily life for just-in-time support.

Non-invasive neuromodulation technologies are one class of interventions with the potential to counteract stress-induced changes in ANS activity during daily life [3]. Transcutaneous cervical vagus nerve stimulation (tcVNS), non-invasive electrical stimulation of a parasympathetic ("rest and digest") nerve to counteract the sympathetic ("fight or flight") response of the ANS, has been investigated longitudinally for PTSD in daily life settings and has been shown to reduce the effects

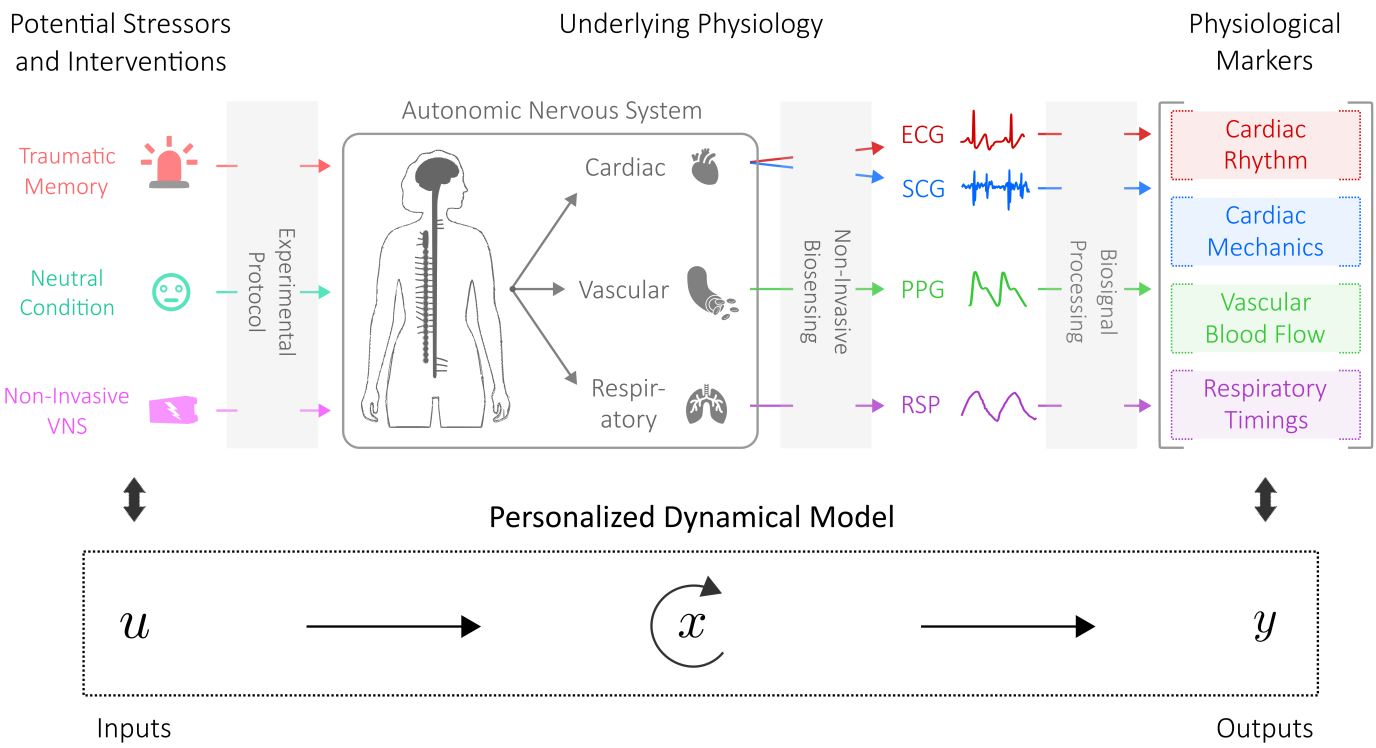

Fig. 1. Mapping between psychophysiological context and a personalized dynamical model. Potential interventions such as non-invasive vagus nerve stimulation (VNS) can be modeled as inputs, $\boldsymbol{u}$, to an underlying dynamical system. This system's latent state, $\boldsymbol{x}$, represents underlying changes in autonomic nervous system activity and varies both autonomously and in response to exogenous inputs. Variations in $\boldsymbol{x}$ produce measurable physiological changes, including cardiovascular and respiratory changes mediated by the autonomic nervous system. Via biosensing and biosignal processing, these changes can be quantified using physiological markers, modeled as the outputs, $\boldsymbol{y}$, of the dynamical system.

of stress on peripheral physiological markers of ANS activity [10]–[12]. Prior work by Gurel et al. demonstrated that static snapshots (e.g., minute-long averages) of markers such as heart rate were reduced minutes or hours after tcVNS was administered. Our recent dynamic modeling efforts built on the work of Gurel et al. by constructing person-specific, single-input, single-output state space models to capture the dynamics of heart rate and, separately, a vascular marker of the ANS in response to tcVNS [13]. Our person-specific approach accounted for heterogeneous effects of both VNS on the ANS and the ANS on peripheral physiology [1], [14]. We thereby addressed the need to characterize the timecourse of ANS responses to tcVNS. This led to findings such as the detection of tcVNS effects on peripheral physiology within 15-20 seconds.

Although our recent work was an important step in dynamically modeling ANS responses to tcVNS [13], two limitations remain. First, our prior work only considered cardiovascular markers of the ANS and modeled changes in each of the two markers separately. The ANS dictates multiple physiological changes throughout the body that all covary during changes in stress [1]. To better quantify changes in ANS activity, the dynamics of multiple physiological markers should be modeled simultaneously. Second, our prior work compared tcVNS against sham stimulation but did not model the effects of stressors in comparison to tcVNS, leaving the magnitude of tcVNS's effects with respect to stressors unknown. Comparisons between stressors and potential interventions are needed to quantify how much of a counteracting effect that a potential intervention can have in the context of such stressors.

The objective of the current work is to advance the state of the art by devising a novel dynamic modeling approach for multiple physiological markers of the ANS in response to acute stressors and non-invasive neuromodulation. Specifically, we use latent state space models to characterize the covarying dynamics of cardiovascular and respiratory markers in response to traumatic memories and tcVNS (Fig. 1). We use latent state space models to mirror the psychophysiological context of interest. Changes in peripheral physiological markers are observable manifestations of changes in underlying or latent ANS activity [1]. Our approach is related to work on latent variable and dynamic models of acute stress physiology [15]–[17]. Prior state space models of ANS activity have primarily considered cardiovascular markers or perspiratory markers on their own [18]–[22]. A few nascent methods better quantify latent ANS dynamics in response to stressors by considering both cardiovascular and perspiratory markers [23]–[25]. However, no work to date has modeled the exogenous impact of both stressors and stress-reducing interventions (needed to account for intervention effects during state estimation), nor the covarying dynamics of cardiovascular and respiratory markers.

In this paper, we detail our person-specific, multiple-input, multiple-output (MIMO) state-space modeling approach that we validate on data from 50 traumatized individuals, 24 with PTSD. We show that our modeling approach satisfies established validation criteria and produces models that are significantly better at predicting physiological dynamics up to

20 s into the future compared to baseline forecasting methods. We then use the MIMO state-space models to quantify the effects of tcVNS on latent ANS dynamics and compare these effects against those of traumatic memories and a neutral condition (for reference). We thereby address the limitations of prior work and make the following contributions:

1) Validate a personalized dynamical modeling approach for multiple cardiovascular and respiratory markers of the ANS

2) Elucidate the opposing effects of tcVNS and traumatic memories on the dynamics of multiple markers of the ANS

3) Demonstrate that multiple markers of the ANS covary in response to stressors and interventions, and this covariation can be principally captured in a lower dimensional subspace

II. Methods

A. Study Cohort and Protocol

The data we analyzed in this work was curated from a double-blind, randomized, sham-controlled trial of tcVNS for individuals with histories of psychological trauma (ClinicalTrials.gov NCT0299289). Institutional Review Board (IRB) approval was obtained from the Emory University School of Medicine (#IRB00091171), Georgia Institute of Technology (#H17126), SPAWAR Systems Center Pacific, and Department of Navy Human Research Protection Program. From the $N = 50$ participants in the analytical sample, 24 had posttraumatic stress disorder (PTSD) at the time of the trial. Thirteen participants with PTSD received active tcVNS throughout the study, and 11 received sham stimulation. From the 26 participants without PTSD, 12 received active stimulation, and 14 received sham stimulation. The sample and portion of the trial we analyzed match that of our prior work [26].

The electrocardiogram (ECG), photoplethysmogram (PPG), seismocardiogram (SCG), and respiratory effort (RSP) signals were recorded throughout the protocol at a 2 kHz sampling rate. The protocol included trauma recall, neutral conditions, and active or sham stimulation, as shown in Fig. 2A. During trauma recall, participants listened to 60-s audio recordings via headphones that each detailed an aspect of one of their prior traumatic events, following established methods [27]. For neutral conditions, participants listened to 60-s audio recordings describing pleasant scenery meant to elicit neutral or positive affect, following established methods [27]. For each administration of stimulation, participants received 120 s of either active tcVNS or sham stimulation. Further details are provided in the supplement.

B. Extraction of Physiologically Meaningful Markers

A total of **nine peripheral markers** of the ANS were extracted from each participant's ECG, RSP, PPG, and SCG data using the signal quality assessment and processing pipelines of prior work [26], [28]. Cardiovascular markers were extracted on a beat-by-beat basis (i.e., one value per heartbeat if signal quality was satisfactory). From the ECG, **heart rate (HR)** was extracted [26]. HR is an extensively studied marker that generally increases during acute stress [2]. From the PPG, **PPG amplitude ($\mathbf{PPG_{amp}}$)** was extracted [28]. $\mathrm{PPG_{amp}}$ is a vascular marker that varies as a function of vasoconstriction and pulse pressure. $\mathrm{PPG_{amp}}$ generally decreases during acute stress when measured at the periphery due to vasoconstriction [11]. From the ECG and PPG, **pulse arrival time (PAT)** was extracted [28]. PAT is a cardiovascular timing that varies as a function of vascular stiffness and the heart's ejection process, generally decreasing during acute stress [12], [29]. From the ECG and SCG, **pre-ejection period (PEP)** was extracted [28]. PEP is a cardiac timing inversely related to cardiac contractility, generally decreasing during acute stress [11], [29]. Using the ECG, PPG, and SCG, the **reciprocal of the pulse transit time (1/PTT)** was extracted. 1/PTT is a correlate of blood pressure and generally increases during acute stress [29].

Respiratory markers were extracted on a breath-by-breath basis. Using the RSP signal along with two ECG-derived respiration signals, **respiration rate (RR)**, **inspiration time ($\mathbf{T_i}$)**, **expiration time ($\mathbf{T_e}$)**, and the **$\mathbf{T_i/T_e}$ ratio** were extracted [26]. RR is an extensively studied marker that tends to increase during acute stress [2]. $\mathrm{T_i}$ is the time it takes for inhalation, and $\mathrm{T_e}$ is the time it takes for exhalation, which is the phase of breathing where vagal tone is greater. $\mathrm{T_e}$ generally decreases during acute stress [30]. $\mathrm{T_i/T_e}$ is the inspiration to expiration ratio and generally increases during acute stress [30].

C. Time Series Preparation for Dynamic Modeling

In preparation for dynamic modeling, the physiological marker (i.e., "physiomarker") time series were processed to produce output vectors, $\{y_k\}_{k=0}^{T_{total}}$, for each participant-specific model shown in Fig. 2B. Before and after each of the following time series preparation steps, the time series were inspected to ensure desired results. For each participant's data, each physiomarker time series was resampled to 0.2 Hz with an anti-aliasing filter and time-aligned via linear interpolation. We selected 0.2 Hz to align with adult respiratory rates [31], preferring cardiovascular markers be downsampled [32]. This ensures that subsequent dynamic modeling focuses on the latent ANS dynamics that drive shared changes across the cardiovascular and respiratory time series, rather than the millisecond-level or second-to-second changes specific to cardiovascular time series. Resampling and filtering time series to focus system dynamics at a particular timescale of interest are well-established approaches [33]. Resampling and time alignment produced nine physiomarker time series that could now be described as a single time series of nine-dimensional feature vectors for each participant, $\{f_k\}_{k=0}^{T_{total}}, f_k \in \mathbb{R}^M$, where $T_{total} + 1$ is the number of total timesteps available for the participant's data, separated by 5-s intervals. $M = 9$ is the number of physiomarkers considered in this work. To produce each y_k, each f_k was then centered relative to $\bar{f}_b = \frac{1}{b+1}\sum_{i=0}^{b} f_i$, the mean vector during the baseline period. $\forall k \;\; y_k \equiv f_k - \bar{f}_b$. As shown in Fig. 2A, the baseline period was 60 s (i.e., $b = 12$).

The inputs shown in Fig. 2B, $\{u_k\}_{k=0}^{T_{total}}, u_k \in \mathbb{R}^V$, represent the presence or absence of the protocol conditions shown in Fig. 2A. In this work, $V = 3$ to represent the three types of protocol conditions. Let $\forall k \;\; u_k \equiv [u_{s,k}, u_{t,k}, u_{n,k}]^T \in \{0, 1\}^3$

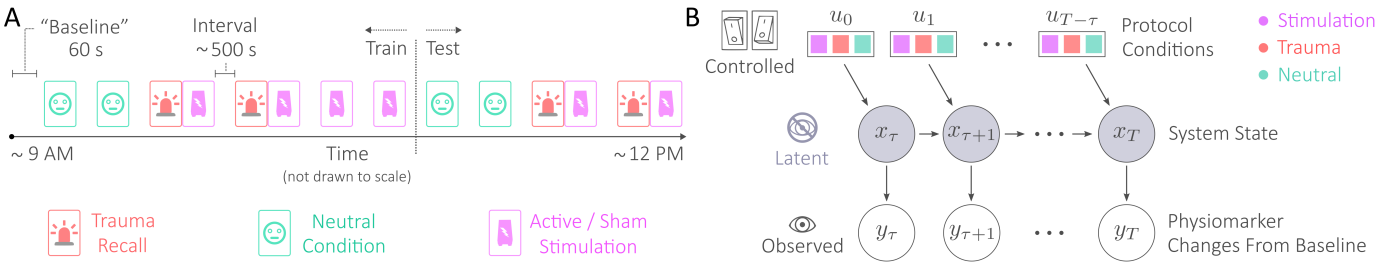

Fig. 2. Data collection and person-specific modeling methods. (A) The double-blind, randomized controlled trial consisted of trauma recall, neutral conditions, and active or sham stimulation. The protocol conditions were ordered as shown. For modeling purposes, each participant's data were split for training and testing at the midpoint between the stimulation and neutral conditions shown. (B) For each participant, a multi-input, multi-output state space model with time delays was trained on data prior to the train-test split and tested on held-out data after the train-test split. The inputs to the model were on-off signals corresponding to the controlled protocol conditions. The outputs of the model were changes in the physiological markers ("physiomarkers") from their baseline averages. These changes were modeled as observable mappings of the system's latent state.

be constructed as one-hot vectors. $u_{s,k}, u_{t,k}, u_{n,k} \in \{0,1\}$ represent the input dimensions at timestep k associated with active or sham stimulation, trauma recall, and neutral conditions, respectively. Let the sets $\mathcal{N}, \mathcal{T}, \mathcal{S}$ include all timesteps k that fall within neutral conditions, trauma recall conditions, or stimulation conditions, respectively. $\forall k \in \mathcal{N}\ u_k = [0,0,1]^T$. $\forall k \in \mathcal{T}\ u_k = [0,1,0]^T$. $\forall k \in \mathcal{S}\ u_k = [1,0,0]^T$. $\forall k \notin \mathcal{N} \cup \mathcal{T} \cup \mathcal{S}$, $u_k = [0,0,0]^T$.

D. Train-Test Split for Each Participant's Data

Each participant's input and output time series, $\{u_k\}_{k=0}^{T_{total}}, \{y_k\}_{k=0}^{T_{total}}$, were split into training data, $\{u_k\}_{k=0}^{T_{train}}, \{y_k\}_{k=0}^{T_{train}}$, and held-out testing data, $\{u_k\}_{k=T_{train}+1}^{T_{total}}, \{y_k\}_{k=T_{train}+1}^{T_{total}}$. As shown in Fig. 2A, the data were split chronologically between the fourth stimulation condition's end time, ζ, and the third neutral condition's start time, γ (i.e., $T_{train} \approx (\zeta+\gamma)/2$; rounded to integer). Splitting the data at T_{train} ensured that both train and test data included trauma, neutral, and stimulation conditions, while allocating more data for training.

The training data were used to fit the models and optimize parameters, as described in II-F and II-G, while the testing data were used to assess predictive capacity of the trained models, as described in II-H. The mean number of training timesteps across participants was 650 (SD 103), with a range of 455-985. The mean number of testing timesteps was 422 (SD 77), with a range of 97-618. The mean proportion of training relative to total timesteps was 61% (SD 5%), with a range of 50%-82%.

E. Person-Specific, MIMO State-Space Models

As illustrated in Fig. 1, the dynamic model graphically shown in Fig. 2B models the effects of exogenous stimuli on changes in physiomarkers, mediated by latent dynamics of ANS state. MIMO linear state space models were fit to the inputs, $\{u_k\}_{k=0}^{T_{train}}$, and outputs, $\{y_k\}_{k=0}^{T_{train}}$, described in II-D. We model the effects of $\{u_k\}_{k=0}^{T_{train}}$ on $\{y_k\}_{k=0}^{T_{train}}$ via a lower dimensional latent state, $\{x_k\}_{k=0}^{T_{train}}$, $x_k \in \mathbb{R}^O$. O is referred to as the model order or dimensionality of the latent state.

$$x_{k+1} = Ax_k + Bu_{k-\tau} + Ke_k \tag{1}$$

$$y_k = Cx_k + e_k \tag{2}$$

The state space model structure shown in (1) and (2) is defined in innovations form with input delays and no feedthrough effects [33]. $A \in \mathbb{R}^{O\times O}$ represents the system matrix that dictates the state's autonomous dynamics, $B \in \mathbb{R}^{O\times V}$ quantifies the input's impact on state dynamics, and $C \in \mathbb{R}^{M\times O}$ represents the output matrix that maps latent state to observable output. $e_k \in \mathbb{R}^M$ represents an innovation term at timestep k that can be viewed as a residual, and $K \in \mathbb{R}^{O\times M}$ represents a Kalman gain matrix used to update next-step predictions of state based on the current timestep's residual. τ encapsulates delays between changes in input and changes in state. O and τ are treated as hyperparameters in this work.

In this work, $V = 3$, and the input can be separated into condition-specific dimensions, $u_k \equiv [u_{s,k}, u_{t,k}, u_{n,k}]^T$. Hence, we can rewrite $u_{k-\tau}$ and separate B into terms quantified for each protocol condition, as in (3) and (4).

$$u_{k-\tau} \equiv [u_{s,k-\tau_s}, u_{t,k-\tau_t}, u_{n,k-\tau_n}]^T \tag{3}$$

$$B \equiv [b_s, b_t, b_n] \tag{4}$$

$b_s, b_t, b_n \in \mathbb{R}^O$, and τ_s, τ_t, τ_n are input delays for stimulation, trauma recall, and neutral conditions, respectively.

Two constraints were placed on A to ensure physiological plausibility and reduce the number of parameters to be learned. First, A was constrained to modal form with no repeated eigenvalues. Modal form entails a block diagonalized matrix, with eigenvalues, $\{\lambda_i\}_{i=1}^O$, placed along the diagonal. The number of parameters that needed to be learned in A was thereby reduced from O^2 to a maximum of $2O - \mathrm{mod}(O,2)$, where $\mathrm{mod}(O,2)$ denotes O modulo 2 (proof in supplement). Second, the eigenvalues of A, $\{\lambda_i\}_{i=1}^O$, were constrained to lie within the unit circle (i.e., $\forall i\ |\lambda_i| < 1$). This enforces asymptotic stability. An asymptotically stable system converges back to the origin if no exogenous inputs are present, which aligns with physiological homeostasis. Recall that the output, y_k, is defined as deviations from a baseline mean (II-C).

F. Model Training

For a given model order and set of input delays, $(O, \tau_s, \tau_t, \tau_n)$, each participant-specific model was trained using a two-stage routine to learn the parameters in A, B, C, K and produce latent states, $\{x_k\}_k$, from the input and output

data. Models were initialized using subspace identification (N4SID) [33]. We then employed prediction error minimization (PEM) to mitigate limitations of N4SID [33]. Our combined N4SID and PEM approach guarantees replicability. For the same data and model structure, the same model will be produced each time N4SID and PEM are employed. This is unlike approaches such as expectation maximization or stochastic gradient descent that are popular but yield differing results depending on initialization or random seed [34].

G. Hyperparameter Optimization

To determine an optimal model order and set of time delays for each participant-specific model, $(O^*, \tau_s^*, \tau_t^*, \tau_n^*)$, we employed a grid search approach. For each participant's training data, the final 30% were set aside as "validation" data (i.e., subset of the training data used to evaluate candidate hyperparameter combinations), $\{u_k\}_{k=T_{fit}+1}^{T_{train}}, \{y_k\}_{k=T_{fit}+1}^{T_{train}}$. We define the subset of the training data used only for training and no hyperparameter tuning as "fitting" data for clarity, $\{u_k\}_{k=0}^{T_{fit}}, \{y_k\}_{k=0}^{T_{fit}}$. $T_{fit} \approx 0.7 \cdot T_{train}$ (rounded to integer). Model order, $O \in \mathbb{N}$, was optimized over the set $\{3, 4, ..., M\}$ to keep $O \leq M = 9$. The stimulation delay term, $\tau_s \in \mathbb{N}$, was optimized over the set $\{1, 2, ..., 5\}$ timesteps (i.e., 5-25 s) based our prior work [13]. The neutral and trauma delay terms, $\tau_n, \tau_t \in \mathbb{N}$, were optimized over the set $\{0, 1, 2\}$ timesteps.

The task used for hyperparameter optimization was the same used for testing: next-step prediction. Next-step prediction entails predicting y_i, having observed all output values until y_{i-1}. All input values, $\{u_i\}_i$, remain known. For a linear state space model, optimal predictions can be produced via Kalman filtering. An advantage of the innovations form shown in (1) and (2) is that the model is directly fit in a manner that facilitates Kalman predictions [33] (details in supplement). The metric used for hyperparameter optimization was the coefficient of determination, R^2. Models were trained on fitting data, and R^2 values were computed for next-step predictions on validation data. For each combination of $(O, \tau_s, \tau_t, \tau_n)$, a scalar R^2 was computed by taking the average of the physiomarker-specific coefficients of determination, $R^2 = \frac{1}{M}\sum_{i=1}^{M} R_i^2$.

Once an optimal $(O^*, \tau_s^*, \tau_t^*, \tau_n^*)$ was determined by maximizing R^2, $(O^*, \tau_s^*, \tau_t^*, \tau_n^*)$ was used to produce the final participant-specific model by training on the entire training set, $\{u_k\}_{k=0}^{T_{train}}$, $\{y_k\}_{k=0}^{T_{train}}$. For robustness, we performed additional tests and found that alternative methods to optimize $(O^*, \tau_s^*, \tau_t^*, \tau_n^*)$, such as minimizing the Akaike Information Criterion with a small sample correction (AICc) on the training data [13], also yielded predictively valid models.

H. Model Testing

Each trained model was tested on participant-specific, held-out data, $\{u_k\}_{k=T_{train}+1}^{T_{total}}, \{y_k\}_{k=T_{train}+1}^{T_{total}}$. Testing was performed via next-step prediction, which is an established approach to dynamic model testing [13], [35]. To determine whether testing performance was satisfactory, the R^2 values for next-step prediction were statistically compared against the corresponding R^2 values for two established baselines: the mean predictor ("mean test") and the naive predictor ("naive test") [35]. The mean predictor predicts that the output will remain the mean of the training data (i.e., $\forall k \hat{y}_k = (1/(T_{train}+1))\sum_{j=0}^{T_{train}} y_j$). The naive predictor predicts that the output will remain the same as the most recently observed output (i.e., for next-step prediction, $\forall k\ \hat{y}_k = y_{k-1}$). Naive prediction can be viewed as assuming identity dynamics, $A = I$, and the mean predictor can be viewed as assuming no state dynamics, $A = 0$. The mean and naive predictors thus represent the two trivial ends of the spectrum from $A = 0$ to $A = I$ for autonomous (i.e., $B = 0$), stable, discrete-time linear dynamical systems. Outperforming these two baselines thus demonstrates that the dynamics learned are nontrivial.

For statistical testing, the R^2 values across all $N = 50$ participant-specific models were compared against the corresponding R^2 values for the naive and mean predictors via paired tests. The Shapiro-Wilk test rejected normality, so two-tailed Wilcoxon signed rank tests were used. Statistical significance was set to α = .05 for all tests in this work.

I. Forecasting Experiments

To further evaluate predictive capacity, we investigated the models' forecasting capabilities. Multi-step forecasting is a generalization of the next-step prediction task, where p-step forecasting requires that the output y_k is predicted, having only observed data until y_{k-p}. The state space model equations of (1) and (2) can be recursively solved forward to obtain equation (5) for p-step forecasting, as proven in the supplement.

$$\hat{y}_k = CA^p x_{k-p} + \sum_{i=1}^{p} CA^{p-i} B u_{k-(p-i)-1-\tau} \tag{5}$$

For the first $p+\tau$ timesteps of held-out data, the last timesteps of training data are used to compute forecasts via (5). Inputs and states are all set to 0 for $k < 0$ (i.e., $\forall_{k<0}\ u_k = x_k = 0$).

We again compared model performance against the two baseline forecasting methods, the mean and naive predictors. For multi-step forecasting, nothing changes for the mean predictor. For the naive predictor, y_k is predicted to be the same as the last observed value, y_{k-p} (i.e., $\forall k\ \hat{y}_k = y_{k-p}$). Paired statistical tests were again used to compare the models' forecasting performance against the baseline methods' performance (i.e., R^2 values) for each p investigated. Shapiro-Wilk tests were used to test for normality. If normality was rejected, then two-tailed Wilcoxon signed rank tests were used. Otherwise, two-tailed, one-sample t-tests were used.

J. Simulating Condition-Specific Step Responses

After validating our modeling approach, we investigated the models' step responses to each of the protocol conditions. Similar to prior work [13], this produced interpretable visualizations of what the models captured from the effects of each input. As shown in Fig. 3A, three sets of single-input step functions were constructed and investigated as inputs to each participant-specific model. The stimulation-specific step response consisted of constructing the stimulation input as follows: $u_{s,0} = u_{s,1} = 0$, $\forall_{1<k\leq T_{sim}}\ u_{s,k} = 1$. The trauma recall and neutral inputs were then set as follows:

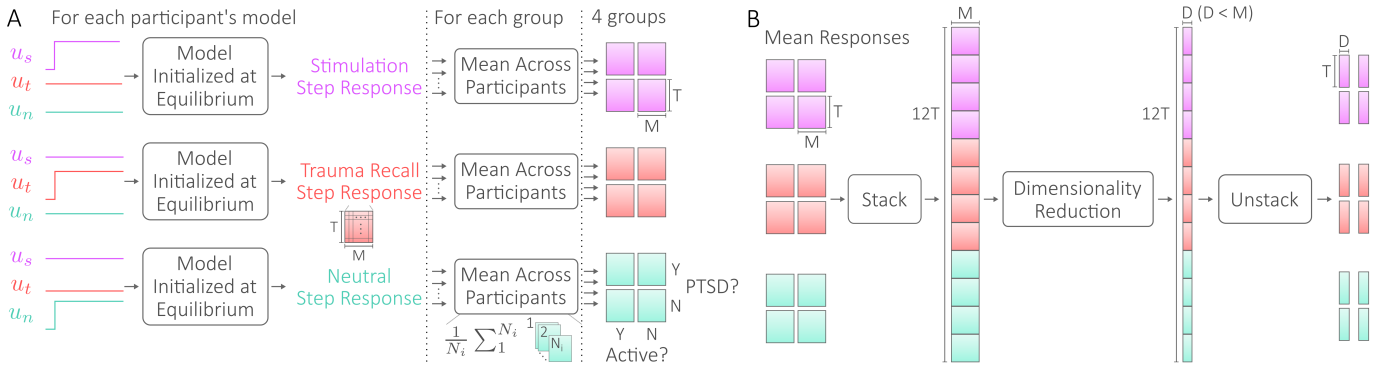

Fig. 3. Step response and dimensionality reduction methods. (A) Each model's condition-specific step responses for stimulation, trauma, and neutral were simulated for T timesteps. The models were initialized at the origin and simulated without noise to isolate the condition-specific effects learned by the models. The simulated responses were then averaged across participants in each of the four groups (active vs. sham x PTSD vs. non-PTSD) for all M physiomarkers. (B) The mean responses were analyzed in lower dimensional principal subspaces of dimension D. The mean responses were vertically stacked to align the M physiomarkers across groups, followed by dimensionality reduction and subsequent analyses.

$\forall_{k\in\{0,1,...,T_{sim}\}}\ u_{t,k} = u_{n,k} = 0$. In a similar way, the trauma-specific step input and neutral-specific step input were constructed by setting all input values to 0 except for the corresponding input signal for timesteps $1 < k \leq T_{sim}$. For the condition-specific step responses, we set $T_{sim} = 13$ to investigate the first minute of the responses, matching the approximate length of the trauma and neutral conditions.

Step responses were simulated by applying (5), setting $p = k$ for $k \in \{0, 1, ..., T_{sim}\}$. We set $x_0 = 0$ for all simulations to initialize the models at the origin (i.e., the equilibrium). Step response simulations produced output arrays of dimension $T_{sim} \times M$, as illustrated in Fig. 3A. As detailed in the supplement, the use of (5) implies that simulation is conducted without noise. For these simulations, we are interested in visualizing what the models learned about the condition-specific responses, in contrast to simulation settings where the inclusion of stochasticity is important (e.g., evaluating the robustness of a control law) [33]. Replicability is also guaranteed by removing stochasticity from these simulations.

1) Physiomarker-Specific Responses

For the purposes of comparison, the step responses were grouped according to active or sham stimulation and non-PTSD or PTSD diagnosis. The responses were averaged across participants within each group, as shown in Fig. 3A. For each protocol condition indexed by $j \in \{1, 2, 3\}$, a three dimensional array of step responses was formed for each of the four groups, $S_{i,j} \in \mathbb{R}^{T_{sim}\times M\times N_i}$, where $i = \{1, 2, 3, 4\}$. N_i is the number of participants in group i. A mean response, $\bar{S}_{i,j}$ is then computed by taking the mean along the third axis of $S_{i,j}$. Performing this grouping and averaging process for every $j \in \{1, 2, 3\}$ and $i \in \{1, 2, 3, 4\}$ produced a total of 12 group-averaged, condition-specific step responses. Standard errors of the mean (SEMs) were also calculated across participants in each group to understand the variability of the responses.

To compare the simulated step responses with real data, we employed "transient focusing" methods to the original physiomarker time series [36]. Transient focusing is detailed in the supplement and ensures that analyses of physiomarker data are focused on the timescale of interest. For each participant, the segmented and transient focused responses of the same protocol condition were first averaged to produce three condition-specific response arrays [36], analogous to the three step response arrays shown for each participant in Fig. 3A. The participant-specific, condition-specific responses were then grouped and averaged in the same way as the step responses shown in Fig. 3A, producing $\bar{R}_{i,j}$ (and corresponding SEMs). Comparisons between real and simulated responses were made for each physiomarker (i.e., each column of $\bar{S}_{i,j}$ and $\bar{R}_{i,j}$). Note that these comparisons are not made for the purposes of model validation but instead to gain intuition as to which aspects of the transient-focused, real responses were captured by the models in their simulated step responses.

2) Lower Dimensional Responses

To investigate how the physiomarker responses covaried, we applied principal component analysis (PCA) to $\bar{S}_{i,j}$, $\bar{R}_{i,j}$ to capture covariation in a linear subspace. The process used for dimensionality reduction is illustrated in Fig. 3B. Dimensionality reduction was performed separately for the simulated and real responses. Let $\bar{Q}_{i,j}$ represent either $\bar{S}_{i,j}$ or $\bar{R}_{i,j}$. The $\bar{Q}_{i,j}$ arrays were first stacked vertically for all i, j to form $Y \in \mathbb{R}^{12T_{sim}\times M}$. Each column of Y was scaled to unit variance but not mean-centered to preserve the origin, (0, 0). We then applied PCA to Y to learn a principal component (PC) coefficients matrix, $P \in \mathbb{R}^{M\times M}$. To compute the lower dimensional representation (i.e., PC scores) of Y, we computed $Z = YP$, where $Z \in \mathbb{R}^{12T\times M}$ represents the PC scores. The first D columns of Z were then isolated to analyze the reduced D-dimensional responses. The D columns of Z are then unstacked and regrouped according to the 12 group-condition combinations. Throughout, we keep chronological ordering intact for temporal analyses, as in prior work [37].

To understand the relationships between physiomarkers as they covary, we investigated the coefficients (i.e., the "recipe") of the principal axis of variance (i.e., PC1) for the real and simulated responses. We also quantified how fast the responses to each protocol condition manifested along PC1. Velocities were computed by dividing the differences between successive response values by the difference in time (5 s in this work).

Additional post hoc analyses were performed after observing differences between the trauma recall and stimulation

response directions along PC1, as well as agreement between the real and simulated response directions. Each participant's condition-specific simulated responses were projected along the PC1 vector computed for the mean simulated responses. The same was done for the real responses with the respective PC1. To statistically test whether trauma recall responses existed primarily on one half of the PC1 axis while stimulation responses existed on the opposite half, we computed the proportion of PC1 projections with a positive mean value for the trauma recall and stimulation responses and tested whether these proportions were different from 50% using two-tailed binomial tests. We next evaluated the agreement between the real and simulated responses in their ordering of the trauma recall, neutral, and stimulation conditions along the PC1 axis. The mean PC1 projections for the three conditions were rank ordered for each participant's simulated data and real data. A two-tailed repeated measures correlation was performed on these ranks to test whether the ordering of the conditions along PC1 agreed between real and simulated data [38].

K. Simulating Just-in-Time Neuromodulation

We conducted a final set of simulations to investigate the potential for tcVNS as a just-in-time stress-reducing intervention in response to trauma recall. Focusing on the PTSD group for clinical relevance, three sets of simulation inputs were constructed for each model in the active and sham PTSD groups. As before (II-J), a neutral step function was constructed for reference, and a trauma recall step function was constructed for comparison. The third set of input signals was formed by setting the neutral input to 0, (i.e., $\forall_{0\leq k\leq T_{JIT}}\ u_{n,k} = 0$, the trauma recall input to the trauma recall step function, $u_{t,0} = u_{t,1} = 0$, $\forall_{1<k\leq T_{JIT}}\ u_{t,k} = 1$, and the stimulation input to a delayed step function. Specifically, we set $\forall_{0\leq k<\kappa_{start}}\ u_{s,k} = 0$ and $\forall_{\kappa_{start}\leq k\leq T_{JIT}}\ u_{s,k} = 1$ – effectively turning stimulation on after κ_{start} timesteps. In this work, $\kappa_{start} = 6$. We set the length of time for these just-in-time simulations to 90 s (i.e., $T_{JIT} = 18$) to investigate the responses for 60 s following the start of stimulation.

This final set of simulations produced three response arrays for each participant in the active or sham PTSD groups: neutral alone, trauma recall alone, and trauma recall combined with stimulation (i.e., "just-in-time stimulation"). The three response arrays were then grouped according to active and sham stimulation groups (i.e., $i \in \{1, 2\}$ in this case). All aggregation, physiomarker-by-physiomarker analyses, and lower dimensional analyses described in II-J were repeated. The only difference is that no real responses were available for comparison since tcVNS (or sham) delivered during trauma recall has not been investigated experimentally to date.

An additional post hoc analysis was performed after observing differences between the active and sham groups' responses. Similar to the post hoc analyses described in II-J, each participant's simulated responses were projected along the PC1 computed for the mean simulated responses. To investigate whether just-in-time stimulation was simulated to have a greater attenuating effect on trauma recall for the active group compared to sham, we computed the difference

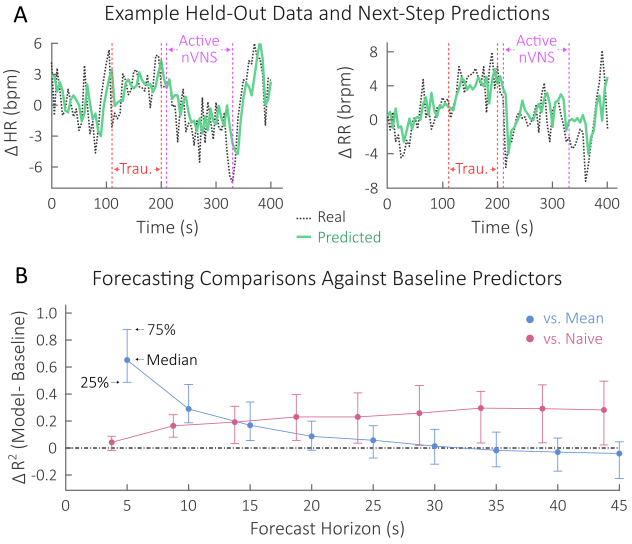

Fig. 4. Model prediction performance. (A) Example held-out heart rate (HR) and respiration rate (RR) data, shown as changes from baseline (**Δ**), for a patient with PTSD in the active group. The data shown include periods of trauma recall (traum.) and non-invasive vagus nerve stimulation (nVNS). The next-step predictions of this patient's model are overlaid for comparison. (B) Comparisons of the models' forecasting performance against two established baselines. The differences in coefficient of determination are plotted against forecast horizon (i.e., the time into the future that data are forecasted). The models significantly ($p < 0.05$) outperform both baseline methods for horizons up to 20 s.

in velocities between the just-in-time stimulation response and the response to trauma recall alone for each participant. To account for delayed stimulation effects, we considered velocities 15 s after stimulation started and computed a mean velocity difference for each participant. These mean velocity differences were compared between active and sham groups using a two-tailed, two-sample t-test (normality and equality of variance were tested using Shapiro-Wilk and Levene's tests).

III. Results

A. Model Testing and Forecasting Performance

On held-out test data, our modeling approach significantly outperformed the baseline naive (P = .02; test statistic z = 2.28; effect size r = .32) and mean prediction methods (P < .001; z = 5.84; r = .83). Differences between our models' R^2 values and the two baseline predictors' R^2 values are shown in Fig. 4B. Next-step prediction corresponds to a forecast horizon of 5 s in this work. Example next-step predictions for HR and RR are shown in Fig. 4A, overlaid on ground truth data. Δ in Fig. 4A denotes changes in the physiomarkers from a baseline mean. The data shown in Fig. 4A are from a participant with PTSD who received active tcVNS.

Multi-step forecasting results on held-out data are shown in Fig. 4B. The models significantly outperformed the naive predictor for all forecasting horizons investigated (P < .05). The models significantly outperformed the mean predictor up to 20 s into the future (P < .05). Although the median difference between the models' and mean predictor's performances were nonzero for horizons of 25 to 40 s, these differences were not statistically significant. The mean predictor outperformed the models for forecast horizons of 45 s and greater (P < .05).

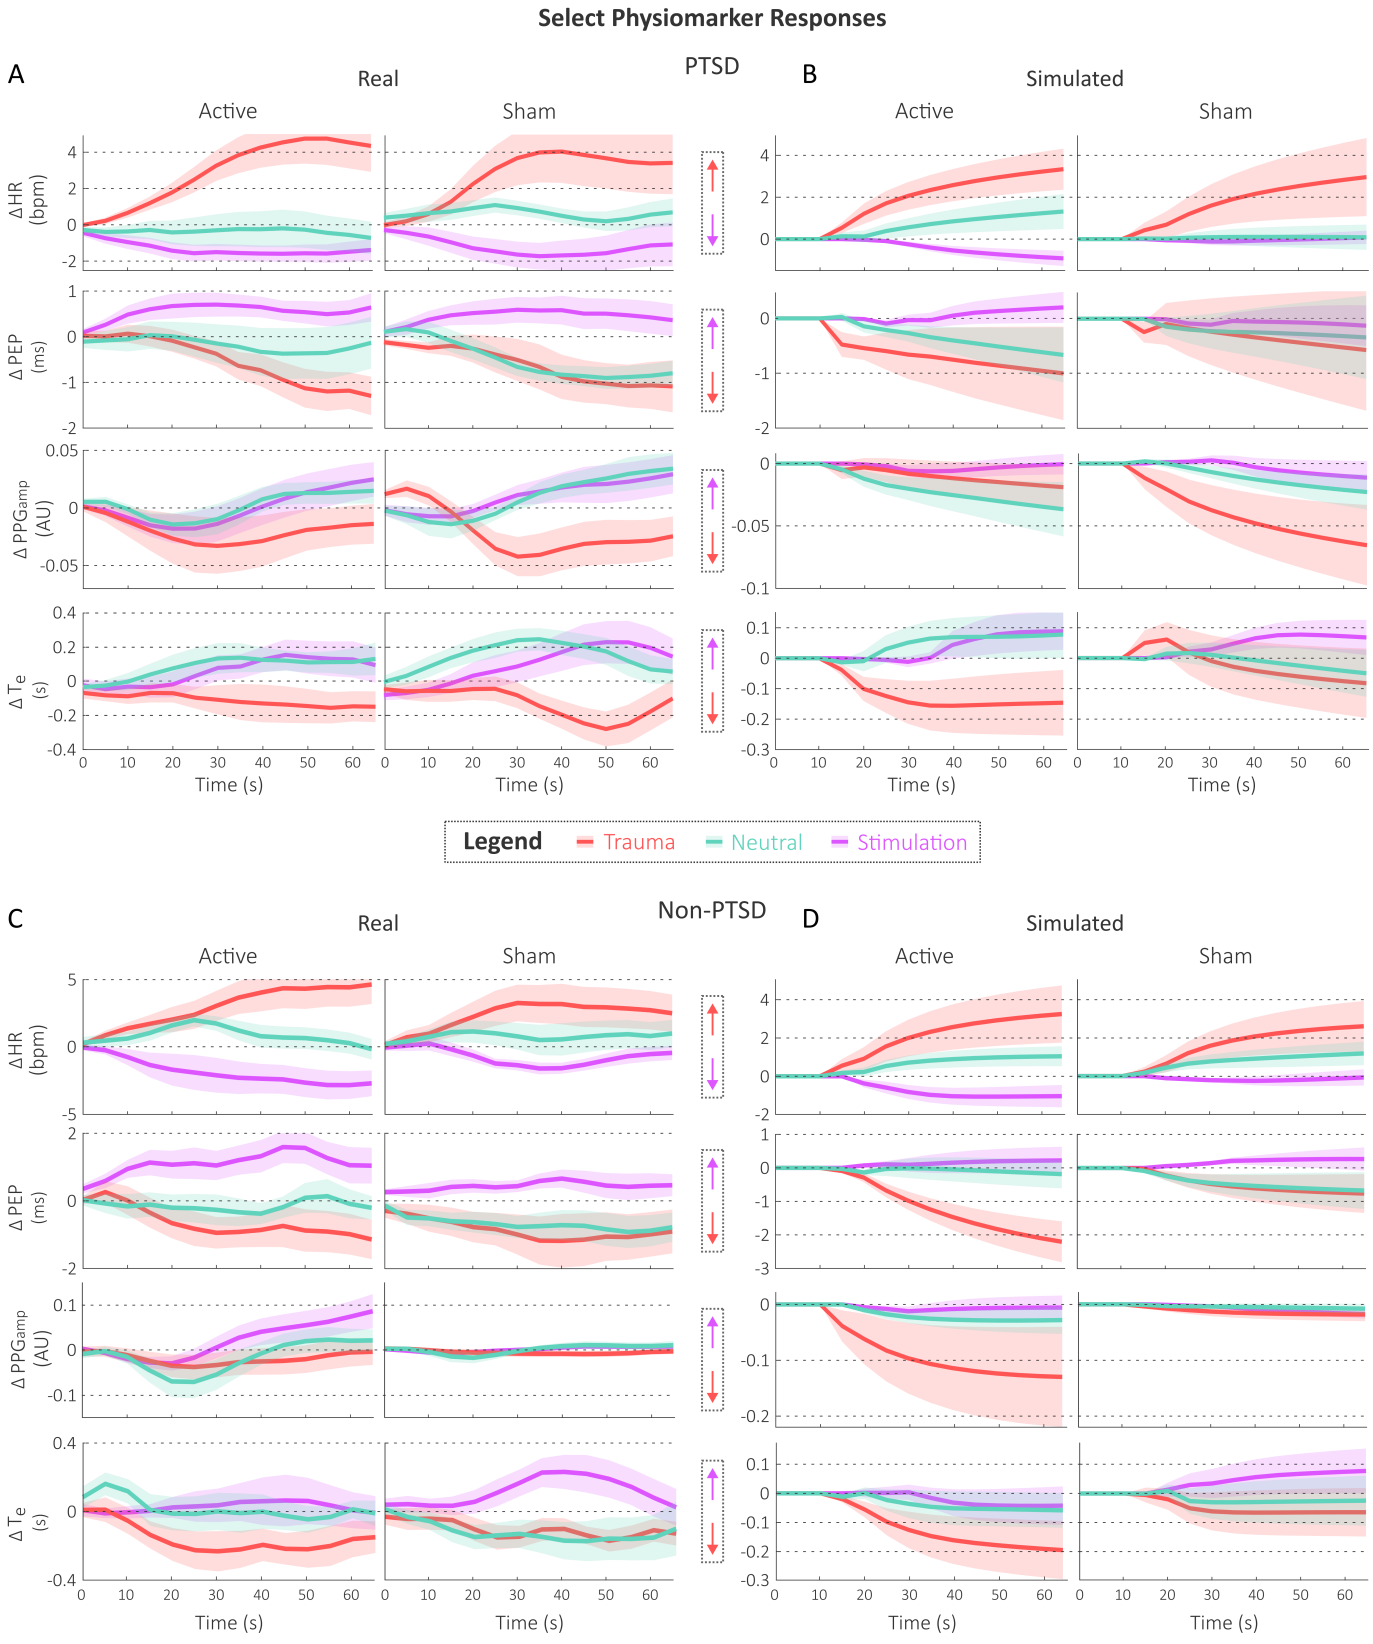

Fig. 5. Select mean physiological marker ("physiomarker") responses to neutral, trauma, and active or sham stimulation. Δ denotes changes from baseline. Shaded regions represent ± standard error of the mean (SEM). The remaining five sets of physiomarker responses are included in the supplement. The sets of arrows included in the middle of the figure highlight the directionality differences captured between trauma recall and stimulation. Physiomarker acronyms are detailed in the main text. (A) Real data for the group of participants with posttraumatic stress disorder (PTSD). (B) Simulated step responses of the models for participants with PTSD. (C) Real data for participants without posttraumatic stress disorder (PTSD). (D) Simulated step responses of the models for participants without PTSD.

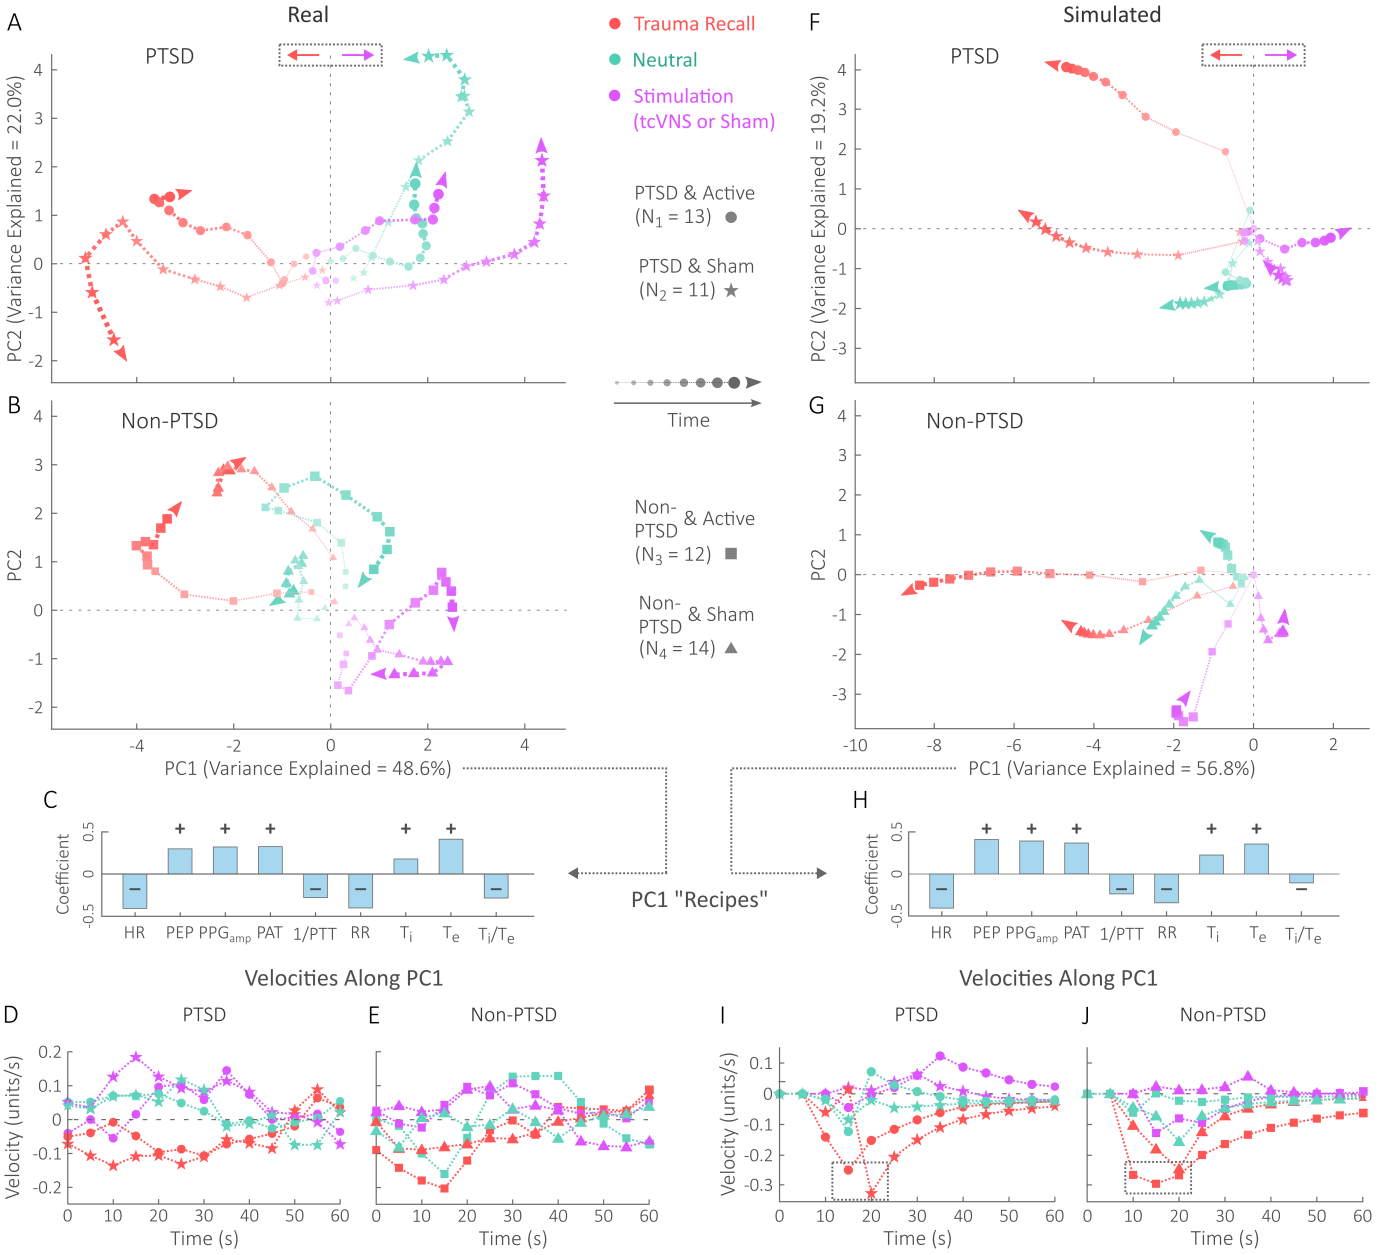

Fig. 6. Physiological responses to the neutral, trauma, and stimulation conditions, analyzed in lower-dimensional principal subspaces. Real data and the corresponding first learned principal component (PC1) are shown on the left (A, B, C, D, and E), while simulated data and the corresponding PC1 are shown on the right (F, G, H, I, and J). The shared color legend distinguishes protocol conditions, while the shared marker legend distinguishes participant groups. Opposing arrows in dotted boxes highlight the opposite directions of trauma and stimulation responses. (A and B) Mean responses to each condition mapped into a 2-D principal subspace for the PTSD and non-PTSD groups, respectively. Trajectories start near the origin and terminate with an arrow. The PC1 axis is shared. (C) The PC1 "recipe," or vector of coefficients that form PC1, for the real data. + and − symbols highlight sign differences between coefficients. (D and E) Velocities along PC1 for responses in subplots A and B, respectively. The velocity axis is shared. (F and G) Group-averaged step responses to each condition, mapped into a 2-D subspace for the PTSD and non-PTSD groups, respectively. Trajectories start at the origin. The PC1 axis is shared. (H) The PC1 recipe for the step responses. (I and J) Velocities along PC1 for responses in subplots F and G, respectively. The velocity axis is shared. Dotted boxes highlight the speed of trauma responses.

B. Condition-Specific Responses

1) Physiomarker-Specific Comparisons

Fig. 5 includes all four participant groups' condition-specific responses for select cardiac (HR and PEP), vascular (PPG_{amp}), and respiratory (T_e) markers. The remaining five physiomarkers' responses are included in the supplement. The overall positive or negative directions of the simulated and real responses agree for both the PTSD (comparing Fig. 5A and 5B) and non-PTSD (comparing Fig. 5C and 5D) groups. Generally, responses to trauma recall and stimulation oppose another, with neutral responses remaining in between. Responses to stimulation are generally greater in magnitude for the active group compared to the sham group.

2) Lower Dimensional Comparisons

Fig. 6 shows the condition-specific real and simulated responses analyzed in lower dimensional principal subspaces. The first two PCs explain nearly 75% or more of the variance. Although the real data exhibit more variability as expected, the

real and simulated responses agree with respect to their overall directions along PC1. In Fig. 6A, 6B, 6F, and 6G, we see that the real and simulated trauma recall responses initially have decreasing values along PC1 as time passes, while stimulation responses generally have increasing values along PC1. This observation is supported by our binomial test results. Of the 50 participants' real responses, 7 trauma recall responses had mean PC1 values greater than 0 (14%, $P < .001$, 95% confidence interval [CI]: [5.8%, 27%]); and 39 stimulation responses had mean PC1 values greater than 0 (78%, $P < .001$, 95% CI: [64%, 88%]). Of the 50 models' simulated responses, 12 trauma recall responses had mean PC1 values greater than 0 (24%, $P < .001$, 95% CI: [13%, 38%]); and 34 stimulation responses had mean PC1 values greater than 0 (68%, P = .02, 95% CI: [53%, 80%]). Comparing Fig. 6A and 6B to Fig. 6F and 6G, the real and simulated responses generally agree in the ordering of the trauma recall, neutral, and stimulation responses along PC1. In particular, stimulation responses have the greatest PC1 values, and trauma recall responses have the lowest PC1 values. Neutral responses remain in between. This agreement is supported by our repeated measures correlation analysis. The rank orders of the three conditions for the real and simulated responses were statistically significantly correlated (r = .39, 95% CI: [.21, .54], $P < .001$). Video versions of Fig. 6A and 6B overlaid and Fig. 6F and Fig. 6G overlaid are included as supplementary material.

Comparing Fig. 6D and Fig. 6E to Fig. 6I and Fig. 6J, the velocity directions agree for the first 40 s. The velocities for the real and simulated trauma recall responses are negative, while the velocities for the stimulation responses are generally positive, with neutral velocities remaining closer to 0. Transient focusing does not re-center the real responses' initial velocities closer to 0 [36], obfuscating comparisons. Note that the velocities of the real responses indicate equilibration after 40 s and confounding carryover effects from the recovery after trauma recall to the stimulation response, as can be seen for the PTSD sham group. We see from the simulated responses that responses to trauma recall are greater in speed compared to stimulation, while neutral responses are the slowest. We also see that the effects of active stimulation are stronger in magnitude than sham.

The orthant in which PC1 lies (i.e., direction of PC1) is the same for the real and simulated data, as shown in Fig. 6C and Fig. 6H, respectively. The PC1 obtained from the real data explains 49% of the variance for the responses in Fig. 6A and Fig. 6B. The PC1 obtained from the simulated data explains 57% of the variance for the responses in Fig. 6F and 6G.

C. Simulations of Just-in-Time Neuromodulation

Fig. 7 shows the simulation results for just-in-time stimulation in response to trauma recall for participants with PTSD. In Fig. 7A, the effects of just-in-time stimulation are most apparent in reducing HR and increasing T_e and PEP compared to trauma recall in the active group. Just-in-time stimulation's effects are less apparent for PPG_{amp}. The delayed effects of tcVNS are clear in HR, where tcVNS is turned on at 30 s, but reduction in HR is not evident until approximately 15 s later. Changes are not as evident for the other physiomarkers and even less so for the sham group. The remaining five physiomarkers' responses are included in the supplement.

The simulated effects of just-in-time tcVNS are more evident along PC1. In Fig. 7B, just-in-time tcVNS attenuates the response to trauma recall within 30 s of stimulation. No attenuation along PC1 is observed for the sham group. The PC1 recipe is shown in Fig. 7C. The direction of PC1 aligns almost exactly with the direction of PC1 shown for the simulated and real responses of Fig. 6. PC1 explains 71% of the variance. The velocities along PC1 are shown in Fig. 7D. The velocities indicate that just-in-time tcVNS has an attenuating effect on the response to trauma recall within 15 s. In fact, just-in-time stimulation reverses the response to trauma recall for the active group, where the velocity along PC1 switches from negative to positive for $\sim$20 s. Just-in-time sham stimulation shows no reversal in velocity along PC1. This observation is supported by our post hoc analysis. The mean differences between the velocities for the just-in-time stimulation responses and the responses to trauma recall were significantly larger (P = .04, t = 2.2, Cohen's d = 0.91) for the active group (mean $\pm$ standard deviation: 0.01 $\pm$ 0.01) compared to the sham group (0.00 $\pm$ 0.01). A video version of Fig. 7B is included as supplementary material.

IV. Discussion

In this paper, we present a novel personalized dynamic modeling approach to quantify the time course of multiple physiomarkers of the ANS in response to acute stressors, such as traumatic memories, and potential stress-reducing interventions, such as tcVNS. Using data from a double-blind, randomized, sham-controlled trial of tcVNS involving 50 individuals with histories of psychological trauma (24 with PTSD) [26], we trained and validated state space models to capture latent ANS dynamics from observed changes in cardiovascular and respiratory physiomarkers. Our findings and simulations provide the first quantitative evidence to suggest that although the stress responses to trauma recall are stronger than the stress-reducing effects of tcVNS, just-in-time tcVNS may potentially attenuate – and briefly reverse – the response to trauma recall. Importantly, our results show that multiple ANS markers covary in response to stressors and interventions, and this covariation is largely explainable by a principal axis of variance. The principal axis is learned without supervision, yet it mirrors changes that are expected psychophysiologically. Our work advances the state of the art in biomedical engineering by developing a replicable, person-specific dynamic modeling approach for multiple physiomarkers of the ANS in response to stressors and interventions. Our findings further psychophysiology and neuromodulation science by quantifying the opposing effects of tcVNS and traumatic memories on ANS dynamics. Overall, this work highlights the potential for closed-loop systems that employ tcVNS to mitigate acute stress for individuals with PTSD. Future work is necessary to replicate these findings in daily life settings and validate our modeling approach in practice.

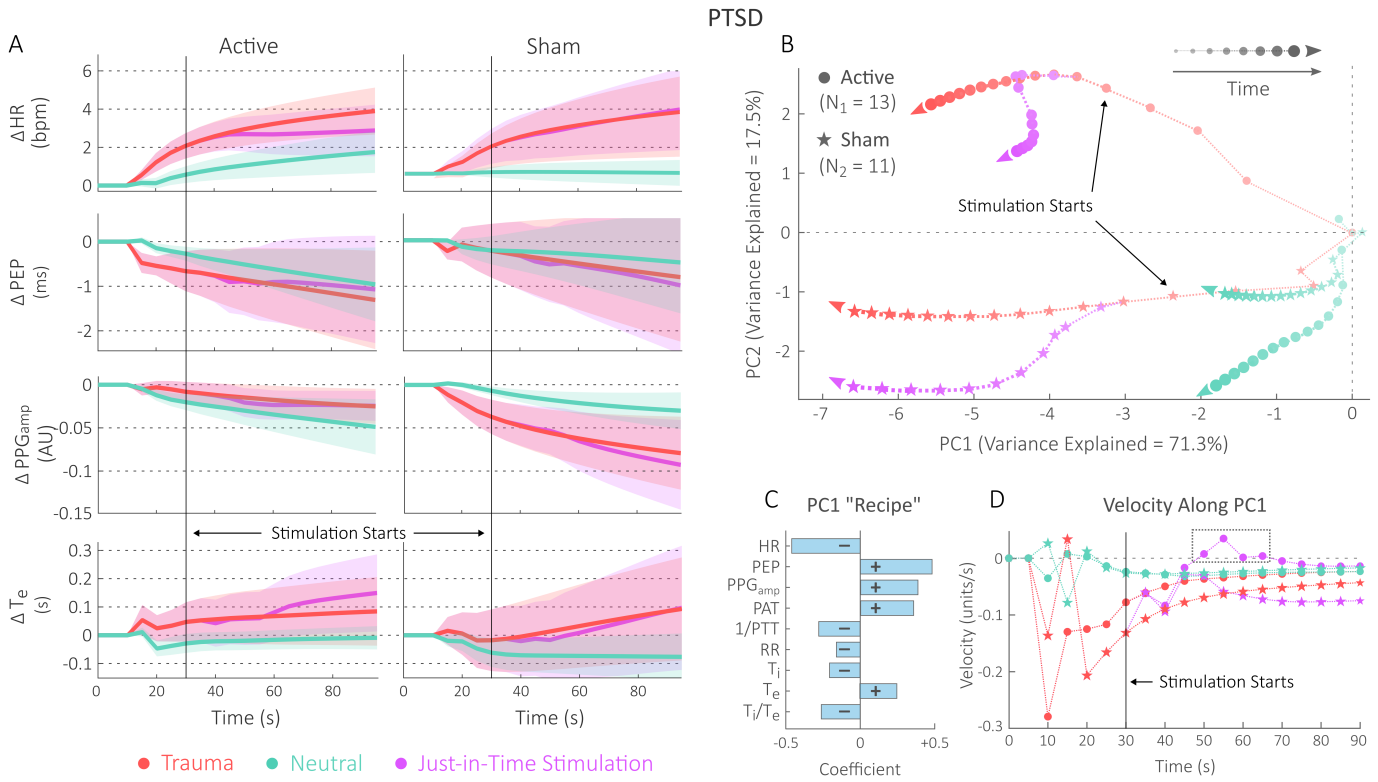

Fig. 7. Simulated responses to just-in-time active or sham stimulation, with responses to trauma recall and the neutral condition included for context. Simulated responses are shown for the models learned on data from patients with posttraumatic stress disorder (PTSD), for both the active and sham stimulation groups. The color legend is shared across subplots, and the marker legend is shared between subplots B and D. (A) Select mean physiomarker responses. Shaded regions represent ± standard error of the mean (SEM). The black vertical lines at 30 s indicate the onset of active or sham stimulation. Physiomarker acronyms are detailed in the main text. The remaining five physiomarker responses are included in the supplementary material. (B) Simulated step responses, averaged within each group, mapped into a 2-D principal subspace. Trajectories start at the origin and terminate with an arrow. The black arrows indicate the point at which stimulation started for both the active and sham groups. (C) The recipe for the principal axis of variation (PC1) learned from the simulated responses in subplot B. The + and − symbols highlight sign differences. (D) Velocities along PC1 for the simulated responses shown in subplot B, plotted against time. The dotted box highlights the reversal in directionality of the active group's just-in-time tcVNS response.

A. TcVNS Effects Counteract Stress Due to Trauma Recall, Though Trauma Recall's Effects are Stronger

This work provides the first quantitative evidence that active tcVNS can attenuate the stress-inducing effects of trauma recall for individuals with PTSD. The literature to date has not quantitatively compared the timecourse of responses to tcVNS and trauma recall in the context of acute stress [11]–[13], [39], [40]. We compare the timecourses in this work and find that both the real and simulated responses to trauma recall and tcVNS show opposite effects to one another. The responses to trauma recall involve increases in HR and reductions in PEP, $\mathrm{PPG_{amp}}$, and $\mathrm{T_e}$, among other physiomarker-specific changes included in the supplement. Increases in HR and decreases in $\mathrm{T_e}$ are indicative of acute stress via decreased parasympathetic activity and/or increased sympathetic activity [2], [30]. Reductions in PEP and $\mathrm{PPG_{amp}}$ indicate increased sympathetic activity [11]. On the other hand, the responses to stimulation capture decreases in HR, increases in PEP, increases in $\mathrm{PPG_{amp}}$ and increases in $\mathrm{T_e}$. In agreement with recent work indicating less tcVNS effects for healthy populations [41], [42], the stress-reducing effects of tcVNS are most evident for participants with PTSD. Along the principal axis of variance, the simulated responses to active tcVNS are opposite to trauma recall responses for the PTSD group. However, the simulated responses to trauma recall are stronger in magnitude and manifest much faster than the responses to tcVNS. By simulating the delivery of tcVNS seconds into trauma recall, we show that although the overpowering stress response to trauma recall is not cancelled by tcVNS, tcVNS has the potential to briefly reverse and attenuate the overall stress response to trauma recall. Sham stimulation did not attenuate the response to trauma recall, indicating that the stress-reducing effects of tcVNS are not placebo effects. These results suggest that tcVNS may be worth investigating as an approach to non-pharmacologically reduce sympathetic or enhance parasympathetic activity to help manage stress. Feasibility studies are necessary to evaluate self-administered tcVNS during episodes of acute stress during daily life.

Our simulations of just-in-time tcVNS motivate continued engineering of closed-loop, non-invasive neuromodulation systems for stress mitigation. Current mobile health interventions for stress management rely primarily on push notification reminders to self-regulate [43]. Although some interventions are founded on evidence-based therapies such as techniques from dialectical behavior therapy for distress tolerance (e.g., TIPP) [44], closed-loop behavioral interventions still require an individual to be receptive to interactions with a mobile

device [45]. The time to interact with a mobile device before an individual's self-awareness or receptivity diminishes due to distress may be insufficient in some cases. We show herein that the stress response to trauma recall peaks quickly (within ∼30 s), and it is well-established that cognitive emotion regulation is significantly less effective once an individual is sufficiently stressed [46]. Closed-loop, non-invasive neuromodulation interventions such as tcVNS (or tVNS delivered auricularly, which has been investigated in daily life settings and closed-loop systems for motor rehabilitation [47], [48]), could be an approach to unconsciously reduce stress when an individual is less receptive to reminders or less conscious of the need to self-regulate. The potential benefits and drawbacks of neuromodulation interventions need to be investigated and compared against behavioral interventions in mobile health settings, paying attention to how biobehavioral context moderates the effects of interventions. In addition, some clinical groups may benefit more from neuromodulation than others, where factors influencing efficacy could include comorbidities, the predominance of physiological symptoms during trauma recall, baseline autonomic function, proficiency in self-regulation, and access, receptivity or resistance to evidence-based treatment [49]. Ultimately, we hypothesize that closed-loop systems that combine behavioral and biological interventions will be optimal, similar to how the combination of pharmacological and behavioral therapies is superior to either alone [50]. Although commercial products exist for self-administered or open-loop tVNS, continued engineering is necessary to actualize closed-loop systems for stress mitigation. For example, our simulations assumed that the acute stress response to trauma recall can be detected within 30 s (i.e., once the stress responses surpass the neutral responses shown in Fig. 6D and 7B). However, detecting acute stress episodes within a minute is nontrivial during daily life [51].

B. Dynamic Models Provide Value Distinct from Static Models and the Physiomarker Responses Themselves

Our dynamic modeling approach enables closed-loop control of stress-reducing interventions in ways that static models cannot. Static models quantify intervention effects at a single timepoint, while dynamic models quantify the timecourse of intervention effects. Interventions should be delivered in a way that is not only favorable for a single timepoint, but rather for the future as well. Optimal control methods often leverage dynamic models to optimize intervention delivery over a sequence of time into the future [52]. This work's forecasting experiments align exactly with optimal control methods such as model predictive control (MPC) [53]. MPC leverages an approximate dynamic model, often a linear state space model, to forecast the outputs of a system in response to potential changes in the inputs. In this work, we show that our approach produces models that outperform baseline forecasting methods up to 20 s into the future. The physiomarkers used as output can all be computed within 5-10 s (depending on breathing rate [31]) using current wearables and real-time signal processing methods [54], enabling measurement updates of state within 10 s. Thus, an example MPC approach could use a prediction horizon of 20 s into the future and a control horizon of 10 s.

The state space models described in this work offer scientific insights that cannot be obtained from the physiomarker time series directly. Our simulation analyses provide two examples. First, the state space models allowed us to precisely control the initial conditions for the condition-specific step responses. We thereby initialized all step response velocities to zero. Forcing all initial velocities to zero helped us quantitatively compare the overpowering responses to trauma recall to the responses to stimulation and neutral conditions. Although transient focusing allowed us to approximately re-center the real responses to zero [36], re-centering velocities would amount to a modified form of linear detrending. De-trending would undermine our ability to compare the directions (i.e., trends) of the condition-specific responses. The second example involves our simulations of just-in-time stimulation. To the best of our knowledge, no trial to date has administered tcVNS or sham stimulation during trauma recall. The state space models allowed us to simulate the delivery of stimulation in response to trauma recall. Although simulations cannot replicate the real world, simulations can inform planning and decision making for future investigations, as is done using digital twins [55].

C. Unsupervised Methods Learn Relationships That Precisely Align With Psychophysiological Literature

An unexpected finding from our analyses is the exact agreement between psychophysiological literature and the relationships between variables quantified via unsupervised machine learning methods. This agreement is supported by two observations. First, the direction of PC1 learned from both real and simulated responses aligns exactly with expected physiomarker changes during reductions in acute stress. Specifically, HR, 1/PTT, RR, and $\mathrm{T_i/T_e}$ decrease as stress diminishes, while PEP, $\mathrm{PPG_{amp}}$, PAT, and $\mathrm{T_e}$ increase [2], [11], [12], [28]. Remarkably, the directions of these changes match the direction of PC1, even though no inductive biases or supervision were used to guide the learning of PC1. Second, the state-space models learned differences between the protocol conditions' effects that align with psychophysiological literature to date. Responses to stimulation were in the direction opposite to responses trauma recall, while responses to neutral conditions remained in between [11], [12]. Similar differences are also evident in the CB parameters learned by the state space models included in the supplement. Notably, A, B, C, and K were optimized solely to predict future physiomarker values based on prior values and protocol conditions, without any explicit information about expected response directions. While it is surprising for data-driven, unsupervised methods to replicate domain knowledge, this outcome underscores the value of selecting physiologically meaningful markers, which provided an informative foundation for data-driven discovery.

D. Limitations and Future Work

This research has limitations. The dynamic models and dimensionality reduction methods used assume linear relationships. Although resultant strengths of our approach include interpretability (e.g., analyses of PC coefficients), nonlinear

modeling to deepen our understanding of dynamics and improve predictive capacity are avenues of future work. The exogenous influence of trauma recall, neutral conditions, and tcVNS or sham stimulation may not match the pulse-like inputs used for modeling. Future studies can investigate the reshaping of input signals according to domain knowledge. Confounding factors that consistently affected the response to a specific protocol condition will not be accounted for by our modeling approach. For example, if a participant was reminded of an occupational stressor during every trauma recall, then the combined effects of trauma recall and occupational stress would be misconstrued as the effects of trauma recall alone. Simulation results should be interpreted accordingly. Outside laboratory settings, stressor presence is often unknown. While our modeling approach incorporates known trauma recall inputs to focus on scientific insights, responses to unknown instances of trauma recall can still be tracked with our approach (by leaving $u_t = 0$). Future work is necessary to evaluate this in practice. The latent states of our state space models do not hold any psychophysiological meaning and were allowed to vary in dimension from person to person. Hence, we forwent analyses of the state space models' latent states. Our future work will investigate the use of inductive biases informed by domain knowledge and random effects terms to weakly enforce shared meaning onto portions of the latent state. Additional peripheral physiomarkers of the ANS (e.g., perspiratory markers) are known to be relevant to acute stress [2]. Unfortunately, the electrodermal activity signals collected in the clinical trial were of unsatisfactory signal quality for automated modeling and analyses. Variability markers (e.g., heart rate variability) computed using established, standard techniques produce time series with significantly slower dynamics than the beat-by-beath and breath-by-breath markers included in this work [7], [56]. We found that this mismatch in timescales results in unsatisfactory model validation results. Our future work will investigate the use of instantaneous variability metrics and more sophisticated but potentially less interpretable modeling approaches to overcome this limitation [57]. Note that our attention to signal quality, interpretability, and replicable methods are strengths of this work.

V. Conclusion

In this work, we modeled the effects of traumatic memories and tcVNS on multiple peripheral physiomarkers of the ANS simultaneously. We demonstrated that our person-specific state-space modeling approach holds predictive value, significantly outperforms baseline forecasting methods, and generates simulated data that match real data from 50 participants with histories of psychological trauma (24 with PTSD). Using the learned models, we simulated the effects of just-in-time tcVNS for participants with PTSD and show that tcVNS has the potential to attenuate the stress response to traumatic memories. These effects are most evident along the principal axis of variance, which explains the majority of nine physiomarkers' variance via a single dimension and replicates the directions of physiomarker changes expected psychophysiologically during changes in stress levels. Our methods advance the state of the art in modeling latent ANS dynamics from changes in peripheral physiomarkers. Our findings motivate continued engineering and research on closed-loop, non-invasive neuromodulation systems for acute stress mitigation. Future studies can evaluate the feasibility of automated nVNS in response to undesirable changes in ANS activity in lab settings and in daily life, as well as the acceptability of such systems for individuals with illnesses such as PTSD. The potential therapeutic value of closed-loop nVNS is compelling. Closed-loop systems for stress mitigation could improve quality of life for millions around the world living with trauma or anxiety disorders.

Data Availability Statement

De-identified biosignals, participant groups, and protocol data will be made available upon reasonable request from the authors. Modeling and analysis code will be made openly available upon publication: `https://github.com/asimgazi/MIMOlinearSSMs_ANSmodeling`.

References

[1] W. Jänig, *Integrative Action of the Autonomic Nervous System: Neurobiology of Homeostasis*. Cambridge University Press, 2006.

[2] G. Giannakakis *et al.*, "Review on psychological stress detection using biosignals," *IEEE Transactions on Affective Computing*, 2019.

[3] J. D. Bremner *et al.*, "Application of noninvasive vagal nerve stimulation to stress-related psychiatric disorders," vol. 10, 2020.

[4] R. K. Pitman *et al.*, "Biological studies of post-traumatic stress disorder," *Nature Reviews Neuroscience*, vol. 13, no. 11, pp. 769–787, 2012.

[5] E. Shvil *et al.*, "Neural, psychophysiological, and behavioral markers of fear processing in PTSD: a review of the literature," *Current psychiatry reports*, vol. 15, 5 2013.

[6] J. Blechert *et al.*, "Autonomic and respiratory characteristics of posttraumatic stress disorder and panic disorder," *Psychosomatic medicine*, vol. 69, pp. 935–943, 11 2007.

[7] A. H. Gazi *et al.*, "Quantifying posttraumatic stress disorder symptoms during traumatic memories using interpretable markers of respiratory variability," *IEEE Journal of Biomedical and Health Informatics*, 2024.

[8] A. J. Shah *et al.*, "Posttraumatic stress disorder and impaired autonomic modulation in male twins," *Biological psychiatry*, vol. 73, pp. 1103–1110, 6 2013.

[9] N. Pole, "The Psychophysiology of Posttraumatic Stress Disorder: A Meta-Analysis," *Psychological Bulletin*, vol. 133, pp. 725–746, 9 2007.

[10] J. D. Bremner *et al.*, "Transcutaneous cervical vagal nerve stimulation in patients with posttraumatic stress disorder (ptsd): A pilot study of effects on ptsd symptoms and interleukin-6 response to stress," *Journal of Affective Disorders Reports*, vol. 6, p. 100190, 2021.

[11] N. Z. Gurel *et al.*, "Quantifying acute physiological biomarkers of transcutaneous cervical vagal nerve stimulation in the context of psychological stress," *Brain stimulation*, vol. 13, pp. 47–59, 2020.

[12] N. Z. Gurel *et al.*, "Transcutaneous cervical vagal nerve stimulation reduces sympathetic responses to stress in posttraumatic stress disorder: A double-blind, randomized, sham controlled trial," *Neurobiology of stress*, vol. 13, 2020.

[13] A. H. Gazi *et al.*, "Digital cardiovascular biomarker responses to transcutaneous cervical vagus nerve stimulation: State-space modeling, prediction, and simulation," *JMIR mHealth and uHealth*, 2020.

[14] H. Y. Liu *et al.*, "Preoperative heart rate variability as predictors of vagus nerve stimulation outcome in patients with drug-resistant epilepsy," *Scientific Reports*, vol. 8, pp. 1–11, 2018.

[15] S. Debnath *et al.*, "A method to quantify autonomic nervous system function in healthy, able-bodied individuals," *Bioelectronic Medicine*, vol. 7, no. 1, 2021.

[16] A. Tavakoli, S. Boker, and A. Heydarian, "Driver state modeling through latent variable state space framework in the wild," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, pp. 1879–1893, 2 2023.

[17] A. H. Gazi *et al.*, "StressFADS: Learning Latent Autonomic Factors of Stress in the Context of Trauma Recall and Neuromodulation," *IEEE 20th International Conference on Body Sensor Networks (BSN)*, 2024.

[18] F. E. Rosas *et al.*, "Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics," *Computers in Biology and Medicine*, 2024.

[19] G. Valenza *et al.*, "Measures of sympathetic and parasympathetic autonomic outflow from heartbeat dynamics," *Journal of Applied Physiology*, vol. 125, 2018.

[20] Y. Kong, H. F. Posada-Quintero, and K. H. Chon, "Sensitive physiological indices of pain based on differential characteristics of electrodermal activity," *IEEE Transactions on Biomedical Engineering*, 2021.

[21] D. S. Wickramasuriya and R. T. Faghih, "A bayesian filtering approach for tracking arousal from binary and continuous skin conductance features," *IEEE Transactions on Biomedical Engineering*, 2020.

[22] R. Amin and R. T. Faghih, "Physiological characterization of electrodermal activity enables scalable near real-time autonomic nervous system activation inference," *PLOS Computational Biology*, 2022.

[23] S. Subramanian *et al.*, "Monitoring surgical nociception using multisensor physiological models," *Proceedings of the National Academy of Sciences of the United States of America*, 2024.

[24] A. L. Callara *et al.*, "Parasympathetic-sympathetic causal interactions assessed by time-varying multivariate autoregressive modeling of electrodermal activity and heart-rate-variability," *IEEE Transactions on Biomedical Engineering*, 2021.

[25] S. Shadi Ghiasi *et al.*, "Assessing autonomic function from electrodermal activity and heart rate variability during cold-pressor test and emotional challenge," *Scientific Reports 2020 10:1*, vol. 10, 2020.

[26] A. H. Gazi *et al.*, "Robust estimation of respiratory variability uncovers correlates of limbic brain activity and transcutaneous cervical vagus nerve stimulation in the context of traumatic stress," *IEEE Transactions on Biomedical Engineering*, vol. 69, 2022.

[27] J. D. Bremner *et al.*, "Neural correlates of memories of childhood sexual abuse in women with and without posttraumatic stress disorder," *American Journal of Psychiatry*, vol. 156, pp. 1787–1795, 1999.

[28] A. H. Gazi *et al.*, "Transcutaneous cervical vagus nerve stimulation inhibits the reciprocal of the pulse transit time's responses to traumatic stress in posttraumatic stress disorder," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1444–1447, 2021.

[29] J. D. Parreira *et al.*, "A proof-of-concept investigation of multi-modal physiological signal responses to acute mental stress," *Biomedical Signal Processing and Control*, 2023.

[30] A. H. Gazi *et al.*, "Transcutaneous cervical vagus nerve stimulation lengthens exhalation in the context of traumatic stress," in *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 1–4, 2021.

[31] C. Chourpiliadis and A. Bhardwaj, "Physiology, respiratory rate," *StatPearls*, 9 2022.

[32] A. D. Jose and D. Collison, "The normal range and determinants of the intrinsic heart rate in man," *Cardiovascular Research*, vol. 4, 1970.

[33] L. Ljung, *System Identification: Theory for the User*. Prentice Hall, second edi ed., 1999.

[34] C. Christophe Biernacki *et al.*, "Choosing starting values for the em algorithm for getting the highest likelihood in multivariate gaussian mixture models," *Computational Statistics and Data Analysis*, 2003.

[35] Y. Yang *et al.*, "Dynamic network modeling and dimensionality reduction for human ecog activity.," *Journal of neural engineering*, 2019.

[36] A. H. Gazi *et al.*, "Leveraging physiological markers to quantify the transient effects of traumatic stress and non-invasive neuromodulation," in *Proceedings of the 45th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2023.

[37] A. H. Gazi *et al.*, "Physiological markers reveal confounding effects of apprehension and habituation during stress protocol," in *Proceedings of the IEEE-EMBS International Conference on Body Sensor Networks (BSN)*, 2023.

[38] J. Z. Bakdash and L. R. Marusich, "Repeated measures correlation," *Frontiers in psychology*, vol. 8, 4 2017.

[39] M. T. Wittbrodt *et al.*, "Noninvasive cervical vagal nerve stimulation alters brain activity during traumatic stress in individuals with posttraumatic stress disorder," *Psychosomatic medicine*, vol. 83, 2021.

[40] J. D. Bremner *et al.*, "Transcutaneous vagal nerve stimulation blocks stress-induced activation of interleukin-6 and interferon- in posttraumatic stress disorder: A double-blind, randomized, sham-controlled trial," *Brain, Behavior, and Immunity - Health*, vol. 9, p. 100138, 2020.

[41] J. A. Sanchez-Perez *et al.*, "Transcutaneous auricular vagus nerve stimulation and median nerve stimulation reduce acute stress in young healthy adults: a single-blind sham-controlled crossover study," *Frontiers in Neuroscience*, vol. 17, p. 1213982, 2023.

[42] S. Sundararaj *et al.*, "Accrued Reductions in Heart Rate Following Transcutaneous Vagal Nerve Stimulation in Adults with Posttraumatic Stress Disorder," *Frontiers in Neuroscience*, 2025.

[43] S. Paganini *et al.*, "Stress management apps: Systematic search and multidimensional assessment of quality and characteristics," *JMIR mHealth and uHealth*, 2023.

[44] M. McKay, J. C. Wood, and J. Brantley, *The dialectical behavior therapy skills workbook : practical DBT exercises for learning mindfulness, interpersonal effectiveness, emotion regulation*. New Harbinger Publications, Inc., 2019.

[45] I. Nahum-Shani *et al.*, "Just-in-Time Adaptive Interventions (JITAIs) in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support," *Annals of Behavioral Medicine*, vol. 52, pp. 446–462, 2018.

[46] C. M. Raio *et al.*, "Cognitive emotion regulation fails the stress test," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 110, pp. 15139–15144, 2013.

[47] B. W. Badran *et al.*, "A pilot randomized controlled trial of supervised, at-home, self-administered transcutaneous auricular vagus nerve stimulation (tavns) to manage long covid symptoms," *Bioelectronic Medicine*, vol. 8, pp. 1–10, 12 2022.

[48] B. W. Badran *et al.*, "Motor activated auricular vagus nerve stimulation as a potential neuromodulation approach for post-stroke motor rehabilitation: A pilot study," *Neurorehabilitation and Neural Repair*, vol. 37, pp. 374–383, 6 2023.

[49] R. A. Bryant, "Post-traumatic stress disorder: a state-of-the-art review of evidence and challenges," *World Psychiatry*, vol. 18, p. 259, 10 2019.

[50] P. Cuijpers *et al.*, "A network meta-analysis of the effects of psychotherapies, pharmacotherapies and their combination in the treatment of adult depression," *World Psychiatry*, 2020.

[51] K. Hovsepian *et al.*, "cstress: Towards a gold standard for continuous stress assessment in the mobile environment.," *Proceedings of the 2015 ACM International Conference on Pervasive and Ubiquitous Computing*, vol. 2015, pp. 493–504, 2015.

[52] A. E. Bryson and Y.-C. Ho, "Applied optimal control : Optimization, estimation and control," *Applied Optimal Control*, 5 2018.

[53] C. E. García, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice—a survey," *Automatica*, vol. 25, 1989.

[54] D. J. Lin *et al.*, "Real-time seismocardiogram feature extraction using adaptive gaussian mixture models," *IEEE Journal of Biomedical and Health Informatics*, 2023.

[55] E. Katsoulakis *et al.*, "Digital twins for health: a scoping review," *npj Digital Medicine*, 2024.

[56] F. Shaffer and J. P. Ginsberg, "An overview of heart rate variability metrics and norms," *Frontiers in Public Health*, vol. 5, p. 258, 9 2017.

[57] R. Barbieri *et al.*, "A point-process model of human heartbeat intervals: new definitions of heart rate and heart rate variability.," *American journal of physiology. Heart and circulatory physiology*, vol. 288, 2005.