

UNITED STATES PATENT AND TRADEMARK OFFICE

PROVISIONAL APPLICATION FOR PATENT

Under 35 U.S.C. § 111(b) and 37 C.F.R. § 1.53(c)

**SYSTEM AND METHOD FOR CLOSED-LOOP
BIOMETRIC-DRIVEN ENVIRONMENTAL CONTROL
IN NETWORKED HABITATION UNITS WITH
FLEET-LEVEL LONGITUDINAL LEARNING**

Sole Inventor:

Sophie Defauw

Quintana Roo, Mexico

Applicant / Assignee:

BEYOND HUMAN STUDIO, INC.

1111B S Governors Ave # 97566

Dover, DE 19904

Date of Filing: April 21 2026

TABLE OF CONTENTS

I.	Title of the Invention
II.	Cross-Reference to Related Applications
III.	Statement Regarding Federally Sponsored Research
IV.	Field of the Invention
V.	Background of the Invention
VI.	Brief Summary of the Invention
VII.	Definitions and Terminology
VIII.	Brief Description of the Figures
IX.	Detailed Description of the Preferred Embodiments
X.	Sensor Fusion and Artificial Intelligence Inference Architecture
XI.	Safety-Preemption Control Hierarchy
XII.	Dual-Logging Architecture and Algorithmic Independence Pathway
XIII.	Wearable-Agnostic Gateway Embodiment
XIV.	Fleet-Level Learning and Causal Dataset Architecture
XV.	Behavioral-State Input Layer
XVI.	Alternative Embodiments and Equivalents
XVII.	Illustrative Examples
XVIII.	Claims
XIX.	Abstract of the Disclosure
XX.	Figures (FIG. 1 through FIG. 12)

I. TITLE OF THE INVENTION

System and method for closed-loop biometric-driven environmental control in networked habitation units with fleet-level longitudinal learning.

II. CROSS-REFERENCE TO RELATED APPLICATIONS

Not applicable

III. STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

Not applicable. The invention disclosed herein was conceived and developed without federal funding, and no U.S. Government agency has rights in the invention under 35 U.S.C. §§ 200–212.

IV. FIELD OF THE INVENTION

The present invention relates generally to the technical field of closed-loop cyber-physical systems integrating wearable biometric sensing with multi-actuator environmental control, and more particularly, to networked modular habitation units equipped with embedded edge-compute processors that execute sensor fusion and artificial intelligence (AI) inference across biometric, environmental, and behavioral data streams to dynamically modulate thermal, ventilation, and photic subsystems, while contributing labeled interventional observations to a fleet-level training corpus distributed across a plurality of geographically distinct deployment sites.

V. BACKGROUND OF THE INVENTION

A. The Technical Problem

Human physiological recovery — comprising sleep architecture, autonomic nervous system homeostasis, and cognitive load dissipation — is governed by a continuous interaction between endogenous physiological state and exogenous environmental variables including, without limitation, ambient temperature, relative humidity, airborne particulate concentration, carbon dioxide concentration, volatile organic compound concentration, photic spectral composition, and acoustic pressure level. Despite widespread recognition of this coupling in the clinical, sleep-medicine, and chronobiology literature, existing technical systems fail to close the feedback loop between the physiological state of an occupant and the environmental parameters of the occupant's habitation.

B. Deficiencies of the Prior Art

The prior art may be grouped into three principal categories, each of which suffers from structural limitations that the present invention is designed to overcome.

First, consumer-grade wearable sensing devices, including without limitation wrist-worn activity monitors, smart rings, and chest-strap heart-rate monitors, provide observational physiological data streams to the wearer via a companion application, but possess no capacity to modify the ambient environment in which the wearer resides. Such devices thereby generate observations without interventions, yielding datasets that are structurally incapable of supporting causal inference.

Second, conventional environmental-control systems, including programmable thermostats, heating-ventilation-and-air-conditioning (HVAC) controllers, building management systems, and connected lighting platforms, actuate environmental parameters pursuant to fixed schedules, pre-configured setpoints, or coarse-grained occupancy detection, wholly without reference to the physiological state of the occupant. Such systems optimize a static or demographic-level comfort model, not the real-time autonomic state of a particular individual.

Third, recent closed-loop systems combining biometric sensing with environmental actuation — including, by way of example and not limitation, temperature-regulating sleep surfaces employing hydronic or thermoelectric cooling coupled to sleep-stage inference (as exemplified by U.S. Patent Application Publication No. 2019/0320808 A1); wearable-to-HVAC control architectures that adjust ambient conditions upon detecting threshold deviations in biometric parameters (as exemplified by U.S. Patent Application Publication No. 2026/0036317); remote biometric monitoring systems coupling physiological sensing to environmental actuators (as exemplified by U.S. Patent Application Publication No. 2016/0364617 A1); and wearable thermal-stimulation devices delivering closed-loop warming or cooling stimuli in response to physiological signals — are each constrained to a single-device or single-room deployment architecture, operate as isolated control loops without inter-deployment learning, and are structurally incapable of accumulating a causal interventional dataset at a population scale.

C. The Unmet Need

There accordingly exists an unmet need in the art for a system and method that: (i) continuously integrates physiological, environmental, and behavioral data streams within a naturalistic residential context; (ii) actuates thermal, ventilatory, and photic subsystems in real time pursuant to fused multi-modal inference; (iii) operates with sub-second deterministic latency at a locally instantiated edge-compute processor; (iv) contributes labeled interventional feedback to a coordinated learning fleet spanning a plurality of deployment sites; (v) supports programmatic substitution of third-party biometric signal-processing components with proprietary on-device inference models trained on accumulated longitudinal data, without disruption to downstream consumers of the structured data schema; and (vi) enforces hierarchically prioritized safety constraints, including without limitation a minimum

blood-oxygen-saturation floor and a maximum ambient carbon-dioxide ceiling, that preempt lower-priority objectives such as acoustic quietness and energy efficiency.

The present invention addresses and resolves each of the foregoing unmet needs, as set forth in the detailed description below.

VI. BRIEF SUMMARY OF THE INVENTION

The present invention provides a system, a method, and an apparatus for closed-loop biometric-driven environmental control in networked habitation units with fleet-level longitudinal learning. The invention comprises, in its preferred and broadest embodiments, the following seven cooperating aspects, each of which independently contributes to the technical advance over the prior art and each of which is separately claimable.

Aspect I. Four-Layer Data Pipeline Architecture

In a first aspect, the invention provides a four-layer data pipeline architecture comprising (1) a Body-Proximate Sensing Device (hereinafter the "BPSD") worn by or positioned proximate to an occupant and producing a continuous physiological data stream; (2) a Middleware Layer performing data synchronization, parsing, normalization, and structured-schema translation of the physiological data stream; (3) a Controlled Habitation Volume (hereinafter the "CHV") incorporating an ambient sensor array, an actuator ensemble, and an embedded Edge Compute Processor (hereinafter the "ECP") executing real-time sensor fusion and AI inference; and (4) a Fleet Management Layer in data communication with a plurality of CHV instances and configured to aggregate labeled feedback observations, refine a shared inference model, and redistribute refined model parameters across the fleet.

Aspect II. Autonomic-State-Aware Thermal Control Model

In a second aspect, the invention provides an Autonomic-State-Aware Thermal Control Model (hereinafter the "ASTCM") executing at the ECP that computes a personalized thermal setpoint as a function of, at minimum, the occupant's heart-rate-variability-derived autonomic state, skin temperature, ambient temperature, and ambient relative humidity. The ASTCM thereby generates different thermal setpoints for a first occupant in a state of elevated sympathetic tone and a second occupant in a state of parasympathetic dominance under identical ambient conditions, representing a specific technical improvement over ambient-temperature-only or demographic-preference thermostatic control.

Aspect III. Zonal Thermal Gradient Shaping

In a third aspect, the invention provides a zonal thermal gradient shaping subsystem comprising a plurality of independently addressable solid-state thermal elements distributed spatially within the CHV interior. The ECP commands each element independently pursuant to the ASTCM output, producing a

shaped spatial thermal gradient across the occupant body — including without limitation a head-cooled/foot-warmed gradient promoting sleep-onset physiology — determined at least in part by occupant body position inferred from an inertial sensor of the BPSD.

Aspect IV. Safety-Priority Control Hierarchy (SICRH)

In a fourth aspect, the invention provides a Sensor-Input Conflict Resolution Hierarchy (hereinafter the "SICRH") governing arbitration among competing actuator directives generated by the AI inference engine. The SICRH enforces a deterministic priority ordering in which a blood-oxygen-saturation safety floor preempts all other objectives; an ambient carbon-dioxide ceiling preempts all non-safety objectives; and remaining comfort, acoustic, and energy-efficiency objectives are arbitrated in descending priority order. The SICRH enables predictable, locally executable safety behavior without dependence on cloud arbitration or network connectivity.

Aspect V. Dual-Logging Algorithmic Independence Pathway

In a fifth aspect, the invention provides a dual-logging architecture that simultaneously records raw sensor waveforms from the BPSD and processed biometric outputs emitted by a biometric signal-processing component (whether a third-party software development kit or a proprietary inference model), generating a paired raw-plus-processed dataset across each occupant-night that constitutes a structured training corpus for subsequent development of substitute on-device inference models. The substitute inference models, upon training and validation, are deployed in schema-preserving substitution for the incumbent signal-processing component without modification to downstream consumers of the biometric data stream.

Aspect VI. Two-Tier Longitudinal Personalization Engine

In a sixth aspect, the invention provides a Longitudinal Personalization Engine (hereinafter the "LPE") implementing a two-tier model architecture: (a) a population-level prior model trained on anonymized aggregated data from a plurality of CHV units and deployment sites; and (b) an individual-level posterior model maintained per occupant and updated from the occupant's accumulated biometric-to-actuator-to-outcome records via Bayesian updating or functionally equivalent online learning methods. The LPE supports portability of the individual-level posterior across CHV instances and deployment sites.

Aspect VII. Causal Fleet Dataset and Continuous-Experimentation Protocol

In a seventh aspect, the invention provides a fleet-level causal dataset architecture in which each environmental actuation decision emitted by the ECP constitutes a controlled experimental intervention with a measured downstream physiological response. Aggregated records across a plurality of CHV instances form a ground-truth causal dataset linking environmental independent variables to physiological

recovery outcome dependent variables at longitudinal population scale — a dataset architecturally inaccessible to observational wearable-only systems, observational clinical-registry systems, or uncoordinated smart-home deployments.

The foregoing aspects cooperate synergistically. The ASTCM and zonal gradient control (Aspects II–III) produce differentiated actuation decisions; the SICRH (Aspect IV) ensures safety-bounded execution; the dual-logging pathway (Aspect V) captures the training corpus; the LPE (Aspect VI) distills the corpus into personalization state; and the fleet dataset architecture (Aspect VII) provides the cross-site statistical power required to separate causation from correlation. The four-layer pipeline (Aspect I) instantiates the system in a manner capable of operating, maintaining, and scaling across arbitrary numbers of CHV instances and deployment sites.

VII. DEFINITIONS AND TERMINOLOGY

For purposes of the present disclosure and the appended claims, and except where the context clearly requires otherwise, the following terms shall have the meanings ascribed to them in this Section VII. Terms defined herein shall carry the defined meaning throughout the specification, in the claims, and in any subsequent non-provisional, continuation, continuation-in-part, or divisional application claiming priority hereto. Where a defined term is followed by a specific hardware reference or commercial component, such reference shall be understood as illustrative of a preferred embodiment and shall not be construed to limit the scope of the defined term or of any claim employing such term.

"Body-Proximate Sensing Device (BPSD)" — any wearable, contact-based, or non-contact sensing apparatus configured to continuously measure a plurality of physiological parameters of a human occupant, including without limitation wrist-worn devices, finger-worn devices, chest-strap devices, in-ear devices, sleep-surface-embedded ballistocardiographic sensors, millimeter-wave or ultra-wideband radar sensors, and thermal-infrared imaging sensors.

"Controlled Habitation Volume (CHV)" — a bounded interior spatial volume sized to accommodate at minimum one human occupant and equipped with an ambient sensor array and an actuator ensemble under the control of an embedded Edge Compute Processor; also referred to herein, and in the specification of Beyond Human Studio, Inc., as the "capsule." The CHV need not be fully enclosed, need not be structurally independent, and may comprise a sub-volume within a larger structure.

"Edge Compute Processor (ECP)" — a microprocessor or system-on-chip physically co-located within or with the CHV, configured to receive data from the BPSD and the ambient sensor array, execute sensor fusion and inference computations, and emit actuator control signals, all without requiring real-time network connectivity to a cloud service for primary control functions.

- "Fleet Management Layer (FML)"** — a cloud-hosted or self-hosted computing infrastructure in data communication with a plurality of CHV instances, configured to aggregate labeled feedback records, train and refine inference models, and redistribute refined model parameters to the plurality of CHV instances.
- "Autonomic-State-Aware Thermal Control Model (ASTCM)"** — a computational model, executing on the ECP, that computes a thermal setpoint or setpoint vector as a function of at minimum the occupant's heart-rate-variability-derived autonomic state, skin temperature, ambient temperature, and ambient relative humidity.
- "Ventilation Demand Model (VDM)"** — a computational model, executing on the ECP, that computes a commanded airflow rate for a filtered ventilation subsystem as a function of at minimum the occupant's respiration rate and blood-oxygen saturation, the ambient carbon-dioxide concentration, and the ambient volatile-organic-compound concentration.
- "Sensor-Input Conflict Resolution Hierarchy (SICRH)"** — a deterministic priority ordering, executing on the ECP, that arbitrates among competing actuator directives generated by the AI inference engine, enforcing a descending priority ordering in which safety-critical objectives preempt comfort-oriented objectives and comfort-oriented objectives preempt energy-efficiency objectives.
- "Longitudinal Personalization Engine (LPE)"** — a two-tier inference architecture comprising (a) a population-level prior model and (b) an individual-level posterior model updated from occupant-specific accumulated feedback records, such that control decisions issued to a particular occupant reflect both population-level invariants and individual idiosyncrasies.
- "Occupant-Night"** — a unit of longitudinal data corresponding to a continuous interval during which a single identified occupant is resident within a single CHV instance and the biometric-to-actuator feedback loop is operative. Aggregated occupant-nights constitute the primary unit of accounting for the fleet-level dataset.
- "Labeled Feedback Tuple"** — a structured record comprising at minimum: (i) a timestamp; (ii) an anonymized occupant identifier; (iii) a fused input vector derived from the BPSD, the ambient sensor array, and optionally a behavioral-state vector; (iv) a control decision emitted by the AI inference engine; (v) any safety-preemption overrides applied by the SICRH; (vi) an ambient-response trajectory observed subsequent to the control decision; (vii) a biometric-response trajectory observed subsequent to the control decision; and (viii) one or more downstream recovery metrics computed at the close of the occupant-night.
- "Schema-Preserving Substitution"** — the act of replacing a first biometric signal-processing component with a second biometric signal-processing component wherein both components emit

processed biometric outputs conforming to an identical structured data schema, such that substitution is transparent to downstream consumers of said outputs.

"Hardware-Agnostic Operating Mode" — an operating mode of the CHV in which the ECP ingests biometric data via a wireless radio gateway embedded within the CHV, said gateway configured to discover, pair with, and ingest raw radio profile payloads from any BPSD conforming to a standardized wireless biometric profile, without dependence on a specific BPSD vendor or a specific signal-processing component thereof.

VIII. BRIEF DESCRIPTION OF THE FIGURES

The accompanying figures, which are incorporated in and constitute a part of this specification, illustrate several embodiments of the invention and, together with the description, serve to explain the principles of the invention. The figures are to be prepared as formal patent drawings, in substantial conformance with 37 C.F.R. § 1.84, prior to conversion of this Provisional Application to a non-provisional application. The textual descriptions set forth in this Section VIII are of sufficient particularity to constitute written description of the embodiments depicted under 35 U.S.C. § 112(a).

FIG. 1 depicts a system-level block diagram of the four-layer data pipeline architecture, including the BPSD Layer, the Middleware Layer, the CHV Layer incorporating the ECP, and the Fleet Management Layer, together with the directed data flows among said layers.

FIG. 2 depicts a functional block diagram of a preferred BPSD embodiment, including the microcontroller unit integrating a wireless radio, the photoplethysmographic (PPG) optical sensor assembly, the three-axis accelerometer, the skin-temperature sensor, the local flash memory, the rechargeable power cell, and the wireless radio interface.

FIG. 3 depicts a functional block diagram of a preferred CHV embodiment, including the ambient sensor array (comprising carbon-dioxide, volatile-organic-compound, temperature, relative-humidity, inertial-measurement, and surface-thermal sensors), the ECP, and the actuator ensemble (comprising the solid-state thermal element array, the filtered ventilation subsystem, and the tunable-spectrum photic subsystem).

FIG. 4 depicts a detailed data-flow diagram of the Middleware Layer, illustrating the translation of raw BPSD transmissions into the structured biometric data schema consumed by the ECP.

FIG. 5 depicts a flowchart of the sensor fusion and AI inference process executed by the ECP at each control interval, including the ingestion of biometric, ambient, and behavioral state vectors; formation of the fused input vector; application of the inference model; and emission of the control decision.

FIG. 6 depicts a flowchart of the SICRH, illustrating the preemption relationships among the blood-oxygen-saturation safety floor, the ambient carbon-dioxide ceiling, the thermal comfort objective, the acoustic-quietness objective, and the energy-efficiency objective.

FIG. 7 depicts a data-flow diagram of the dual-logging architecture, showing the parallel capture of raw sensor waveforms (Stream A) and processed biometric outputs (Stream B), the indexing of both streams by occupant-night, and the subsequent use of the paired dataset for training a substitute on-device inference model.

FIG. 8 depicts a state diagram illustrating the transition between a first operating mode, in which the ECP consumes biometric data via a third-party signal-processing component hosted on the Middleware Layer, and a second Hardware-Agnostic Operating Mode, in which the ECP consumes biometric data via an embedded wireless radio gateway within the CHV.

FIG. 9 depicts a schematic cross-sectional view of the CHV interior, illustrating the spatial distribution of the solid-state thermal element array and the resulting shaped thermal gradient across the occupant body along a head-to-foot axis.

FIG. 10 depicts a functional block diagram of the Fleet Management Layer, illustrating the aggregation of Labeled Feedback Tuples across a plurality of CHV instances, the model-refinement pipeline, and the over-the-air redistribution of refined model parameters.

FIG. 11 depicts a flowchart of the BPSD asset-rotation and sanitation protocol executed upon occupant transition within a given CHV instance.

FIG. 12 depicts a functional block diagram of an alternative BPSD embodiment employing non-contact occupant sensing modalities, including without limitation ballistocardiographic, millimeter-wave radar, ultra-wideband radar, and thermal-infrared imaging modalities.

IX. DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The following detailed description sets forth preferred embodiments of the present invention. The description is provided to enable any person skilled in the relevant art to practice the invention. Specific embodiments are described by way of illustration and not of limitation; various modifications, substitutions, and functional equivalents are expressly contemplated and fall within the scope of the invention. Where a specific hardware component, material, dimensional parameter, or operational value is identified in connection with a preferred embodiment, functionally equivalent alternatives shall be understood as expressly contemplated within the scope of the invention, and no such reference shall be construed to limit the scope of any claim.

A. System Overview

Referring to FIG. 1, a system (100) in accordance with the present invention comprises a Body-Proximate Sensing Device (200), a Middleware Layer (300), a Controlled Habitation Volume (400) incorporating an Edge Compute Processor (420), and a Fleet Management Layer (500) in data communication with a plurality of Controlled Habitation Volumes (400a, 400b, ..., 400n). The BPSD (200) produces a continuous biometric data stream. The Middleware Layer (300) normalizes said stream into a structured schema. The ECP (420) ingests the normalized stream together with data from an ambient sensor array (410) disposed within the CHV (400), executes sensor fusion and AI inference, and emits control signals to an actuator ensemble (430). The FML (500) aggregates Labeled Feedback Tuples from the plurality of CHVs, refines a shared inference model, and redistributes refined parameters to the fleet.

B. Body-Proximate Sensing Device

Referring to FIG. 2, in a preferred embodiment the BPSD (200) comprises: a microcontroller unit (210) integrating a Bluetooth Low Energy radio of version 5.4 or functional equivalent; a photoplethysmographic sensor assembly (220) including at least one light-emitting diode operating in a green, red, or infrared wavelength band and at least one photodiode; a three-axis accelerometer (230); a skin-temperature sensor (240); a flash memory (250) for local buffering of sensor data during transient wireless disconnection and for firmware over-the-air update staging; a rechargeable power cell (260); a housing (270); and a band or strap (280). In the preferred embodiment, the microcontroller unit (210) comprises a Nordic Semiconductor nRF54L20 or functional equivalent; this reference is illustrative of a preferred commercial embodiment and shall not be construed to limit the scope of the invention. Functionally equivalent microcontroller units from other vendors, including without limitation devices from the Espressif ESP32 family, the Dialog Semiconductor DA14xxx family, and the STMicroelectronics STM32WB family, are expressly contemplated. The wireless radio is not limited to Bluetooth Low Energy; functionally equivalent protocols including IEEE 802.15.4, Thread, Wi-Fi (IEEE 802.11 family), and ultra-wideband (UWB) are expressly contemplated.

B.1 Substitutable Biometric Signal-Processing Component

The BPSD (200) further comprises a biometric signal-processing component operable to convert raw sensor waveforms into processed biometric outputs conforming to a structured data schema. In a first operating mode, the biometric signal-processing component comprises firmware, a software development kit, or a library provided by a third-party original design manufacturer of the BPSD and executing upon the microcontroller unit (210), upon the Middleware Layer (300), or partitioned between both. In a second operating mode, said component is replaced in whole or in part by a proprietary on-device inference model, trained pursuant to the methodology set forth in Section XII *infra*. The structured data schema emitted by the biometric signal-processing component is preserved identical across the first and second operating modes, thereby effecting Schema-Preserving Substitution.

C. Controlled Habitation Volume

Referring to FIG. 3, in a preferred embodiment the CHV (400) comprises a structural enclosure (401) defining an interior volume sized to accommodate a single adult occupant in a recumbent position; a door or closure (402); an electrical power interface (403); a data interface (404); the ambient sensor array (410); the ECP (420); and the actuator ensemble (430). Nothing in this Detailed Description shall be construed to limit the CHV to a free-standing structure; sub-volumes within larger structures, including without limitation bedroom sub-volumes, private-office sub-volumes, and partitioned hospitality-accommodation sub-volumes, are expressly contemplated.

C.1 Ambient Sensor Array

The ambient sensor array (410) comprises, in a preferred embodiment, the following sensing elements:

- A carbon-dioxide sensor (411), preferably a non-dispersive infrared (NDIR) type sensor;
- A volatile-organic-compound sensor (412);
- An ambient-temperature sensor (413);
- A relative-humidity sensor (414);
- An inertial measurement unit (415) for occupancy, door-event, and structural-event detection;
- A plurality of surface-temperature sensors (416a – 416n) for thermal-gradient mapping; and
- Optionally, an acoustic-level sensor (417), a particulate-matter sensor (418), and an illuminance sensor (419).

C.2 Edge Compute Processor

The ECP (420) comprises a microprocessor (421), a memory (422), a data-interface subsystem (423), and an actuator-driver subsystem (424). In a preferred embodiment, the microprocessor (421) comprises an ARM Cortex-M4F class microcontroller with integrated floating-point unit, selected for real-time determinism, low idle power consumption, direct peripheral control over actuators, and cost efficiency at fleet scale. Alternative embodiments employing higher-capability microprocessors, including without limitation ARM Cortex-M7, ARM Cortex-A class application processors, x86 single-board computers, and neural processing units, are expressly contemplated where inference workload requirements exceed the computational envelope of the preferred embodiment.

C.3 Actuator Ensemble

The actuator ensemble (430) comprises, in a preferred embodiment: (i) a thermal subsystem (431) comprising a plurality of independently addressable solid-state thermal elements; (ii) a filtered ventilation subsystem (432) comprising a variable-speed fan and a high-efficiency particulate-air (HEPA) filter of H13 rating or higher; and (iii) a tunable-spectrum photic subsystem (433) comprising a plurality of

light-emitting diodes of differing correlated color temperatures and operable, in combination, to vary correlated color temperature continuously across a range of at least 2700 kelvin to 6500 kelvin and to vary intensity continuously from a minimum to a maximum level.

C.4 Zonal Thermal Gradient Shaping

In a preferred embodiment, the solid-state thermal elements of the thermal subsystem (431) comprise thermoelectric (Peltier-effect) devices. The plurality of thermoelectric devices is distributed spatially within the CHV interior such that a first subset is positioned proximal to a head region of an occupant in a recumbent position, a second subset is positioned proximal to a torso region, and a third subset is positioned proximal to a foot region. The ECP commands each subset independently pursuant to the ASTCM (see Section X.A), producing a shaped spatial thermal gradient across the occupant body. Thermal insulation, directional airflow channels, or functionally equivalent means may be employed to reduce inter-zone thermal cross-talk. Alternative solid-state thermal actuators, including without limitation flexible thermoelectric films, magnetocaloric modules, and electrocaloric modules, are expressly contemplated, as are non-solid-state thermal actuators employing vapor-compression cycles, absorption cycles, or circulating-fluid systems in form factors compatible with the CHV.

X. SENSOR FUSION AND ARTIFICIAL INTELLIGENCE INFERENCE ARCHITECTURE

At each control interval, the ECP (420) executes a sensor fusion and AI inference procedure comprising, in sequence: (1) ingestion of the normalized biometric state vector from the Middleware Layer (300); (2) ingestion of the ambient state vector from the ambient sensor array (410); (3) optional ingestion of a behavioral-state vector as described in Section XV; (4) formation of a fused input vector; (5) application of an inference model to said fused input vector; (6) emission of a control decision comprising at least a per-zone thermal setpoint vector, a ventilation flow-rate command, and a photic setpoint comprising correlated color temperature and intensity; (7) application of the SICRH to the control decision; (8) transmission of the final control decision to the actuator-driver subsystem; and (9) persistence of a Labeled Feedback Tuple to local storage for subsequent synchronization with the FML. The inference model may be instantiated as a rule-based heuristic system, a feed-forward neural network, a one-dimensional convolutional neural network, a recurrent neural network, a transformer-based sequence model, a gradient-boosted decision-tree ensemble, a Kalman-filter-based state estimator, or any ensemble or combination of the foregoing, and is stored in the ECP memory (422) in a quantized representation suitable for microcontroller-class inference.

A. Autonomic-State-Aware Thermal Control Model (ASTCM)

The ASTCM produces a personalized thermal setpoint, or a vector of zonal thermal setpoints, as a function of the fused input vector. Without limitation, the ASTCM accounts for

heart-rate-variability-derived autonomic state features (including without limitation the root-mean-square of successive differences (RMSSD), the proportion of successive NN-intervals exceeding fifty milliseconds (pNN50), and the ratio of low-frequency to high-frequency spectral power of the inter-beat-interval series), skin temperature, recent activity state, ambient temperature, and ambient relative humidity. The ASTCM thereby produces, in general, distinct thermal setpoints for two occupants under identical ambient conditions but differing in autonomic state — a determination structurally unattainable by ambient-only thermostatic control or by demographic-level comfort modelling.

Without admitting the necessity of any particular functional form, the ASTCM may be represented schematically as:

$$T_{\text{setpoint}} = f_{\text{ASTCM}}(\text{HRV_features}, T_{\text{skin}}, A_{\text{state}}, T_{\text{ambient}}, \text{RH}_{\text{ambient}})$$

wherein the functional form f_{ASTCM} is learned from the longitudinal dataset described in Section XIV, and the specific parametric form, feature weighting, and training methodology are not limited to any particular embodiment.

B. Ventilation Demand Model (VDM)

The VDM produces a commanded airflow rate for the filtered ventilation subsystem (432) as a function of the occupant's respiration rate and blood-oxygen saturation, the ambient carbon-dioxide concentration, and the ambient volatile-organic-compound concentration. The VDM operates subject to the SICRH described in Section XI infra, such that a measured decline in blood-oxygen saturation toward or below a configured safety floor causes immediate elevation of the commanded airflow rate, notwithstanding any concurrent acoustic-quietness objective.

C. Sleep-Stage-Aware Photic and Environmental Modulation

In a further embodiment, the inference engine modulates the thermal, ventilation, and photic commands as a function of an inferred sleep stage of the occupant. The inference engine learns, over successive occupant-nights, which environmental configurations correlate with extended deep-sleep duration, extended REM-sleep duration, and elevated overnight heart-rate-variability in the particular occupant, and biases subsequent control decisions toward said configurations subject always to the SICRH.

D. Recovery Scoring and Anticipatory Pre-Conditioning

In a further embodiment, the inference engine computes a daily recovery index from the cumulative biometric history of the occupant and executes anticipatory pre-conditioning of the CHV environment in advance of forecasted occupancy events. Pre-conditioning may include, without limitation, pre-cooling the CHV in advance of expected sleep-onset, or adjusting photic setpoints in advance of expected wake.

XI. SAFETY-PREEMPTION CONTROL HIERARCHY

Referring to FIG. 6, the ECP applies the Sensor-Input Conflict Resolution Hierarchy (SICRH) to every control decision prior to transmission to the actuator-driver subsystem. The SICRH enforces the following deterministic priority ordering, in descending order of precedence:

- (1) Blood-oxygen-saturation safety floor. Upon detection of a measured blood-oxygen-saturation value at or below a configured floor, or trending toward said floor at a rate exceeding a configured rate threshold, the ECP commands the filtered ventilation subsystem to an airflow rate sufficient to restore and maintain blood-oxygen saturation to within configured tolerance, without regard to any acoustic-quietness objective otherwise in force.
- (2) Ambient carbon-dioxide ceiling. Upon detection of an ambient carbon-dioxide concentration at or above a configured ceiling, the ECP commands the filtered ventilation subsystem to an elevated airflow rate.
- (3) Ambient volatile-organic-compound ceiling. Upon detection of an ambient volatile-organic-compound concentration at or above a configured ceiling, the ECP commands the filtered ventilation subsystem to an elevated airflow rate.
- (4) Thermal comfort objective. Subject to the preceding three priorities, the ECP actuates the thermal subsystem pursuant to the ASTCM.
- (5) Sleep-stage-aware acoustic-quietness objective. Subject to the preceding four priorities, the ECP minimizes actuator acoustic emissions during inferred sleep intervals.
- (6) Energy-efficiency objective. Subject to the preceding five priorities, the ECP minimizes electrical power consumption.

The SICRH is implemented deterministically at the ECP and does not require network connectivity to the FML for its operation. The SICRH therefore continues to enforce safety-critical constraints during transient or sustained loss of network connectivity. The SICRH is separately and independently claimable as a method and as an apparatus embodiment.

XII. DUAL-LOGGING ARCHITECTURE AND ALGORITHMIC INDEPENDENCE PATHWAY

Referring to FIG. 7, the BPSD (200), the Middleware Layer (300), and the FML (500) are jointly configured to perform dual logging of two parallel data streams, designated herein as Stream A and Stream B.

Stream A (Raw Sensor Waveforms): raw analog-to-digital converter sample sequences from each sensor of the BPSD, including without limitation the photoplethysmographic optical sensor, the three-axis accelerometer, and the skin-temperature sensor, captured at the native sampling rates of the respective sensors and stored in the BPSD flash memory (250) and/or the Middleware Layer for subsequent synchronization with the FML.

Stream B (Processed Biometric Outputs): structured biometric outputs generated from Stream A by the incumbent biometric signal-processing component, including without limitation heart rate, heart rate variability, blood-oxygen saturation, sleep stage classification, stress score, skin temperature, and respiration rate, conforming to the structured data schema.

Stream B feeds the real-time ECP control system and the user-facing application layer. Stream A and Stream B are stored in paired form, indexed by occupant-night and by timestamp, constituting a labeled training corpus in which Stream B outputs serve as ground-truth labels for corresponding Stream A inputs.

Upon accumulation of a training corpus of sufficient cardinality and diversity, the corpus supports the development and validation of substitute on-device biometric inference models trained to reproduce the Stream B outputs from the Stream A inputs. Suitable model architectures include, without limitation, one-dimensional convolutional neural networks for temporal waveform processing, compact transformer architectures for sequence modeling, and recurrent neural network architectures. The substitute models are compressed for microcontroller-class deployment via quantization, pruning, knowledge distillation, or functionally equivalent techniques, and may be implemented using frameworks including without limitation TensorFlow Lite Micro, CMSIS-NN, ONNX Runtime for microcontrollers, or functionally equivalent frameworks.

Upon validation, the substitute on-device inference model is deployed to the BPSD firmware, to the Middleware Layer, or to both, via over-the-air firmware update, in Schema-Preserving Substitution for the incumbent biometric signal-processing component. Because the structured data schema is preserved, downstream consumers of Stream B, including the ECP inference engine (Section X) and the FML (Section XIV), require no modification whatsoever. This Schema-Preserving Substitution property constitutes a distinct and independently claimable aspect of the invention.

A. Optional Clinical-Grade Validation

In a further embodiment, a subset of the paired Stream A / Stream B corpus is collected in parallel with reference measurements obtained from independent medical-grade reference instruments, including without limitation twelve-lead electrocardiography, clinical-grade pulse oximetry, and polysomnography. The substitute on-device inference model is then validated against said reference measurements to establish independent accuracy benchmarks, thereby avoiding exclusive reliance upon the incumbent

signal-processing component as ground truth and establishing a documentary basis for subsequent regulatory submissions, should such submissions be pursued.

XIII. WEARABLE-AGNOSTIC GATEWAY EMBODIMENT

Referring to FIG. 8, an alternative embodiment eliminates the dependency of the system on any particular BPSD or its signal-processing component. In said embodiment, the CHV (400) incorporates an embedded wireless radio gateway (440) configured to discover, pair with, and ingest raw radio profile payloads from any BPSD conforming to a standardized wireless biometric profile. Suitable standardized profiles include, without limitation, the Bluetooth Heart Rate Profile, the Bluetooth Pulse Oximeter Service, the Bluetooth Environmental Sensing Service, and functionally equivalent profiles defined under the Bluetooth Special Interest Group specifications or under other wireless-standards-body specifications.

In the Hardware-Agnostic Operating Mode, the embedded wireless radio gateway (440) terminates the wireless session at the CHV, demodulates and decodes the raw profile payloads, and delivers the demodulated payloads to the ECP (420) for processing by the sensor fusion and AI inference engine. The Middleware Layer (300) hosted on the occupant's smartphone is optional or absent in this operating mode. Transition between the first operating mode (Middleware Layer hosted) and the Hardware-Agnostic Operating Mode is effected via over-the-air firmware update to the ECP (420), without modification to the actuator ensemble (430) or the ambient sensor array (410).

XIV. FLEET-LEVEL LEARNING AND CAUSAL DATASET ARCHITECTURE

Referring to FIG. 10, the FML (500) maintains data communication with a plurality of CHV instances (400a, 400b, ..., 400n) via the Internet or another wide-area network. The FML performs: (a) aggregation of Labeled Feedback Tuples transmitted from each CHV instance at a configured synchronization interval; (b) refinement of the shared inference model using the aggregated corpus; and (c) redistribution of refined model parameters to the CHV fleet via over-the-air firmware update, preferably scheduled during periods of low occupancy to avoid interruption of actuator control functions.

A. Two-Tier Longitudinal Personalization Engine

The FML instantiates the Longitudinal Personalization Engine (LPE) comprising: (a) a population-level prior inference model trained on anonymized aggregated Labeled Feedback Tuples across the fleet; and (b) a per-occupant individual-level posterior model indexed to an anonymized occupant identifier. The posterior model encodes occupant-specific offsets, residuals, or adapter weights learned from the Labeled Feedback Tuples associated with that occupant. Upon assignment of an occupant to a CHV instance — including without limitation an instance different from any prior CHV instance occupied by said occupant, and at a deployment site geographically distinct from any prior

deployment site — the FML transmits the occupant’s personalization state to the ECP of the CHV instance, such that personalization persists across instances and sites.

B. Causal Dataset Construction

Each control decision emitted by the ECP constitutes a controlled experimental intervention, the effect of which upon the occupant’s subsequent biometric-response trajectory is captured in the corresponding Labeled Feedback Tuple. Aggregated across the fleet, the Labeled Feedback Tuples constitute a ground-truth causal dataset linking environmental independent variables to physiological recovery outcome dependent variables at longitudinal population scale. This dataset architecture is structurally inaccessible to (i) observational wearable-only systems, which lack environmental-variable control; (ii) observational clinical-registry systems, which lack both environmental control and longitudinal continuity; and (iii) uncoordinated smart-home deployments, which lack biometric ground truth, cross-site statistical power, or both.

C. Continuous-Experimentation Protocol

In a further embodiment, the FML executes a continuous-experimentation protocol across the fleet. At a controlled fraction of occupant-intervals selected pursuant to an experimental design procedure, the ECP is instructed to emit a perturbed control decision deviating from the model-recommended decision along a specified dimension and within a tolerance range preserving occupant safety and comfort. The ambient-response trajectory and biometric-response trajectory associated with the perturbed decision are recorded and contributed to the Labeled Feedback Tuple. Over a sufficient cardinality of occupant-intervals, the aggregated perturbation data yield causal estimates of the effect of each environmental dimension on physiological recovery outcomes, rather than correlational estimates. Perturbations are at all times constrained by the SICRH.

D. Federated and Privacy-Preserving Learning Options

The FML may, in further embodiments, implement federated learning protocols in which raw Labeled Feedback Tuples are retained at the respective CHV edge, and only model-parameter gradients or functionally equivalent aggregated summary statistics are transmitted to the FML. Differential privacy mechanisms, k-anonymity mechanisms, secure multi-party computation mechanisms, or functionally equivalent privacy-preserving techniques may be applied to the aggregated statistics prior to cross-fleet release.

XV. BEHAVIORAL-STATE INPUT LAYER

A companion application executing on the occupant's smartphone, on a kiosk device within a facility housing the CHV, or on a web interface, captures occupant interactions with a broader ecosystem of which the CHV forms a part. Captured interactions include, without limitation:

- Meal ordering, nutritional macronutrient composition, and meal timing;
- Caffeine beverage orders, optionally gated by the application such that access to a caffeine dispenser requires application-mediated authentication;
- Scheduled and completed utilization of recovery-equipment modalities including without limitation sauna, cold plunge, infrared exposure, and compression therapy;
- Exercise sessions, including intensity, duration, and modality, captured by the application or synchronized from external fitness-tracking sources;
- Attendance at social events hosted within the facility;
- Selection of physical workspaces within the facility, including without limitation focused-work zones, collaborative zones, quiet zones, and outdoor zones; and
- Time-of-day and temporal-pattern features derived from the foregoing.

The behavioral-state input layer derives, at each control interval, a behavioral-state vector summarizing the foregoing features over a preceding time window. Said vector is provided as an input to the ECP inference engine pursuant to Section X, enabling the system to modulate the CHV environment in anticipation of predictable biometric consequences of recent behavioral events.

XVI. ALTERNATIVE EMBODIMENTS AND EQUIVALENTS

A. Non-Contact BPSD Embodiment

Referring to FIG. 12, alternative embodiments substitute or supplement the worn BPSD with one or more non-contact occupant-sensing modalities, including without limitation: ballistocardiographic sensors embedded in a sleep surface within the CHV; millimeter-wave radar sensors; ultra-wideband radar sensors; thermal-infrared imaging sensors; depth-camera sensors; and ambient-acoustic sensors operable to derive respiration rate and heart rate from low-amplitude micro-motion or acoustic cues. The inference engine accepts biometric-proxy inputs from any of the foregoing modalities in substitution for, or in combination with, the worn BPSD.

B. Facility-Level Deployment and Multi-Unit Coordination

In a preferred facility-level deployment, a plurality of CHV instances are installed within a single structure, sharing building-level utilities. In said deployment, the FML additionally performs building-level coordination including, without limitation: arbitration of aggregate make-up-air demand

such that simultaneous high-demand ventilation requests do not exceed building HVAC capacity; anti-phase coordination of thermal-actuator cycling across adjacent CHV instances to reduce inter-unit thermal interference; aggregate quiet-hours enforcement; and predictive pre-conditioning of individual CHV instances based on occupant-reported schedules.

C. BPSD Asset-Rotation and Sanitation Protocol

Referring to FIG. 11, in a further embodiment the BPSD is assigned at the level of the CHV rather than to an individual occupant in perpetuity. Upon transition from a first occupant to a second occupant within the same CHV, a sanitation protocol is executed, comprising: removal of the BPSD from the first occupant; separation of the housing from the band or strap; laundering of the band or strap at a temperature at or above sixty degrees Celsius; sanitization of the housing by application of an antimicrobial agent, preferably an isopropyl-alcohol solution of seventy percent or greater concentration; visual inspection for wear and fastener degradation; retirement of any BPSD exhibiting such condition and replacement from an inventory reserve maintained at a configured surplus fraction; and reassignment to the next occupant of the CHV. By such assignment, longitudinal environmental and biometric data indexed to the CHV accumulate across successive occupants, supporting unit-level calibration and diagnostics independent of occupant turnover.

D. Dual Regulatory-Posture Architecture

In a further embodiment, the BPSD and its signal-processing component are architected to support two distinct regulatory postures, namely (i) a general-wellness posture, in which the device outputs are presented as wellness estimates and the device is not marketed for diagnostic or therapeutic purposes; and (ii) a medical-device posture, in which the same underlying hardware, supplemented by firmware validation, documentation, and validation data collected pursuant to Section XII.A, supports submission under applicable medical-device regulatory pathways. The dual-posture architecture requires no hardware modification between postures; substitution is accomplished through firmware configuration and regulatory documentation supplementation.

E. Data Privacy and Consent Framework

In a further embodiment, biometric data collected from the BPSD is classified under a heightened data-protection standard. On-wearable and on-middleware preprocessing is structured to minimize transmission of raw physiological waveforms where processed features are sufficient for downstream control. The FML implements per-data-type granular opt-in consent, segregates individual-level data from anonymized aggregate analytics, supports occupant-initiated data export and deletion, and applies differential-privacy, k-anonymity, secure aggregation, or functionally equivalent techniques to any aggregate analytics released to building operators, franchise operators, or research partners.

F. Doctrine of Equivalents

Without limitation of the foregoing, and in accordance with the doctrine of equivalents, the present invention shall be understood to encompass all functionally equivalent apparatus, methods, and systems that perform substantially the same function in substantially the same way to achieve substantially the same result as the preferred embodiments described herein, whether or not such apparatus, methods, or systems are specifically enumerated in this Detailed Description.

XVII. ILLUSTRATIVE EXAMPLES

The following examples are provided by way of illustration of the preferred embodiments and are not to be construed as limitations upon the scope of the invention or of the appended claims.

Example 1. Sleep-Onset Environmental Preparation

An occupant, wearing the BPSD, enters the CHV at approximately 22:00 local time. The companion application indicates an intended sleep-onset time of approximately 22:30. Upon occupant entry, the ECP receives from the BPSD a measured heart rate consistent with an elevated resting level and a short-window HRV value consistent with sympathetic tone. The ECP receives from the ambient sensor array an interior temperature exceeding the predicted ASTCM setpoint. The ECP commands the thermal subsystem to produce a head-cooled, foot-warmed gradient consistent with sleep-onset physiology, and commands the photic subsystem to initiate a gradual dim-down and spectral-shift-toward-warm curve. Subsequently, the BPSD reports a decrease in heart rate, an increase in HRV, and, shortly thereafter, an inferred sleep-stage transition from wake to light sleep.

Example 2. Safety-Preemption Override During Sleep

During an occupant sleep interval, the ambient carbon-dioxide sensor registers a sustained rise above the configured ceiling, concurrent with a measured decline in the occupant's blood-oxygen saturation toward the configured floor. The VDM, having been commanded at a low acoustic-minimizing flow rate prior to said events, detects both safety-threshold crossings. Pursuant to the SICRH, the ECP issues an override command to the filtered ventilation subsystem, elevating fan flow rate notwithstanding the acoustic-quietness objective. Ambient carbon-dioxide and occupant blood-oxygen saturation return to within tolerance; the ECP subsequently reduces commanded flow rate consistent with restored safety margins.

Example 3. Substitute Inference Model Deployment

Upon accumulation of a paired Stream A / Stream B corpus of sufficient cardinality and diversity across a plurality of CHV instances and occupants, a one-dimensional convolutional neural network is trained on said corpus to reproduce the HRV output channel of Stream B from the raw

photoplethysmographic and accelerometer waveforms of Stream A. The trained model is validated against an independent held-out validation subset and against reference electrocardiographic measurements. The model is subsequently compressed and deployed to BPSD firmware via over-the-air update in Schema-Preserving Substitution for the incumbent signal-processing component. Downstream consumers of Stream B, including the ECP inference engine and the FML, operate without modification.

Example 4. Cross-Site Personalization Transfer

An occupant occupies a first CHV instance at a first deployment site for a continuous plurality of occupant-nights. Over said interval, the LPE posterior for the occupant learns idiosyncratic offsets to the population-level ASTCM prior, consistent with the occupant's particular thermal-response characteristics. The occupant subsequently arrives at a second deployment site and is assigned to a second CHV instance. The FML transmits the occupant's LPE posterior to the ECP of the second CHV instance, which commences ASTCM computation using the occupant-specific posterior from the first occupant-night at the second site, rather than from the population prior.

XVIII. CLAIMS

Note: Claims are not required in a Provisional Application under 35 U.S.C. § 111(b); the following claims are nevertheless included pursuant to best practice, in order to establish written-description support for the scope of subject matter intended to be pursued in a subsequent non-provisional application. The Applicant reserves the right to pursue, in any subsequent non-provisional, continuation, continuation-in-part, or divisional application, claims of broader, narrower, or otherwise different scope from those set forth below. Hardware, component, and material references appearing in the claims shall be construed as designating preferred embodiments and shall not be construed to limit the claim scope pursuant to the definitions set forth in Section VII supra.

What is claimed is:

1.A closed-loop biometric-driven environmental control system, comprising:

- (a) a plurality of Controlled Habitation Volumes (CHVs) geographically distributed across one or more deployment sites, each CHV comprising an interior spatial volume sized to accommodate at minimum one occupant; an ambient sensor array including at minimum a carbon-dioxide sensor, a volatile-organic-compound sensor, an ambient-temperature sensor, and a relative-humidity sensor; an actuator ensemble including a thermal subsystem comprising a plurality of independently addressable solid-state thermal elements, a filtered ventilation subsystem, and a tunable-spectrum photic subsystem; and an Edge Compute Processor (ECP) physically co-located

within or with the CHV, in communication with the ambient sensor array and the actuator ensemble, and storing an inference model in a memory thereof;

- (b)a Body-Proximate Sensing Device (BPSD) configured to be worn by or positioned proximate to an occupant, the BPSD including at minimum a photoplethysmographic sensor, an inertial sensor, and a wireless radio, and producing a continuous biometric data stream;
- (c)a Middleware Layer configured to receive said biometric data stream, normalize said stream into a structured schema, and deliver said normalized stream to the ECP; and
- (d)a Fleet Management Layer (FML) in data communication with the plurality of CHVs and configured to (i) aggregate Labeled Feedback Tuples from said plurality, each said tuple comprising at minimum a fused input vector derived from said normalized biometric data stream and said ambient sensor array, a control decision emitted by said inference model, and a downstream physiological response recorded subsequent to said control decision; (ii) refine said inference model using said aggregated Labeled Feedback Tuples; and (iii) redistribute refined model parameters to said plurality of CHVs;

wherein, at each control interval, the ECP of a CHV ingests said normalized biometric data stream and said ambient sensor array data, applies said inference model to a fused input vector derived therefrom, and emits a control decision to said actuator ensemble, said control decision comprising at minimum a per-zone thermal setpoint vector, an airflow command to said filtered ventilation subsystem, and a correlated-color-temperature and intensity command to said tunable-spectrum photic subsystem.

2.The system of claim 1, wherein said plurality of independently addressable solid-state thermal elements is disposed within the interior volume of each CHV such that a first subset is positioned proximal to a head region of the occupant in a recumbent position, a second subset is positioned proximal to a torso region, and a third subset is positioned proximal to a foot region; and wherein said ECP independently commands each subset to produce a shaped spatial thermal gradient across the occupant body, said shaped gradient being determined at least in part by a body position inferred from said inertial sensor of said BPSD and by an autonomic state inferred from heart-rate-variability features derived from said photoplethysmographic sensor.

3.The system of claim 1, wherein said inference model applies a Sensor-Input Conflict Resolution Hierarchy (SICRH) comprising, in descending order of precedence: (i) a blood-oxygen-saturation safety floor; (ii) an ambient carbon-dioxide ceiling; (iii) an ambient volatile-organic-compound ceiling; (iv) a thermal comfort objective; (v) an acoustic-quietness objective; and (vi) an energy-efficiency objective; and wherein, upon detection of violation or trending violation of a higher-priority parameter, the ECP overrides lower-priority objectives, including without limitation the acoustic-quietness objective during an inferred sleep interval, to restore the higher-priority parameter to within tolerance.

4. The system of claim 1, wherein said BPSD further comprises a biometric signal-processing component operable to convert raw sensor waveforms into processed biometric outputs conforming to said structured schema; and wherein said Middleware Layer and said FML are jointly configured to log, in parallel, said raw sensor waveforms as Stream A and said processed biometric outputs as Stream B into a dual-logged paired corpus indexed by occupant-night, such that said dual-logged corpus constitutes a labeled training set wherein Stream B outputs serve as ground-truth labels for corresponding Stream A inputs.
5. The system of claim 4, further comprising a substitute on-device inference model trained on said dual-logged corpus, said substitute model being configured to produce processed biometric outputs conforming to said structured schema and being deployable to said BPSD in Schema-Preserving Substitution for said biometric signal-processing component without modification to downstream consumers of said structured schema, said downstream consumers including said ECP and said FML.
6. The system of claim 1, wherein each CHV further comprises an embedded wireless radio gateway operable, in a Hardware-Agnostic Operating Mode, to discover, pair with, and ingest raw radio profile payloads from any BPSD conforming to a standardized wireless biometric profile, and to deliver said ingested payloads to said ECP, thereby permitting operation of said system without dependence upon any specific BPSD vendor or any specific biometric signal-processing component thereof.
7. The system of claim 1, further comprising a companion application executing on a compute device associated with the occupant, said companion application configured to capture behavioral-state events including at least one of a nutritional-intake event, a caffeine-intake event, a recovery-equipment-utilization event, an exercise event, a social-activity event, and a spatial-utilization event, and to deliver a behavioral-state vector summarizing said events to said ECP as an input to said inference model.
8. The system of claim 1, wherein said FML instantiates a Longitudinal Personalization Engine (LPE) comprising (i) a population-level prior inference model trained on said aggregated Labeled Feedback Tuples across the fleet, and (ii) a per-occupant individual-level posterior model indexed to an anonymized occupant identifier and updated from Labeled Feedback Tuples associated with said occupant; and wherein, upon assignment of said occupant to a CHV instance different from any prior CHV instance occupied by said occupant, said FML transmits said occupant's posterior to the ECP of said different CHV instance, such that personalization state persists across CHV instances and across deployment sites.
9. The system of claim 1, wherein said FML executes a continuous-experimentation protocol comprising, at a controlled fraction of occupant-intervals across the fleet, commanding the ECP to emit a perturbed control decision deviating from a model-recommended decision along a specified

dimension and within a safety-preserving tolerance range; recording ambient-response and biometric-response trajectories associated with said perturbed decision; and incorporating said trajectories into Labeled Feedback Tuples transmitted to said FML, whereby the aggregated perturbation data yield causal estimates of environmental effects on physiological recovery outcomes.

10. The system of claim 1, wherein said plurality of independently addressable solid-state thermal elements comprises a plurality of thermoelectric devices employing the Peltier effect.
11. The system of claim 1, wherein said ECP comprises an ARM Cortex-class microprocessor with an integrated floating-point unit, and wherein said inference model is stored in said memory as a quantized neural-network representation.
12. The system of claim 1, wherein said filtered ventilation subsystem comprises a variable-speed fan and a high-efficiency particulate-air filter of H13 rating or higher, and wherein the airflow command issued thereto is determined at least in part as a function of measured ambient carbon-dioxide concentration, measured ambient volatile-organic-compound concentration, and the occupant's measured respiration rate and blood-oxygen saturation.
13. The system of claim 1, wherein said tunable-spectrum photic subsystem is operable to vary correlated color temperature continuously across a range spanning at minimum from 2700 kelvin to 6500 kelvin, and wherein said ECP commands a sleep-onset dim-down and spectral-warming curve prior to an inferred sleep-onset time, and a wake-optimization brightening and spectral-cooling curve timed to coincide with an inferred transition from a light sleep stage toward a wake state.
14. The system of claim 1, wherein said BPSD is assigned to said CHV rather than to any individual occupant; and wherein upon transition from a first occupant to a second occupant of said CHV, said BPSD undergoes a sanitation cycle comprising laundering of a band or strap portion at a temperature at or above sixty degrees Celsius and sanitization of a housing portion by application of an antimicrobial agent, prior to assignment to said second occupant.
15. A method for closed-loop biometric-driven environmental control in a Controlled Habitation Volume (CHV), said method comprising:
 - (a) receiving, at an Edge Compute Processor (ECP) physically co-located within or with the CHV, a normalized biometric data stream from a Body-Proximate Sensing Device (BPSD) worn by or positioned proximate to an occupant of the CHV, said normalized biometric data stream having been produced by a Middleware Layer from a raw biometric data stream emitted by said BPSD;
 - (b) receiving, at said ECP, ambient sensor data from an ambient sensor array disposed within the CHV;

- (c) forming a fused input vector from at minimum said normalized biometric data stream and said ambient sensor data;
 - (d) applying an inference model, stored in a memory of said ECP, to said fused input vector to produce a control decision;
 - (e) applying a Sensor-Input Conflict Resolution Hierarchy (SICRH) to said control decision, said SICRH enforcing in descending order of precedence at minimum a blood-oxygen-saturation safety floor, an ambient carbon-dioxide ceiling, a thermal comfort objective, an acoustic-quietness objective, and an energy-efficiency objective;
 - (f) emitting said control decision, as modified by said SICRH, to an actuator ensemble of said CHV comprising a plurality of independently addressable solid-state thermal elements, a filtered ventilation subsystem, and a tunable-spectrum photic subsystem;
 - (g) logging a Labeled Feedback Tuple comprising at minimum said fused input vector, said control decision, any override applied by said SICRH, and a downstream physiological response recorded subsequent thereto, to a persistent local store; and
 - (h) transmitting said Labeled Feedback Tuple to a Fleet Management Layer (FML) in communication with a plurality of CHVs, said FML configured to refine said inference model using Labeled Feedback Tuples aggregated across said plurality and to redistribute refined model parameters to said plurality.
- 16.** The method of claim 15, further comprising: logging, in parallel with said normalized biometric data stream, raw sensor waveforms from said BPSD as Stream A and processed biometric outputs produced by a biometric signal-processing component of said BPSD as Stream B; training a substitute on-device inference model on said paired Stream A and Stream B corpus, said substitute model being configured to produce processed biometric outputs conforming to the structured schema of Stream B; and deploying said substitute on-device inference model to said BPSD in Schema-Preserving Substitution for said biometric signal-processing component, without modification to said ECP or said FML.
- 17.** The method of claim 15, further comprising: commanding a first subset of said plurality of independently addressable solid-state thermal elements proximal to a head region of said occupant, and a second subset proximal to a foot region of said occupant, to produce a shaped spatial thermal gradient in which said head region is maintained at a cooler setpoint than said foot region during an interval surrounding an inferred sleep-onset time.
- 18.** The method of claim 15, further comprising: at a controlled fraction of said control intervals, emitting a perturbed control decision deviating from a model-recommended decision along a specified dimension within a safety-preserving tolerance range; recording an ambient-response trajectory and a

biometric-response trajectory associated with said perturbed control decision; and incorporating said recorded trajectories into said Labeled Feedback Tuple transmitted to said FML, whereby causal estimates of environmental effects on physiological recovery outcomes are derivable from the aggregated data.

19.A Controlled Habitation Volume (CHV) configured to operate as a node within a learning fleet of CHVs, said CHV comprising:

- (a) an interior spatial volume sized to accommodate at minimum one occupant;
- (b) an ambient sensor array comprising at minimum a carbon-dioxide sensor, a volatile-organic-compound sensor, an ambient-temperature sensor, and a relative-humidity sensor;
- (c) a thermal subsystem comprising a plurality of independently addressable solid-state thermal elements distributed spatially within the interior volume such that a first subset is positioned proximal to a head region of an occupant in a recumbent position and a second subset is positioned proximal to a foot region;
- (d) a filtered ventilation subsystem comprising a variable-speed fan and a filter;
- (e) a tunable-spectrum photic subsystem operable to vary correlated color temperature across a range of at minimum from 2700 kelvin to 6500 kelvin;
- (f) an Edge Compute Processor (ECP) physically co-located within or with said CHV, configured to receive a normalized biometric data stream from a Body-Proximate Sensing Device (BPSD) worn by or positioned proximate to the occupant, receive ambient sensor data from said ambient sensor array, apply an inference model to a fused input vector derived therefrom subject to a Sensor-Input Conflict Resolution Hierarchy, and emit a control decision to said thermal subsystem, said filtered ventilation subsystem, and said tunable-spectrum photic subsystem; and
- (g) a data interface configured to transmit Labeled Feedback Tuples to a Fleet Management Layer (FML) and to receive refined model parameters therefrom.

20. The CHV of claim 19, further comprising an embedded wireless radio gateway operable to ingest raw radio profile payloads from any BPSD conforming to a standardized wireless biometric profile, such that said CHV is operable in a Hardware-Agnostic Operating Mode independent of any specific BPSD vendor.

XIX. ABSTRACT OF THE DISCLOSURE

A closed-loop system and method for biometric-driven environmental control in a Controlled Habitation Volume (CHV) operating as a node within a learning fleet of CHVs distributed across a plurality of

deployment sites. Each CHV comprises an ambient sensor array, a plurality of independently addressable solid-state thermal elements, a filtered ventilation subsystem, a tunable-spectrum photic subsystem, and an Edge Compute Processor (ECP) executing sensor fusion and artificial intelligence inference. A Body-Proximate Sensing Device (BPSD) worn by or positioned proximate to the occupant delivers a biometric data stream, normalized by a Middleware Layer, to the ECP. An inference engine fuses biometric, ambient, and behavioral state vectors and emits actuator control decisions subject to a Sensor-Input Conflict Resolution Hierarchy (SICRH) enforcing blood-oxygen-saturation safety floor and ambient carbon-dioxide ceiling constraints preempting comfort and acoustic-quietness objectives. A Fleet Management Layer aggregates Labeled Feedback Tuples from the CHV fleet, refines a shared inference model via a Longitudinal Personalization Engine (LPE), and redistributes refined parameters across the fleet. A dual-logging architecture captures paired raw sensor waveforms and processed biometric outputs, supporting subsequent training of substitute on-device inference models deployable in Schema-Preserving Substitution for third-party biometric signal-processing components. Further aspects include a Hardware-Agnostic Operating Mode via an embedded wireless radio gateway; a continuous-experimentation protocol yielding causal estimates of environmental effects on physiological recovery outcomes; cross-site persistence of individual-level personalization posteriors; and a CHV-level BPSD asset-rotation and sanitation protocol.

XX. FIGURES

The following figures (FIG. 1 through FIG. 12) illustrate preferred and alternative embodiments of the invention as described in Section VIII supra. The figures are provided at the Provisional Application stage in informal-drawing form; prior to conversion to a non-provisional application, formal patent drawings shall be prepared in substantial conformance with 37 C.F.R. § 1.84, including without limitation the line-weight, margin, and lettering-size requirements set forth therein. The informal drawings included herein are of sufficient particularity to support the written description of the embodiments depicted under 35 U.S.C. § 112(a).

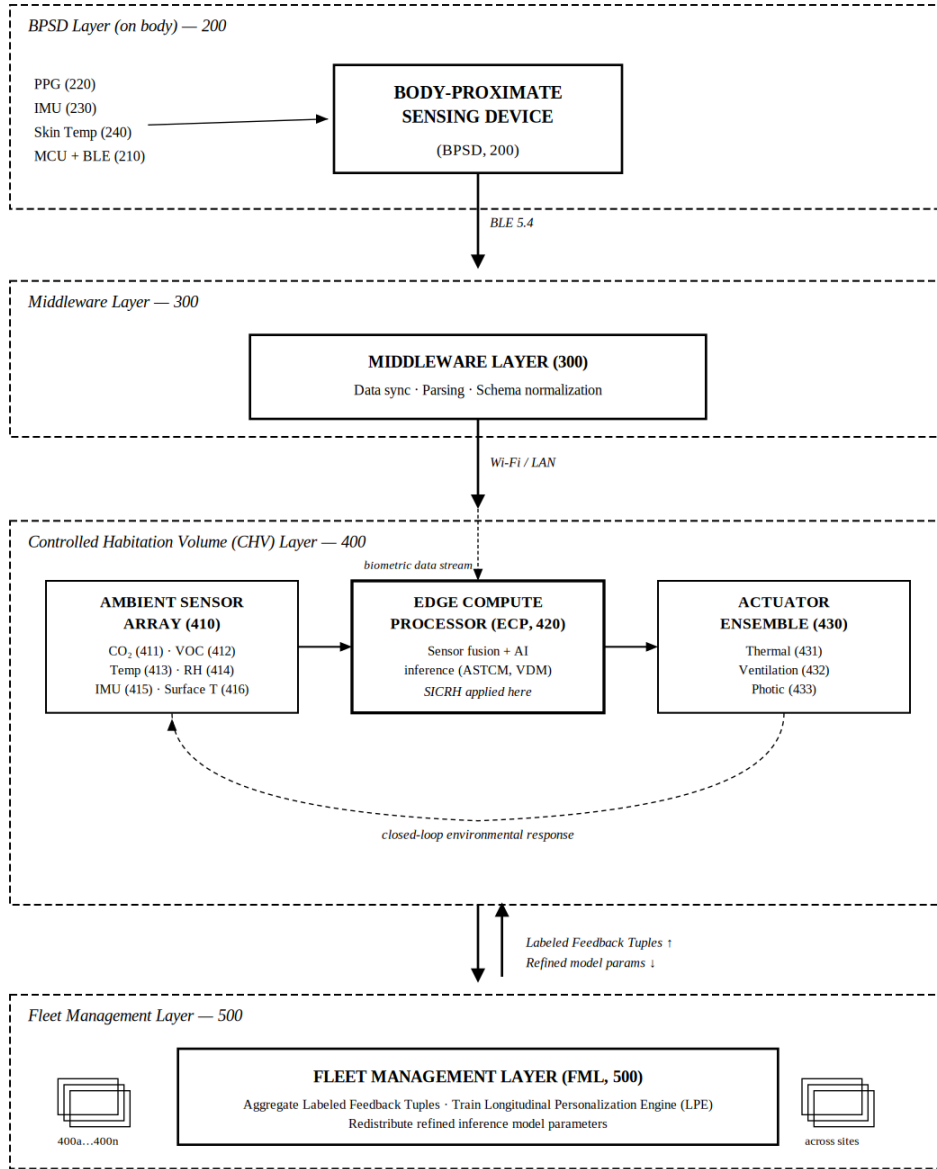


FIG. 1

FIG. 1. Four-layer data pipeline architecture — BPSD, Middleware, CHV, and Fleet Management Layers.

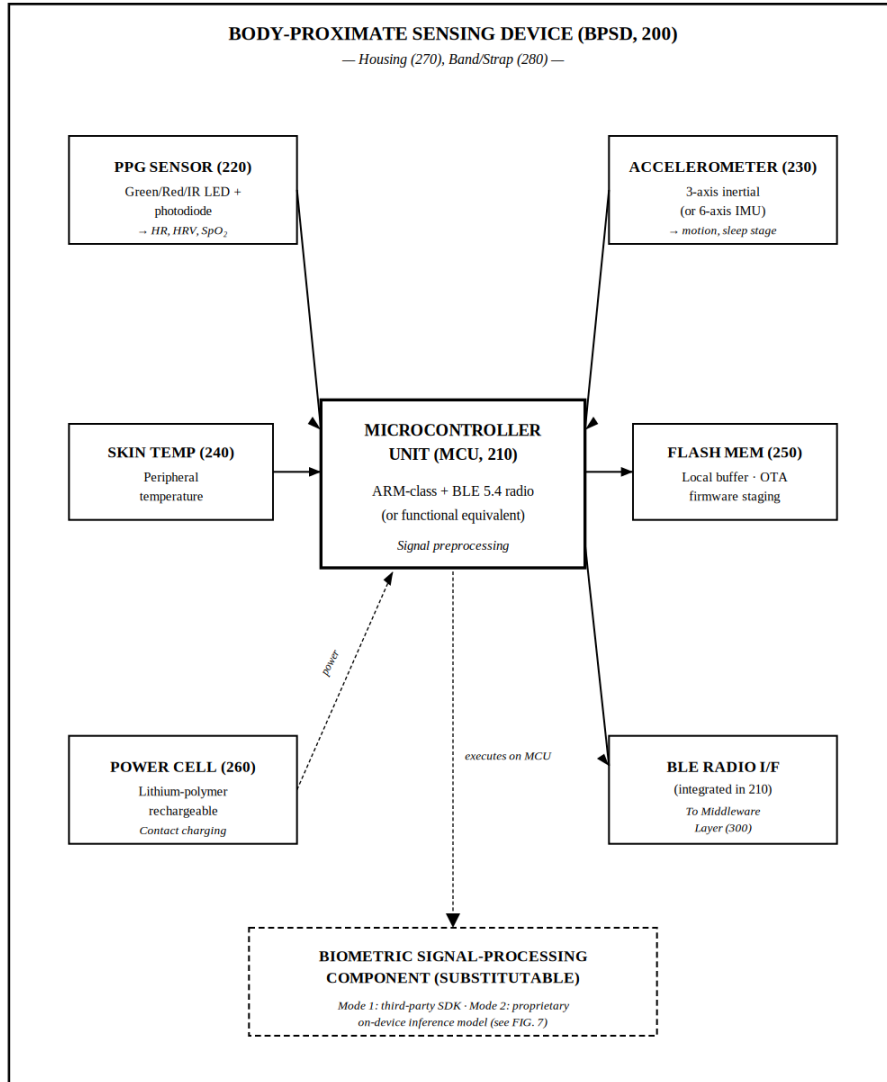


FIG. 2

FIG. 2. *Body-Proximate Sensing Device (BPSD, 200) — functional block diagram of preferred embodiment.*

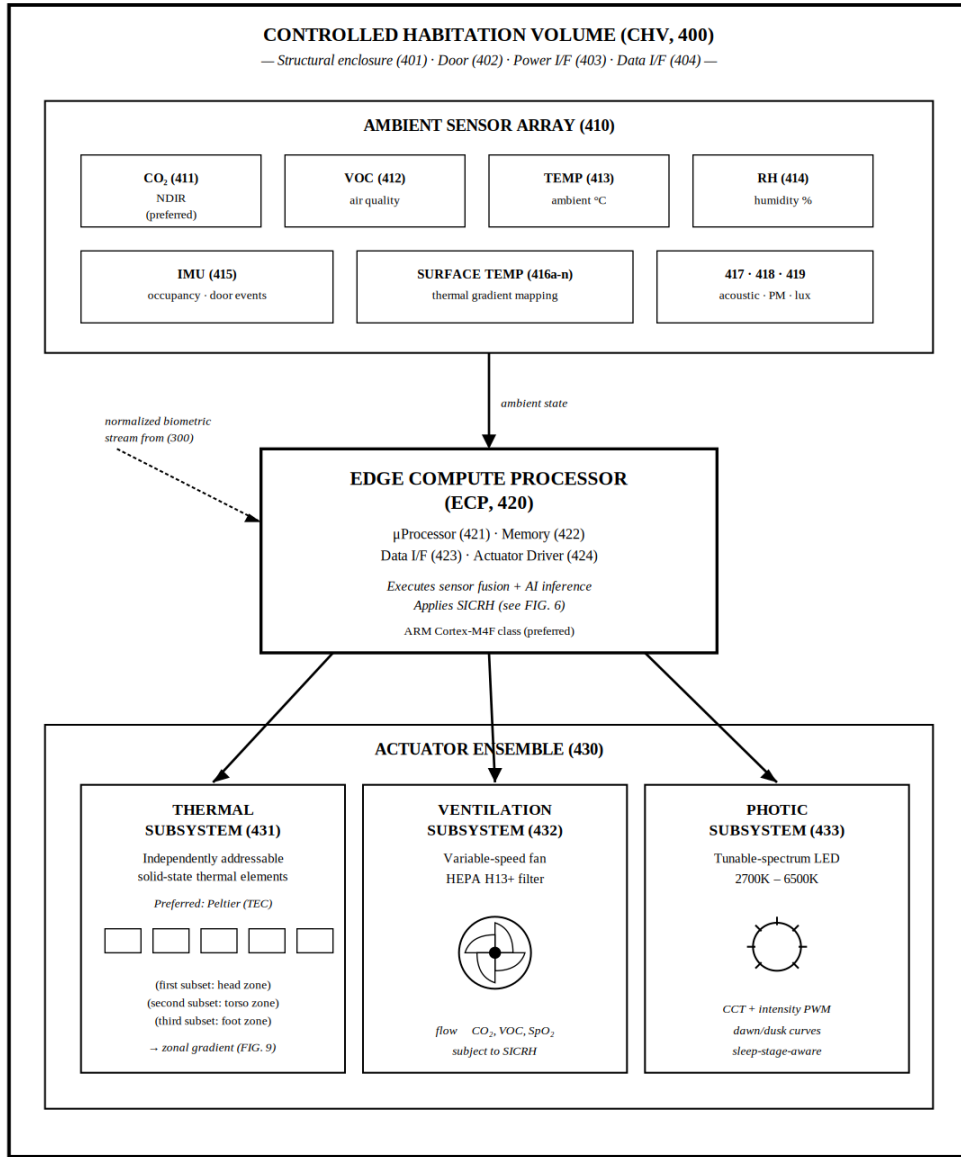


FIG. 3

FIG. 3. *Controlled Habitation Volume (CHV, 400) — functional block diagram of preferred embodiment.*

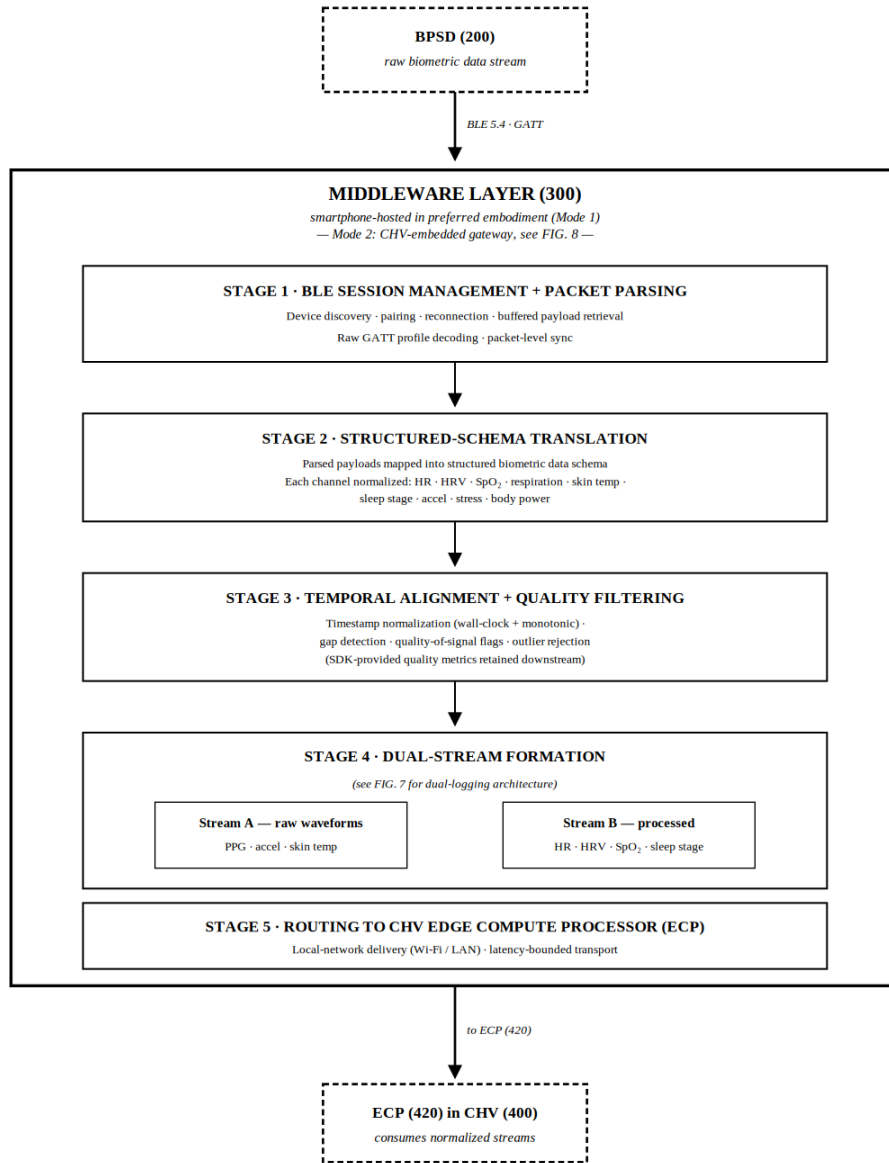


FIG. 4

FIG. 4. *Middleware Layer (300) — five-stage data-flow from raw BPSD transmissions to ECP-consumable streams.*

Sensor Fusion and AI Inference Process executed by ECP (420) at each control interval

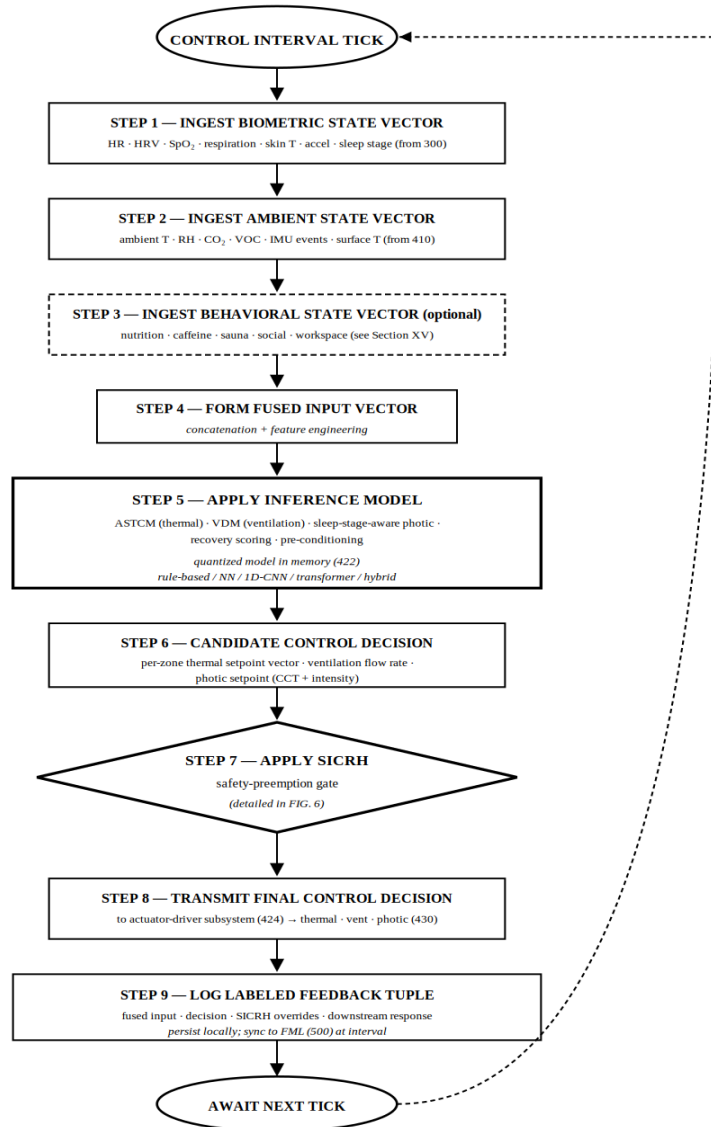


FIG. 5

FIG. 5. Sensor fusion and AI inference process flow — executed by ECP (420) at each control interval.

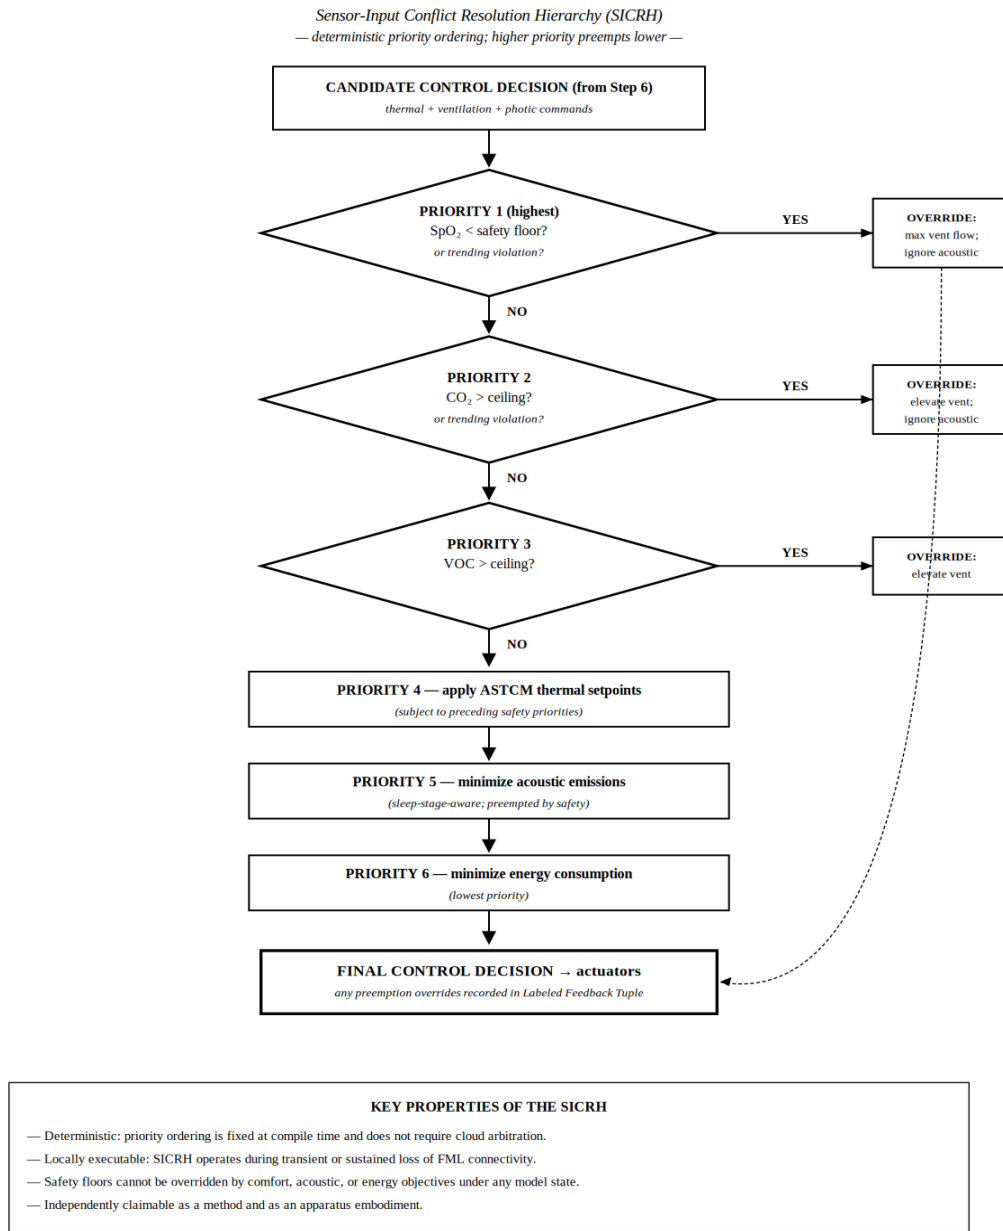


FIG. 6

FIG. 6. *Sensor-Input Conflict Resolution Hierarchy (SICRH) — safety-preemption priority ordering.*

Dual-Logging Architecture enabling Schema-Preserving Substitution

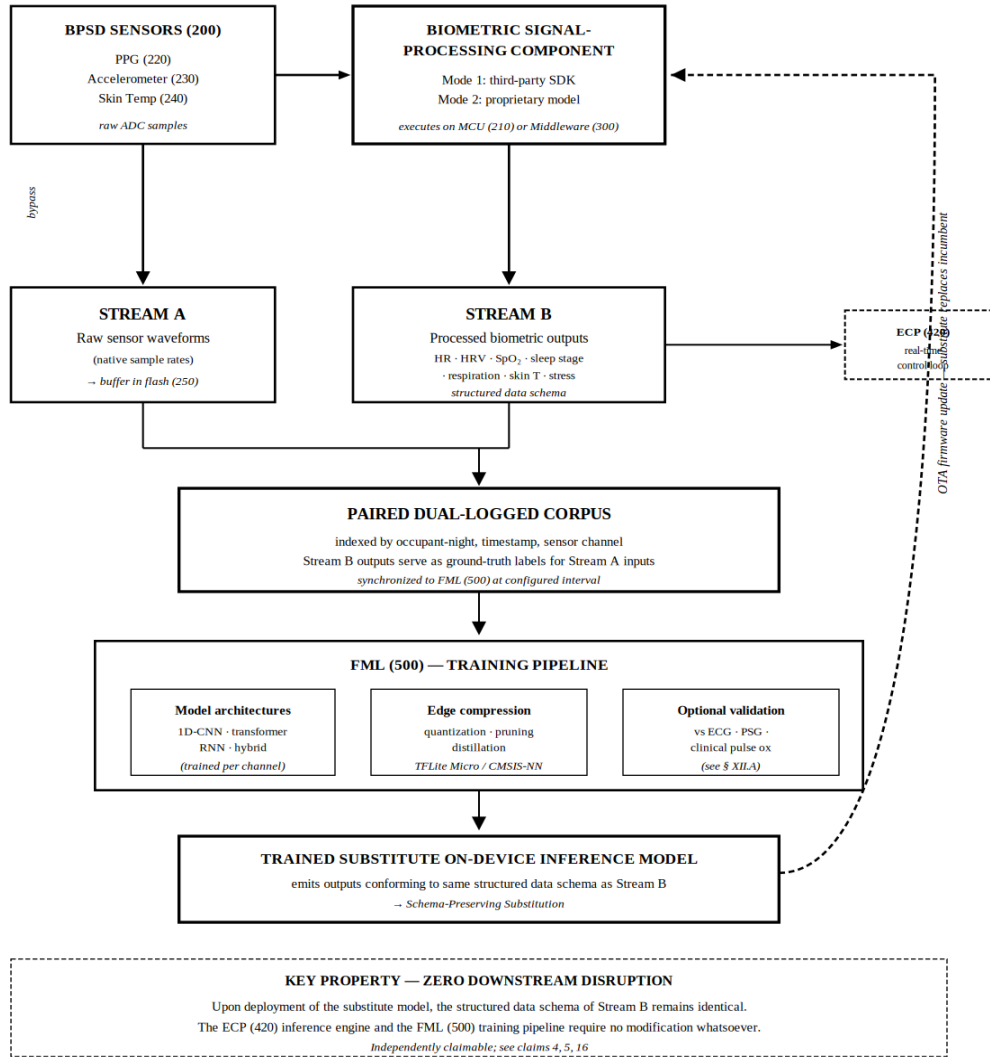


FIG. 7

FIG. 7. Dual-logging architecture — Stream A (raw waveforms) + Stream B (processed outputs) enabling Schema-Preserving Substitution.

Operating-Mode Transition: Middleware-hosted (Mode 1) ↔ Hardware-Agnostic Gateway (Mode 2)

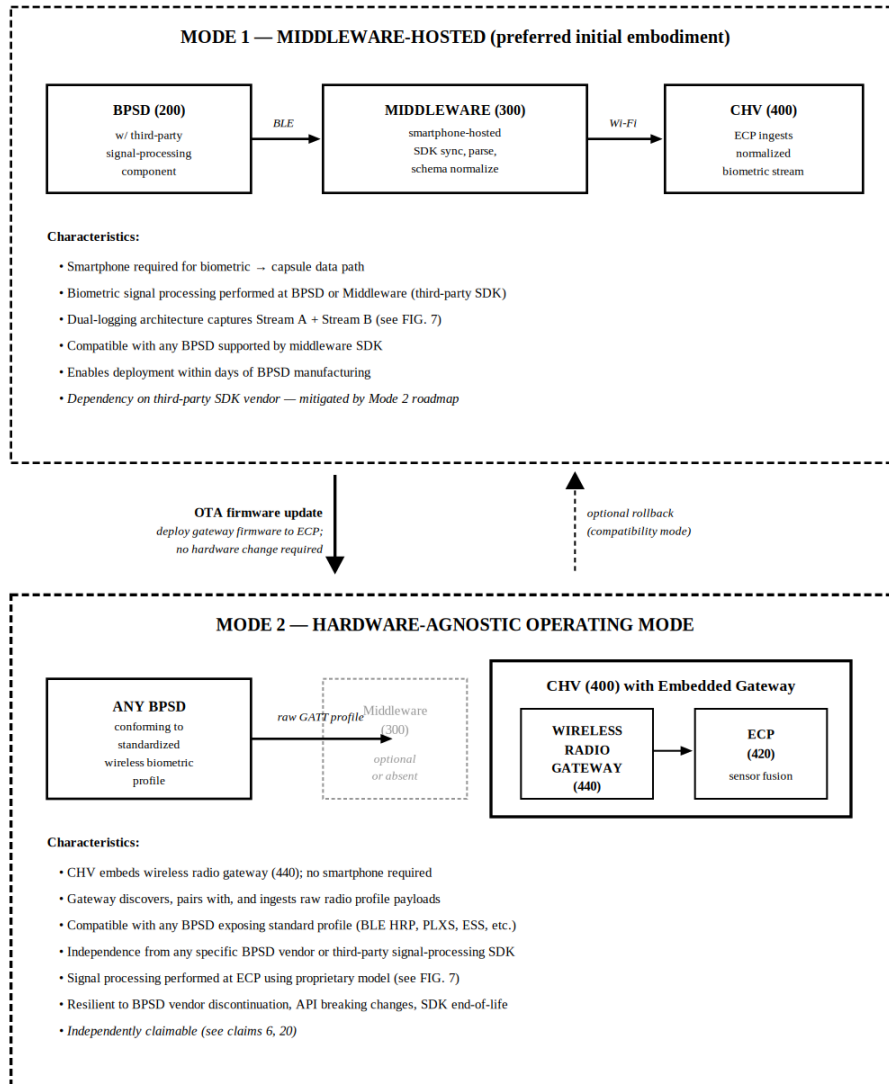


FIG. 8

FIG. 8. Operating-mode state diagram — Middleware-hosted (Mode 1) ↔ Hardware-Agnostic Gateway (Mode 2).

Schematic Cross-Section of CHV Interior — Zonal Thermal Gradient Shaping

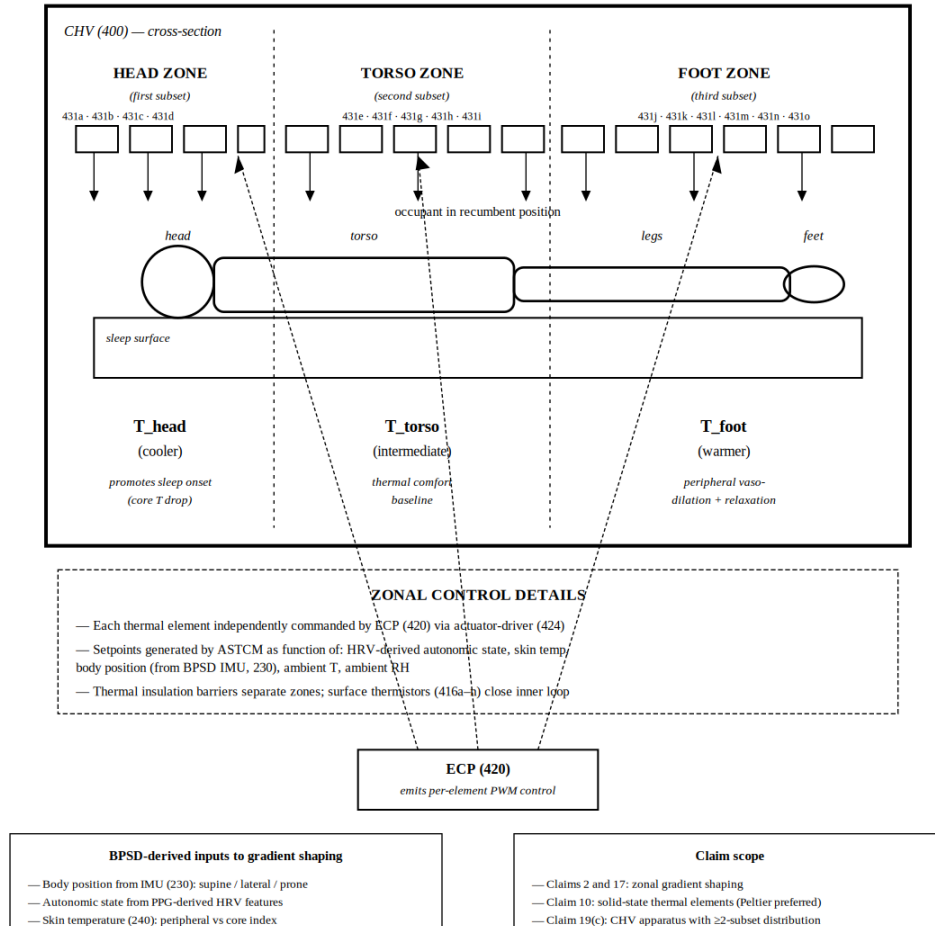


FIG. 9

FIG. 9. CHV interior cross-section — zonal thermal gradient shaping across head, torso, and foot regions.

Fleet Management Layer (FML, 500) — aggregation, training, and redistribution

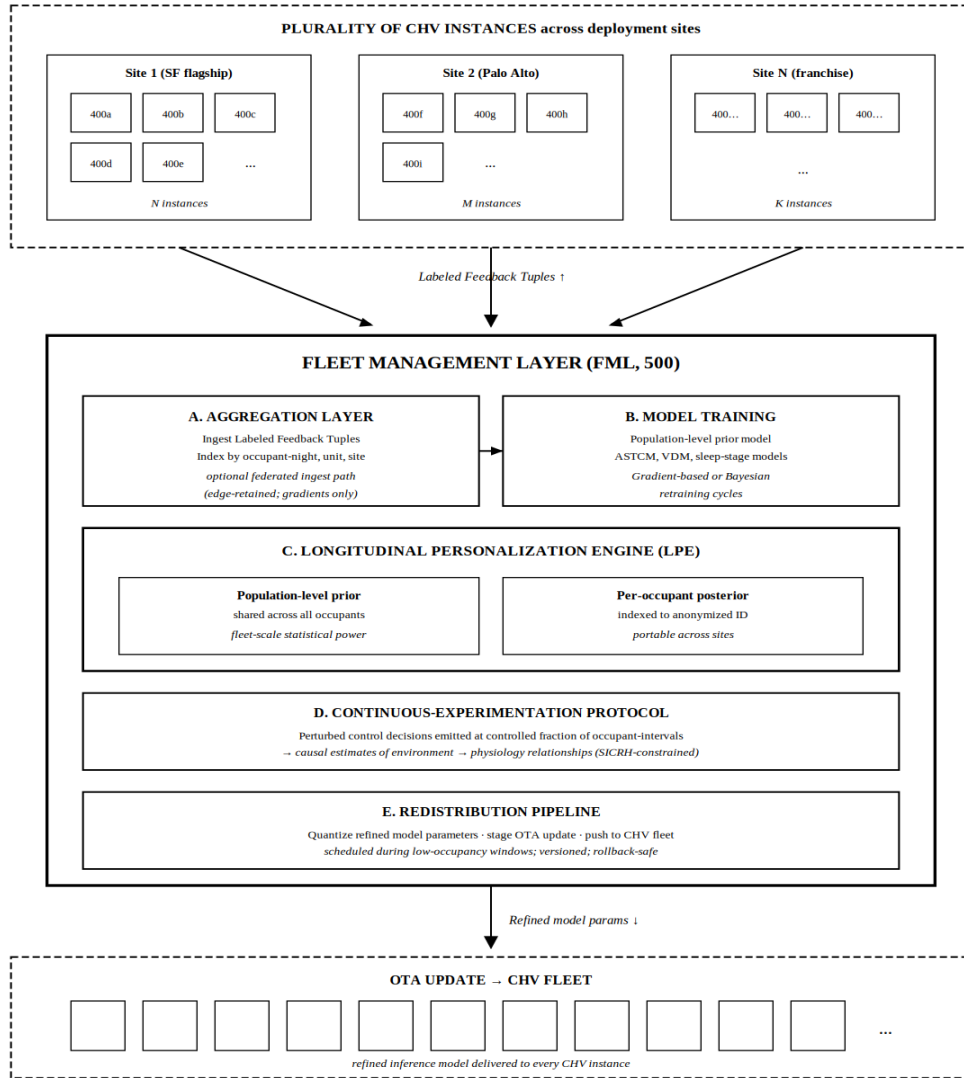


FIG. 10

FIG. 10. Fleet Management Layer (FML, 500) — aggregation, training (including LPE), and redistribution pipeline.

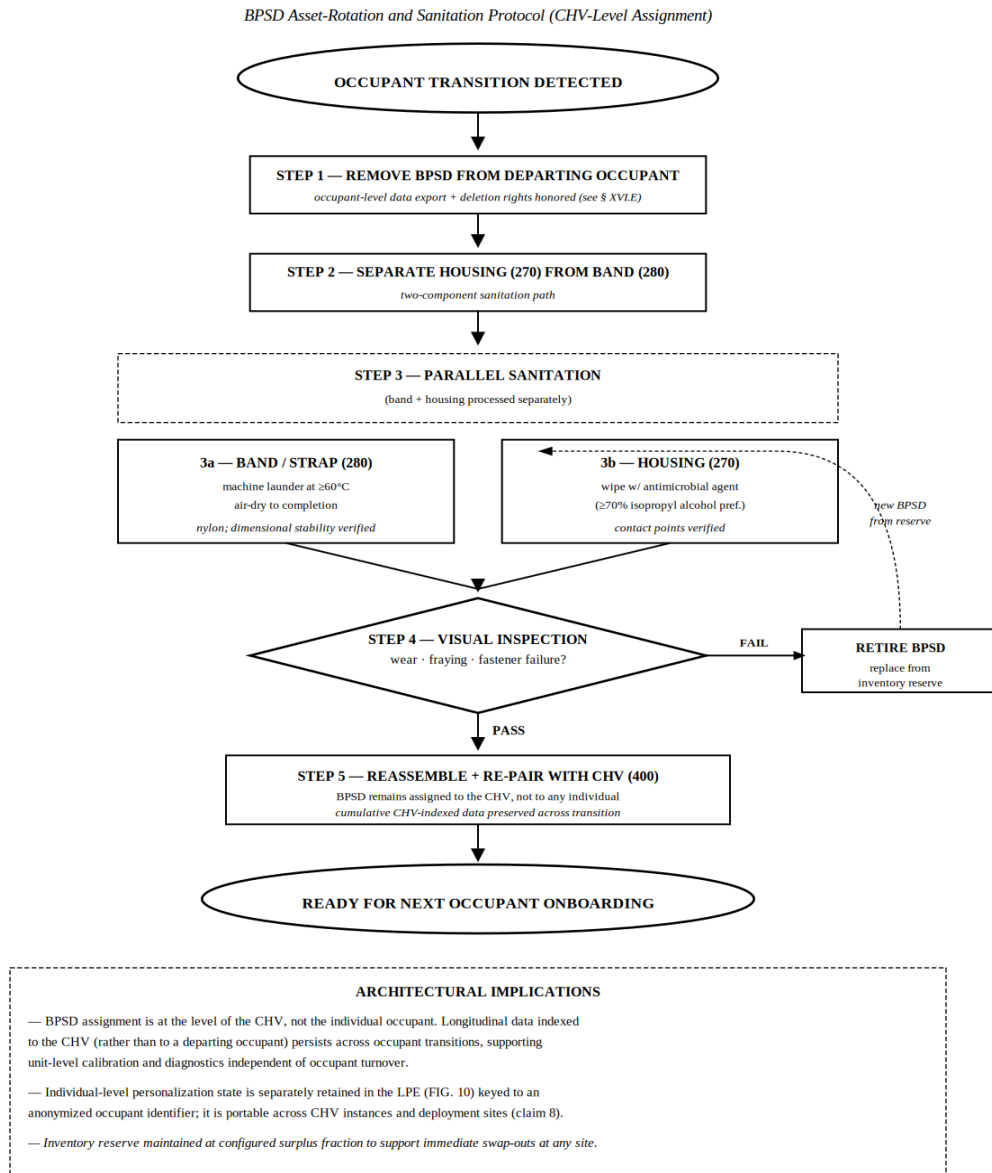


FIG. 11

FIG. 11. *BPSD asset-rotation and sanitation protocol executed upon occupant transition within a CHV.*

Alternative BPSD Embodiment — Non-Contact Occupant Sensing Modalities

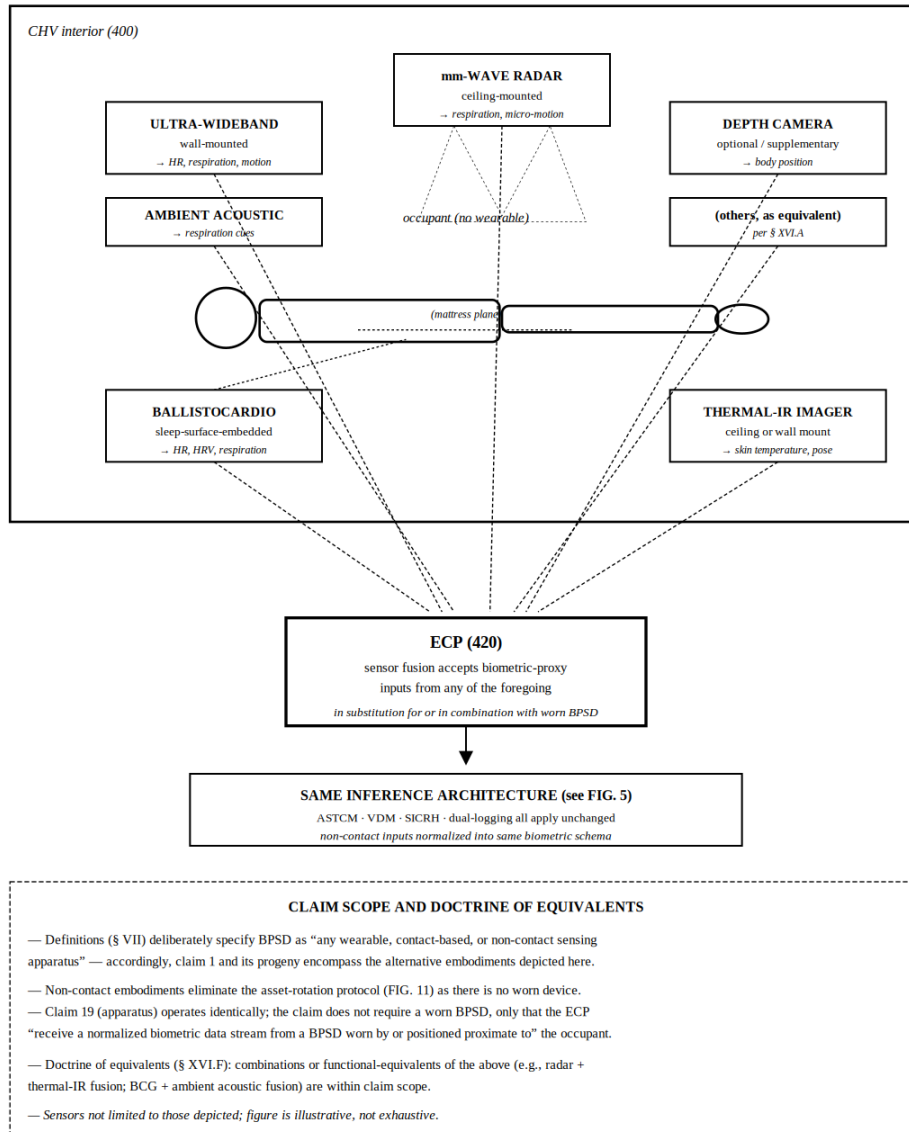


FIG. 12

FIG. 12. Alternative non-contact BPSD embodiment — ballistocardiographic, radar, thermal-IR, depth, and acoustic modalities.