

Claims

1. **[System]** A closed-loop control system for maintaining a target condition in a fluid subject (human or industrial), the system comprising:
 2. a **sensor array** configured to continuously measure at least twenty physiologic or chemical variables of the subject in parallel[49];
 3. a plurality of **replaceable intervention cartridges**, each cartridge including an actuator for delivering or removing a substance to affect the subject and a memory storing at least:
 - (i) a unique cartridge identifier and chemical-class code, (ii) predefined μ -LOW and μ -HIGH threshold values for at least one of said measured variables, and (iii) an incompatibility list of class codes for which simultaneous activation is unsafe[49];
 4. a **controller** operatively connected to the sensor array and cartridges, the controller being programmed with a hierarchical multi-layer control logic that maintains the subject's variables within a homeostatic envelope, wherein:
 - (i) a **vital override layer** automatically disables or adjusts any cartridge when a measured variable exceeds a hard safety limit associated with that cartridge's memory[50];
 - (ii) a **guard-rail layer** suppresses actuation of any cartridge whose class code appears in the incompatibility list of another concurrently active cartridge[50], thereby preventing harmful inter-cartridge interactions;
 - (iii) an **optimizer layer** continuously adjusts one or more cartridge outputs using a feedback control algorithm to reduce deviation of each measured variable from its target range; and
 - (iv) a **dynamic inference layer** that injects probe micro-doses via selected cartridges when all measured variables lie within their respective μ -LOW to μ -HIGH bands, and updates a patient- or process-specific coupling model based on resulting changes in the measured variables[51], the coupling model informing subsequent cartridge adjustments to tailor control to the individual subject.
5. **[Jacobian Learning]** The system of claim 1, wherein the dynamic inference layer is configured to perform adaptive Jacobian learning by applying a probe dose of magnitude proportional to the μ -band width for a given variable and computing a change in that variable[5], updating an estimated sensitivity (partial derivative) via recursive least squares with each probe, and automatically reducing the probe magnitude as the confidence in the sensitivity estimate exceeds a predefined threshold[33].
6. **[Graph AI]** The system of claim 1, wherein the dynamic inference layer comprises a graph-based neural network model that represents the measured variables and cartridge actions as nodes in a graph, infers latent interaction edges between nodes from time-series data, and outputs real-time weight factors or adjustments for each cartridge based on learned relationships among variables, such that the controller prioritizes or limits specific cartridge actions according to the strength of inferred interactions.
7. **[Predictive AI]** The system of claim 1, wherein the dynamic inference layer includes a trained sequence prediction model selected from at least one of: a bi-directional long short-term memory (LSTM) network or a transformer network[9], the model receiving

recent sensor readings and cartridge activity as input and predicting future values of at least one measured variable, and wherein the controller preemptively modifies a cartridge actuation when the model predicts that the measured variable will breach a predefined tolerance (ϵ) within a future time interval[37].

8. **[Dose Ceiling Lockout]** The system of claim 1, wherein each cartridge's memory further stores a maximum cumulative dose per unit time for its substance, and the guard-rail layer automatically **locks out** or limits that cartridge's actuation if a forward-looking dose calculation indicates that the maximum allowable $\text{mg}\cdot\text{kg}^{-1}$ per 24 hours would be exceeded within the next predetermined period (30 minutes)[31].
9. **[Sensor Capacity]** The system of claim 1, wherein the sensor array comprises a plurality of detachable sensor modules (wearable patches, invasive probes, inline sensors) providing, in aggregate, at least 140 simultaneous measurement channels spanning diverse analytes[1], and wherein the controller updates the control outputs at a frequency of at least once every 5 seconds based on the aggregated sensor data[3].
10. **[Adaptive Targets]** The system of claim 1, wherein the dynamic inference layer refines target ranges for one or more variables by analyzing stability of the subject's state over time, such that if the subject consistently remains stable at a variable value near a boundary of a preset range, the controller adapts the target range for that variable to include said value, provided it remains within safe limits.
11. **[Hierarchical Tiers]** The system of claim 1, wherein the controller's hierarchical control logic comprises at least five tiers including, in order of descending priority: an absolute safety interlock tier (emergency stop conditions), a conflict arbitration tier, a dynamic optimization tier (including both feedback and adaptive model inputs), a logging/record-keeping tier, and a manual override tier at the top that can supersede automated actions, such that higher-tier conditions temporarily suspend or modify lower-tier actions until safe conditions are restored[28].
12. **[Manual Override]** The system of claim 1, further comprising a **manual override mechanism** accessible to an operator, which when activated causes the controller to enter a safe manual-control state in which automated adjustments are paused and the operator can directly control the cartridge outputs, and wherein the controller is configured to automatically resume autonomous closed-loop control after a predetermined duration of manual override unless an extension or re-activation of the manual override is made, thereby preventing indefinite suspension of autonomous safety features without confirmation.
13. **[Cartridge Authentication]** The system of claim 1, wherein each replaceable cartridge includes a security element storing cryptographic authentication data, and the controller is configured to verify the authenticity and calibration of a cartridge upon insertion by reading said data[23], and to disable or eject any cartridge that fails authentication or exhibits a calibration error, with the event recorded in the audit log.
14. **[Cartridge Types]** The system of claim 1, wherein the plurality of cartridges comprises a heterogeneous set of intervention types, including at least one **infusion cartridge** for

adding a liquid reagent to the subject, at least one **adsorption or filtration cartridge** for removing or neutralizing an unwanted substance, and at least one **gas exchange or ventilation module** for regulating respiratory gases, all operatively controlled by the controller, such that the closed-loop system can simultaneously manage chemical composition, purity, and gas exchange of the subject fluid[52].

15. **[Audit Logging]** The system of claim 1, further comprising an **immutable audit log** stored in a memory or distributed ledger, wherein each sensor reading, control decision, cartridge actuation, and manual override input is recorded as a timestamped entry that includes a cryptographic hash linking it to a previous entry[11][53], thereby forming a tamper-evident chain of events for the system's operation.
16. **[Audit Interface]** The system of claim 12, wherein the audit log is accessible through a secure read-only interface that allows authorized external systems or personnel to query the log (for example, retrieving data in a standardized format for regulatory review)[54], and wherein any attempt to modify past log entries is detectable via the cryptographic hashes, ensuring retrospective data integrity and compliance with electronic record regulations[13].
17. **[Redundant Controller]** The system of claim 1, wherein the controller comprises a primary control unit and at least one redundant secondary control unit operating in parallel, the system further including a hardware watchdog timer that monitors the performance of the primary control unit[10]; if the primary control unit fails to send a periodic watchdog signal or violates expected operation, the secondary control unit automatically assumes control of the cartridges and the watchdog triggers a fail-safe routine that places the system in a safe operating mode[41], thereby providing continued operation or safe shutdown in the event of a single-point failure.
18. **[Industrial Process Variant]** The system of claim 1, wherein the "subject" is a continuous-flow **industrial process fluid** and the measured variables include chemical properties selected from pH, redox potential, dissolved oxygen, turbidity, metabolite concentrations, or other process parameters, and wherein the cartridges comprise reagent addition modules, filtration units, or other process actuators, the controller being configured to apply the same hierarchical control logic and dynamic inference *mutatis mutandis* to maintain the fluid within a target process envelope[48], such that all references to patient safety and therapeutic intervention in the control logic are interpreted as process safety and industrial process intervention respectively.
19. **[Method]** A method of automatically regulating a biochemical state of a subject using a multi-layer closed-loop system, the method comprising:
 20. **monitoring** a plurality of parameters of the subject via a set of sensors, each sensor producing real-time readings of a biochemical or physical variable;
 21. **detecting** any excursion of a sensor reading beyond a predetermined safe limit and immediately executing a corresponding safety override action, including reducing or stopping a delivery from an associated cartridge or triggering an alarm, to mitigate potential harm;

22. **computing control actions** at a controller by evaluating the sensor readings through a hierarchical algorithm that first filters out any disallowed or unsafe interventions based on stored rules (including inter-therapy incompatibilities and dose limits) and then applies a feedback control law to determine preliminary adjustments for a plurality of therapeutic cartridges;
23. **adapting the control actions** using an adaptive machine-learning model that updates in real time: providing the model with recent sensor trends and applied doses, obtaining adjusted weighting or predictions from the model that reflect the subject's individual response characteristics, and modifying the preliminary cartridge adjustments accordingly to better fit the subject's personalized needs;
24. **administering** one or more doses or other therapeutic outputs to the subject through the respective cartridges based on the adapted control actions, thereby adjusting the subject's parameters toward desired target ranges;
25. during periods when all measured variables of the subject are within their target or μ -band ranges, **injecting a probe micro-dose** via at least one cartridge (of magnitude small enough to stay within the safe range) and measuring the subject's resulting response, then updating the machine-learning model's representation of the subject's dose-response coupling using the measured response to refine future control computations[5];
26. and **logging** each sensor measurement, computed action, model adjustment, and any override event as an immutable record in a time-stamped audit log[34][11], such that a complete trace of the closed-loop control process is preserved for audit and validation.
27. **[Apparatus]** An autonomous therapeutic device for maintaining biochemical homeostasis in a patient, the device comprising the system of claim 1 integrated into a portable unit, including the sensor array and the plurality of cartridges operatively coupled to a common housing and power supply[27], wherein the device is configured to run on battery power for a predetermined duration and to communicate wirelessly with external monitoring systems, thereby providing a self-contained apparatus capable of real-time patient monitoring, adaptive intervention, and secure data logging in compliance with medical device regulations.

Concluding Remarks

The foregoing specification describes illustrative embodiments of a closed-loop multi-layer control system that merges robust rule-based safety with adaptive neural inference to achieve superior regulation of complex biochemical systems. This innovative architecture addresses longstanding challenges in both clinical therapy automation and industrial process control by providing a framework that is at once *adaptive*, *safe*, and *transparent*.

While specific embodiments (including particular sensor types, machine-learning techniques, and use-case scenarios) have been detailed for clarity, these should not be construed as limiting. The invention may be realized with a variety of hardware configurations, control algorithms, and learning models without departing from its core concept: **hierarchically arbitrated closed-loop control augmented with real-time relational learning**. For example, additional layers or sub-layers could be introduced, alternative AI models (e.g. adaptive Bayesian controllers or

federated learning across devices) could be employed, and the technology can be applied to any multi-variable system requiring adaptive regulation.

Modifications and substitutions to the described embodiments will be apparent to those skilled in the art in light of this disclosure. All such variations are intended to be within the scope of the invention as defined in the following claims. The inventors therefore envision that the scope of the present invention be defined by the claims and equivalents thereof, rather than by the specific examples given.