PROVISIONAL APPLICATION FOR PATENT

UNDER 35 U.S.C. 111(b)

CAUSALLY VALIDATED WEARABLE SYSTEM FOR UNCERTAINTY-AWARE COGNITIVE-STATE DETECTION AND BOUNDED SENSORY INTERVENTION

Inventor: Niraj Gohil

Residence: Jersey City, New Jersey, United States

Citizenship: India

Includes specification, 27 claims, 18 figures, and reference-numeral appendix

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] None. This application claims priority to no earlier-filed application. STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] Not applicable. FIELD OF THE DISCLOSURE FIELD OF THE DISCLOSURE

FIELD OF THE DISCLOSURE

[0003] The disclosure relates to wearable computing, physiological signal processing, human-computer interaction, causal control, and bounded sensory actuation. More particularly, it relates to a wearable system that detects a user-defined cognitive condition, including an operationally defined cognitive-noise condition, with an uncertainty-aware state model, decides whether and how to deliver a non-invasive sensory action, verifies the action actually delivered, measures an artifact-qualified proximal response, and adapts an action policy without conflating intervention efficacy with detector correctness.

[0004] The disclosure also relates to active diagnostic probing, transition-hazard estimation, independent output supervision, multi-device dose coordination, and optional stabilization using a psychometrically bounded output. BACKGROUND BACKGROUND

BACKGROUND

[0005] Wearable devices can acquire electroencephalography, electrooculography, photoplethysmography, electrodermal activity, motion, posture, and user-interaction data. Such signals can correlate with attention, autonomic arousal, sleep transition, workload, and movement. A correlation, however, does not establish that a particular cognitive condition is present for a particular user in a particular context.

[0006] A single population threshold or a single spectral ratio can be confounded by motion, muscle activity, ocular activity, fatigue, task difficulty, time of day, electrode contact, or an unfamiliar context. A detector that is forced to classify every input can issue a confident but invalid output when the user or operating context differs from the data on which the detector was trained.

[0007] Closed-loop feedback systems commonly choose a cue after detecting a state and then revise cue parameters according to a later response. That later response may represent spontaneous recovery, regression to the mean, a task change, an artifact caused by the cue, or a true cue effect. Without a deliberately assigned no-output comparator and a recorded assignment probability, the system cannot reliably separate those possibilities.

[0008] A further error occurs when a system treats improvement after an intervention as proof that the original state detector was correct. An intervention can improve performance for reasons unrelated to the inferred state. Conversely, a valid state may be detected even

when a selected intervention has no effect. State classification and treatment-effect estimation are therefore different learning problems.

[0009] Actuation creates another validity problem. A vibration motor can mechanically disturb an EEG electrode. An audio transducer can couple into a microphone or induce an orienting response before an intended outcome window. An optical output can couple into an optical sensor. A fixed waiting interval may end too soon for one mounting condition and waste useful observations in another.

[0010] A software command also does not establish that the commanded output occurred. Battery state, driver saturation, poor contact, mechanical obstruction, clipping, thermal limiting, or a safety rejection can cause the executed dose to differ from the requested dose. Learning from the request alone can associate an outcome with an output that was not delivered.

[0011] Repeated cues can create annoyance, habituation, distraction, or dependence. A policy that maximizes an immediate physiological or behavioral score without burden and dose constraints can learn to intervene too frequently.

[0012] A need therefore exists for an integrated wearable architecture that distinguishes detection evidence, causal intervention evidence, and safety execution evidence. A further need exists for the architecture to preserve an unknown state, to reject invalid sensing windows, to verify the physical output, to qualify a post-output response window from measured artifact, and to restrict online learning to evidence appropriate for the model being updated. SUMMARY SUMMARY

SUMMARY

[0013] In one aspect, a wearable system includes at least one physiological sensor, a sensory output transducer, a main processor, and a safety controller that is operationally independent of the main processor. The main processor produces a state probability, an uncertainty value, and an out-of-distribution score for a user-defined cognitive condition, including an operationally defined cognitive-noise condition. An event becomes eligible only when signal quality, uncertainty, out-of-distribution, context, cooldown, burden, and safety conditions are satisfied.

[0014] For an eligible event, an action randomizer selects among two or more actions that include a no-output action and at least one sensory-output action. The randomizer records a probability assigned to each available action before selection. The selected action and an event identifier are supplied to the safety controller.

[0015] The safety controller checks an immutable or separately protected safety envelope, also referred to herein as a protected output envelope, and authorizes, clips, or rejects the selected action. An executed-dose monitor measures at least one physical or electrical characteristic of the output and produces an executed-dose record that distinguishes a requested action from the output actually produced.

[0016] An artifact monitor estimates residual coupling from the output into one or more sensing channels. A response window opens after the residual remains below a validity threshold for a required number of windows. If the residual does not satisfy the criterion before a maximum wait, the response is marked invalid and is not used to update an action-effect model.

[0017] A causal-effect model, also referred to as the action-effect model, estimates a proximal effect of a sensory-output action relative to the no-output action using eligible events having logged action probabilities, valid executed-dose information, and valid response windows. The effect model updates an action-selection policy subject to safety, burden, habituation, and dose constraints.

[0018] A state detector is updated by a different path. Permitted state evidence includes an independently obtained state label, such as a user thought probe, a controlled task block, or validated task-performance evidence. Permitted evidence may also include the response to an active diagnostic probe if a state-conditional response likelihood for the probe was calibrated from independently labeled data and if probe dose and artifact validity conditions are satisfied.

[0019] In an active diagnostic embodiment, the system computes a state posterior according to P(z given x,r,a_p) proportional to P(r given z,a_p)P(z given x), where z is a candidate state, x is a pre-probe feature vector, a_p is a diagnostic probe, and r is an artifact-qualified response. An ordinary treatment response is excluded from the detector-update path unless the treatment was expressly designated and independently validated as a diagnostic probe.

[0020] In another aspect, the system estimates a hazard that the user will enter a defined cognitive condition within a future horizon. The hazard is trained from independently labeled transition times and is not defined merely as the time difference between a threshold crossing and a later threshold crossing.

[0021] In another aspect, multiple wearable nodes share authenticated timing and a global output token or burden budget. A node can deliver a sensory output only while holding a valid token lease, while each node continues to enforce its own local safety envelope.

[0022] In an optional stabilization aspect, an output is maintained below a user-specific perception criterion only when a psychometric controller has separately established both an upper bound on notice probability and a lower bound on beneficial-response probability. Catch trials and periodic recalibration are used in operation. Failure of either bound disables the optional mode. BRIEF DESCRIPTION OF THE DRAWINGS BRIEF DESCRIPTION OF THE DRAWINGS

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] FIG. 1 is a schematic illustration of a wearer, a head-worn sensing and output device, a wrist sensing node, and a paired task device.

[0024] FIG. 2 is a hardware block diagram showing sensing channels, a main processor, an independent safety controller, an executed-dose monitor, a fixed hardware limiter, and a sensory output.

[0025] FIG. 3 is an end-to-end data-flow diagram that separates state-model updates from causal action-policy updates.

[0026] FIG. 4 is a model-separation diagram for a state detector, a causal-effect model, and a constrained action-selection policy.

[0027] FIG. 5 is a flow diagram of enrollment, calibration, validation, and arming.

[0028] FIG. 6 is a signal-processing diagram for synchronized neural, ocular, autonomic, motion, contact-quality, and behavioral data.

[0029] FIG. 7 is a diagram of uncertainty-aware state inference, an out-of-distribution gate, and a transition-hazard estimate.

[0030] FIG. 8 is a diagram of context, safety, dose, cooldown, and habituation gates followed by randomized action selection.

[0031] FIG. 9 is a timing and logging diagram for randomized action assignments that include a no-output assignment.

[0032] FIG. 10 is a provenance diagram that distinguishes a requested dose from an authorized dose and an executed dose.

[0033] FIG. 11 is a plot of an artifact residual and an adaptively determined response-window opening time.

[0034] FIG. 12 is a block diagram of a causal-effect estimator using randomized action, propensity, executed-dose, context, and artifact-qualified outcome data.

[0035] FIG. 13 is a block diagram of an active diagnostic probe and a state posterior computed from a state-conditional response likelihood.

[0036] FIG. 14 is a diagram of bounded detector learning, bounded policy learning, drift detection, suspension, rollback, and re-enrollment.

[0037] FIG. 15 is a state diagram of the independent safety controller.

[0038] FIG. 16 is a circuit-level block diagram of a fail-low enable path, pulse-width limiter, current-limited driver, and dose-feedback path.

[0039] FIG. 17 is a multi-device coordination diagram with authenticated time synchronization and a global dose token.

[0040] FIG. 18 is a flowchart of an operating method according to an embodiment. DETAILED DESCRIPTION DETAILED DESCRIPTION

DETAILED DESCRIPTION

[0041] The following description enables representative embodiments. Features described for one embodiment can be combined with features described for another embodiment unless the combination is expressly incompatible. The invention is not limited to a particular form factor, sensor family, classifier family, or sensory modality.

[0042] The terms comprising, including, and having are open ended. The term or is inclusive unless the context requires otherwise. A reference to one element can include more than one such element. Numerical values are examples rather than universal biological limits. A range is intended to disclose the stated endpoints and the expressly described intermediate examples, but a later claim should be supported by the specific disclosure on which it relies.

[0043] A cognitive condition is a machine-operational state defined before deployment by a target task, an enrollment protocol, an evidence source, and an output objective. It is not a diagnosis. Examples include a task-attention lapse during a selected computer task, perseverative cognitive arousal during a user-selected recovery period, working-memory overload during a defined training task, and pre-sleep cognitive arousal after a user initiates a sleep-attempt mode.

[0044] As used herein, a cognitive-noise condition is a user-specific machine-operational cognitive condition in which internally generated or externally evoked cognitive activity measurably interferes with a selected task, recovery objective, or sleep objective. Examples include task-unrelated thought, perseverative or repetitive thought, intrusive attentional capture, working-memory interference, and pre-sleep cognitive arousal. The condition is defined through an enrollment protocol, independent label source, context, and predefined outcome rather than by an unobservable semantic content of thought.

[0045] Interruption of a cognitive-noise condition means a validated transition away from the enrolled target state, or a predefined improvement in an artifact-qualified proximal outcome within a specified time window, relative to an eligible no-output comparator or another defined action. Examples include return to task, reduced reaction-time variability, cessation of repetitive interaction, autonomic recovery, or progression toward sleep without an awakening marker.

[0046] Stopping cognitive noise does not require permanent suppression of thought, elimination of ordinary spontaneous cognition, or inference of private semantic thought content. In different embodiments, stopping includes terminating, shortening, reducing the probability of persistence of, or preventing entry into the enrolled cognitive-noise condition, subject to uncertainty, safety, burden, and causal-benefit criteria.

[0047] A state species is one of two or more separately trained cognitive conditions. Different state species can use different sensors, feature subsets, labels, thresholds, hazard horizons, interventions, and proximal outcomes. A multi-head model can share a feature encoder while maintaining a separate calibrated output head for each state species.

[0048] An independent state label is evidence produced without using the success or failure of the ordinary sensory treatment whose effect is being estimated. Examples include a randomized thought probe, a contemporaneous user label, a known controlled task segment, a task error pattern validated against thought probes, or a clinician or experimenter label in an authorized research embodiment.

[0049] A diagnostic probe is a deliberately selected low-burden action whose response distribution has been characterized conditional on independently labeled states. A diagnostic probe can physically resemble a treatment, but a response is used as state evidence only when the probe is identified as diagnostic in the event record and a validation criterion for state separation is satisfied.

[0050] An ordinary treatment is an action selected to change a proximal outcome. Improvement or non-improvement after an ordinary treatment updates the causal-effect model or action policy, not the state detector. This separation prevents a self-confirming loop in which the system declares its own detections correct because its outputs changed the user.

[0051] An executed dose is a measured representation of an output actually produced. Depending on modality, it can include measured current, voltage, back electromotive force, acceleration, displacement, sound pressure proxy, optical flux proxy, onset, offset, pulse count, frequency, duty cycle, ramp, clipping, and reject status.

[0052] An artifact-qualified outcome is an outcome measured only after an output-coupling residual and ordinary signal-quality tests satisfy stored criteria. A fixed blanking time can be used as a minimum, but the opening of a valid window is determined from the measured residual in preferred embodiments.

[0053] A causal effect is an estimated difference between a proximal outcome under one action and the proximal outcome that would have been observed under another action, usually the no-output action, in a comparable context. A causal effect is not the same as the difference between a pre-output value and a post-output value on one occasion.

[0054] FIG. 1 shows a wearer 10, a head-worn housing 100, sensing contacts 104 and 105, a sensory output 106, a local controller 110, a wrist sensor 150, and a paired task device 152. The head-worn housing can be a temple module, eyeglass arm, ear-worn body, headband, cap insert, or adhesive module. The wrist sensor can provide PPG, EDA, temperature, and inertial data. The paired task device can provide task identity, application focus, input timing, error timing, and user labels.

[0055] The local controller can perform the full decision path without a network round trip. A phone or computer can participate in enrollment, model training, review, or task context while the time-critical trigger, safety authorization, and output termination remain local. Local operation reduces latency and prevents a network failure from extending an output.

[0056] FIG. 2 shows neural and ocular sensors 154, autonomic and motion sensors 156, sensor front end 114, main processor 110, memory 112, independent safety controller 116, executed-dose monitor 402, power monitor 124, thermal monitor 128, fixed hardware limiter 118, driver 120, and sensory output 106. The main processor and safety controller

can be different integrated circuits, different cores with a hardware-enforced protection boundary, or a programmable processor and a fixed-function state machine.

[0057] The safety controller receives a bounded action request rather than unrestricted waveform code. For example, the main processor can request an enumerated pattern identifier, amplitude bin, and duration bin. The safety controller resolves the request to a waveform stored in protected memory, applies context-independent hard limits, and produces a fail-low enable signal.

[0058] FIGS. 3 and 4 show the intentional separation of a state detector 502, causal-effect model 500, and action policy 504. State evidence 506 can update detector parameters. Randomized actions and valid outcomes 508 can update causal-effect parameters. Estimated effects, burden measures, and dose limits 510 can update or constrain policy parameters. The data structures can share timestamps and features, but write permissions and update APIs are separated. In FIG. 3, synchronized acquisition 160 supplies the acquired signals, and an execute-or-withhold stage 312 executes or withholds the selected action. Annotation 512 in FIG. 4 denotes the prohibition against relabeling the state detector from an ordinary treatment response.

[0059] In one implementation, the detector is exposed through an inference interface and a detector-update interface. The detector-update interface accepts a label source code. A policy-learning process cannot submit a label source code corresponding to ordinary treatment response. A diagnostic-probe source code is accepted only with a probe validation identifier, an executed-dose validity bit, and an artifact-window validity bit.

[0060] FIG. 5 shows enrollment. At block 600 the user or application selects a target task and a state species. The system collects labeled blocks 602, fits a personal model 604, calibrates probability and uncertainty 606, calibrates any diagnostic-probe likelihood 608, estimates transition hazard 610, establishes context-specific baselines 612, establishes perception and benefit limits 614, and performs hold-out validation before arming 616.

[0061] For a task-attention embodiment, enrollment can alternate intended task performance, rest, deliberately induced distraction, and randomized thought probes. Labels can identify on-task attention, task-unrelated thought, external distraction, fatigue, and unknown. The classifier is not required to reduce all non-target states to a single cognitive-noise label.

[0062] For a perseverative-arousal embodiment, enrollment can include user-labeled neutral, problem-solving, repetitive-thought, and physical-activity intervals. Autonomic signals are interpreted with motion and posture so that exercise-related arousal is not automatically labeled as perseverative cognition.

[0063] For a pre-sleep embodiment, the user explicitly starts a sleep-attempt context. The system can distinguish wakeful cognitive arousal, ordinary wakeful rest, movement, and sleep transition. A daytime classifier is not silently reused for the pre-sleep state species.

[0064] Enrollment data are divided into training, calibration, and hold-out sets by time block rather than by randomly mixing adjacent windows. This reduces leakage from highly correlated neighboring samples. Performance can be required to meet stored criteria for

discrimination, calibration error, false interventions per hour, unknown-state rejection, and latency before automatic operation is armed.

[0065] A population model can initialize a personal model. Personalization can adjust a calibration layer, a state head, feature normalization parameters, or a bounded subset of model weights. The population model remains available for rollback. A user can operate in monitor-only mode until sufficient personal labels have been collected.

[0066] FIG. 6 shows synchronized signal acquisition. EEG and EOG signals 162, PPG and EDA signals 164, and IMU and contact signals 166 are synchronized and resampled at block 200. Time alignment can be maintained by a common clock, timestamp exchange, measured clock offset, and drift correction. Task events from device 152 are placed on the same event time base.

[0067] The sensor front end can acquire one or more EEG channels at 128, 250, 256, 500, 512, 1000, or another suitable number of samples per second. Example electrode locations include frontal, frontopolar, temporal, mastoid, auricular, and reference locations. Dry, semi-dry, wet, capacitive, microneedle, or textile electrodes can be used. Contact impedance, common-mode level, saturation, and electrode motion are monitored.

[0068] Signal conditioning can include DC removal, line-noise suppression, anti-alias filtering, band-pass filtering, robust detrending, and resampling. Ocular, muscle, cardiac, cable, and motion artifact estimates 202 can be produced by regression, adaptive filtering, wavelet methods, source separation, robust subspace methods, a trained artifact model, or combinations thereof.

[0069] A window quality record 204 can include electrode impedance, signal saturation, packet loss, motion energy, ocular energy, muscle energy, line-noise ratio, stationarity, sensor agreement, and a missingness mask. Each derived feature retains provenance to its source window and quality record.

[0070] EEG feature families can include time-domain statistics, Hjorth parameters, line length, zero crossings, entropy, fractal measures, event-related potentials, delta, theta, alpha, beta, and low-gamma band powers, individualized peak frequencies, band-power ratios, cross-frequency coupling, coherence, phase-locking, and graph or connectivity measures. No single listed feature is required in every embodiment.

[0071] Ocular features can include blink rate, blink duration, saccade rate, fixation stability, pupil features when available, and slow eye movement. Autonomic features can include heart rate, RMSSD, SDNN, frequency-domain HRV, Poincare features, pulse amplitude, pulse-arrival measures, tonic EDA, phasic EDA responses, respiration estimates, skin temperature, and recovery slopes.

[0072] Behavioral features can include task accuracy, reaction-time variability, inter-key intervals, pointer trajectory, application switching, scrolling patterns, omission errors, correction rate, gaze-to-task alignment, and resumption latency. Context features can include task identity, posture, motion class, time of day, ambient light, sound level, calendar class with permission, sleep-attempt mode, and vehicle-operation status.

[0073] The feature vector 220 can include raw features, normalized features, temporal derivatives, rolling quantiles, and missingness indicators. Context-specific normalization avoids treating a normal change between sitting and walking as a cognitive event. A feature is not imputed from future data in a real-time decision path.

[0074] FIG. 7 shows an uncertainty-aware inference engine. A personal multi-head classifier 206 produces a state score that is converted by calibrator 606 to state probability 222. Ensemble disagreement, posterior variance, conformal score, Monte Carlo dropout, evidential uncertainty, or another uncertainty estimate is generated at block 208.

[0075] An out-of-distribution score 210 can be produced from feature-space distance, density, reconstruction error, energy score, classifier margin relative to enrolled data, or an ensemble. If uncertainty or out-of-distribution score exceeds a respective limit, the system outputs unknown state 224 and does not create an ordinary treatment eligibility event.

[0076] The detector can require both a probability criterion and temporal persistence. Persistence can be implemented by a dwell counter, a hidden Markov model, a semi-Markov model, a sequential probability ratio, a cumulative-sum statistic, or a state-space filter. Hysteresis can require a lower release threshold before another dwell begins.

[0077] The transition-hazard estimator 610 predicts entry into a labeled state within a horizon H. Suitable models include a discrete-time survival model, proportional-hazard model, recurrent state-space model, temporal point-process model, or calibrated classifier trained on windows preceding independently labeled onset times. Censoring and competing states can be represented.

[0078] A preemptive eligibility event can be created when the transition hazard exceeds a hazard threshold, the expected action benefit is positive for the current context, and all quality and safety gates pass. A lead-time distribution is evaluated on held-out labeled transitions. Merely detecting a signal earlier than a user's delayed button press is not by itself treated as proof of pre-awareness detection.

[0079] FIG. 8 shows the trigger path. Eligible state event 214 passes through task and receptivity gates 230 and dose, cooldown, and habituation gates 232. Receptivity can be estimated from task phase, recent dismissal, motion, conversation, notification burden, and user preference. Automatic cues can be disabled during vehicle operation, machinery use, crossing a street, a presentation, or another configured protected context.

[0080] The gate can require a minimum expected benefit for at least one action. If every action has a non-positive lower confidence bound after burden cost, the available action set contains only no output. The system does not deliver an action solely because a state probability is high.

[0081] The randomizer 300 selects from no output 304, an orienting haptic cue 306, a paced haptic cue 308, and an audio or visual cue 310, or another configured set. Randomization probabilities are logged before the action is executed and are bounded away from zero for actions used in causal estimation.

[0082] An orienting haptic cue can include one to five pulses, each about 20 to 500 milliseconds, separated by about 20 milliseconds to 2 seconds, with a total pattern duration of about 50 milliseconds to 10 seconds. A paced cue can provide repeated tactile pulses corresponding to inhalation, exhalation, task pacing, or a brief reset. Values are examples to enable implementation and are constrained by user comfort and the safety envelope.

[0083] Audio output can be provided by an open-ear speaker, earphone, bone-conduction transducer, or paired device. It can include a tone, click, spoken label, breathing pace, or task-resumption cue. Optical output can include a peripheral icon or brightness change. An action can also change the paired task environment, such as briefly reducing non-target notifications, if the user has authorized that control.

[0084] The preferred consumer-wellness embodiments use mechanical haptic, acoustic, or optical sensory outputs. A transcutaneous electrical output can be included in a separately validated embodiment with electrode-contact monitoring, charge balancing where applicable, fixed current and pulse-width limits, and an appropriate regulatory and safety design. No transcranial electrical output is required by the invention.

[0085] FIG. 9 illustrates repeated eligibility events and assignments along a time axis 320. The event log 302 stores an event identifier, context vector or context hash, available actions, probability of each available action, selected action, randomization-seed commitment or audit value, model version, and time. A no-output assignment is an intentional action, not a safety rejection, as indicated by no-output eligibility annotation 322.

[0086] A rejected action is separately coded. If the safety controller rejects an action because of temperature, contact, or dose history, the rejection may be related to context and cannot automatically be treated as a randomized no-output control. A primary per-protocol effect model can use only honored assignments with valid outcomes. An intention-to-treat analysis can retain rejected assignments to estimate the effect of assignment, provided the analysis and missing-outcome handling are specified.

[0087] FIG. 10 shows requested action 400, safety authorization 116, driver 120, physical dose sensors 402, executed-dose reconstruction 404, and signed or integrity-protected event record 406. The field set 407 can store requested, authorized, and measured values, onset and offset times, clipping, rejection, fault, contact quality, and firmware versions.

[0088] For a haptic motor, the executed-dose monitor can combine driver current, voltage, motor back electromotive force, and a housing accelerometer. For a piezoelectric actuator, it can combine drive voltage, current, and sensed displacement or acceleration. For an audio output, it can combine electrical drive with a microphone or transducer model. For an optical output, it can combine drive current with a photodiode or calibrated emitter model.

[0089] The executed-dose record is associated with the causal observation. If the measured amplitude, duration, or onset differs from its requested value beyond a permitted tolerance, that is, when the executed dose fails a dose-validity criterion, the observation is coded as clipped, partial, or rejected. Dose-response learning can use measured dose. Assignment-effect learning can use the randomized assignment and separately model compliance.

[0090] A cryptographic message authentication code, hash chain, secure counter, or write-once event sequence can protect event provenance. Cryptography is not required for every embodiment, but event identifiers and monotonic sequence checks prevent an outcome from being joined to the wrong action.

[0091] FIG. 11 shows adaptive artifact settling under control of an artifact settling controller 408. During output and immediately after output, the system estimates a residual artifact ratio 420, rho(t). The ratio can compare energy predictable from the actuator reference with background signal energy, compare an artifact classifier score with a limit, or combine multiple coupling metrics.

[0092] A response window 410 opens, at a time indicated by response-window opening marker 426, when rho(t) remains below a stored artifact limit 422 for M consecutive analysis windows following the output offset marker 424, when ordinary sensor-quality criteria pass, and when a minimum blanking time has elapsed. M can be two, three, five, ten, or another validated integer. If no valid opening occurs before a maximum delay, the outcome is missing for causal learning and no imputed improvement is assigned.

[0093] Artifact templates can be learned during doffed self-test, quiet wearing, or controlled low-burden calibration. Template subtraction is followed by a residual-validity test. The existence of a template does not make a contaminated response valid by itself.

[0094] The proximal outcome 412 depends on state species. For a task-attention lapse, it can include task resumption, reaction-time stabilization, error reduction, gaze return, and an artifact-qualified neural marker. For cognitive arousal, it can include autonomic recovery, cessation of repetitive interaction, and user confirmation. For pre-sleep cognitive arousal, it can include reduced arousal markers and progression toward a sleep-transition state without an awakening marker.

[0095] An outcome can be scalar or vector. A composite score uses weights fixed before the event or learned only from independent validation data. Outcome windows, sign conventions, and missingness rules are versioned. A policy is not allowed to redefine an outcome after observing whether a selected action appeared successful.

[0096] FIG. 12 shows causal-effect model 500, which receives a validated causal event tuple 518. For action a relative to no output 0, a context-specific effect can be written tau_a(x) = E[Y(a) - Y(0) given X=x]. The sign of Y is selected so that a positive value represents the desired proximal change.

[0097] One suitable doubly robust pseudo-outcome is phi_a = m_a(X) - m_0(X) + I(A=a)/pi_a(X) times [Y - m_a(X)] - I(A=0)/pi_0(X) times [Y - m_0(X)], where m_a is an outcome model, pi_a is the logged assignment probability, and I is an indicator. Other inverse-propensity, weighted regression, Bayesian hierarchical, randomization-based, or causal-forest estimators can be used. FIG. 12 further shows baseline outcome models 520, a propensity-weighted residual 522, and an invalid-observation exclusion 524 that removes dose-invalid or artifact-invalid observations from the effect-model update.

[0098] Randomization is preferably stratified or blocked by a context summary so that action and no-output observations accumulate in relevant contexts. Propensities remain bounded to maintain positivity. A minimum sample count, effective sample size, confidence interval, or posterior probability is required before an action is declared beneficial for a context.

[0099] The action policy can be a constrained contextual bandit, actor-critic controller, Bayesian decision rule, or table. An example objective selects an action to maximize estimated causal benefit minus weighted annoyance, distraction, energy, and habituation, subject to the hard safety envelope, a rolling output-dose budget, a minimum exploration probability for each action retained for causal estimation including a minimum no-output exploration probability, and user preferences.

[0100] Exploration is not increased during high-risk or protected contexts. The no-output probability can increase when uncertainty about state or effect is high. The policy can select monitor-only operation if the lower confidence bound of benefit is not positive.

[0101] The causal-effect model can include delayed outcomes, but a proximal outcome used for online adaptation is bounded to a defined time window. Long-term claims of improved capability require separate validation. The system can record long-horizon task metrics for offline review without allowing them to override the local safety controller.

[0102] FIG. 13 shows the active diagnostic probe. Candidate latent states 530 include a candidate state z having a pre-probe state prior 540, P(z given x), obtained from pre-probe features. The randomizer selects a low-burden diagnostic probe 532, a_p. An artifact-qualified probe response 534, r, is measured. A calibrated response model supplies the state-conditional response likelihood 536, P(r given z,a_p). The system computes a probe-informed state posterior 538 proportional to P(r given z,a_p)P(z given x).

[0103] Examples of r include an event-related EEG component, an orienting heart-rate change, a blink or saccade response, a microgesture, a task-resumption latency, or a vector of such features. The response likelihood can be represented by a Gaussian model, mixture model, discriminative likelihood ratio, neural density estimator, or calibrated classifier.

[0104] Before the response can update the state detector, validation gate 542 requires that the probe was randomly or protocol-selectively designated as diagnostic, the executed dose was within tolerance, the response window was artifact-qualified, and state-conditional response distributions demonstrated a stored separability criterion on independently labeled data. Suitable criteria include a minimum likelihood ratio, cross-validated area under a receiver-operating curve, calibration error, or mutual information with confidence bounds.

[0105] The probe model is user-specific in preferred embodiments. A population probe model can initialize operation, but detector updates remain disabled until personal validation or a conservative transfer criterion passes. Probe responses obtained in an unknown or out-of-distribution context do not update the detector.

[0106] A diagnostic probe can be interleaved with ordinary treatment trials. Its purpose code, action probability, and analysis window are recorded. If a single physical pattern serves

both purposes, the record states which inference path is permitted for the event. A later observed benefit does not retroactively convert an ordinary treatment into a diagnostic probe.

[0107] FIG. 14 shows bounded learning. Validated state evidence updates the detector by a bounded step 550. Causal action effects update the policy by a bounded step 556. Calibration monitor 552, drift score 554, and burden monitor 558 can suspend operation 560 and cause rollback or re-enrollment 562.

[0108] Drift can be detected from feature-distribution shift, increased unknown rate, calibration error on new labels, rising false-cue reports, loss of probe separability, changed electrode configuration, or degradation of executed-dose consistency. A minor drift can select monitor-only mode. A major drift can invalidate a model and require enrollment.

[0109] Detector thresholds are not continuously moved merely to maintain a desired intervention rate. A desired rate can constrain burden, but detector calibration is assessed against independent labels. This prevents the device from redefining the state to justify a target number of cues.

[0110] User dismissal, suppression, or negative rating is stored as burden and preference evidence. It can reduce selection probability or suspend a modality. It is not automatically a state label unless the user expressly answers a state-label question designed for that purpose.

[0111] A user-initiated cue can provide evidence that the user wanted an action. It is not automatically a false-negative label. If the interface separately asks whether the target state was present, that answer can be stored as an independent state label with a source code and confidence.

[0112] FIG. 15 shows safety states OFF 700, ARMED 702, ENABLED 704, and LATCHED FAULT 706. Output is disabled by default. Transition to ARMED requires self-test, valid protected limits, acceptable power and temperature, and a heartbeat. Transition to ENABLED requires a valid bounded request and a valid sensor or contact condition applicable to the output.

[0113] The safety controller terminates an output at the earlier of requested termination and an independent maximum duration. It enforces maximum amplitude, pulse width, duty cycle, actions per interval, cumulative activation time, rate of amplitude change, temperature, and modality-specific limits. Limits can be lower than hardware maxima and can differ by user-selected comfort profile.

[0114] A watchdog timeout, main-processor reset, invalid request code, message-authentication failure, brownout, overtemperature, overcurrent, driver fault, sensor disconnection, token loss, or manual suppression causes deassertion of enable, indicated by enable-deassertion annotation 708 in FIG. 15. Selected faults are latched until the condition clears and the user acknowledges or power cycles according to the validated safety plan.

[0115] FIG. 16 shows a main processor request, safety controller, fail-low gate 720, waveform source 722, pulse-width limiter 118, current-limited driver 120, sensory output 106, current and motion sense 402, and ADC feedback 724. The fail-low gate can require simultaneous valid levels from the request path, safety controller, watchdog, and hardware limiter.

[0116] The pulse-width limiter can use a retrigger policy that cannot extend output beyond a fixed maximum, a monostable, an RC time constant, a hardware timer clocked independently of the main processor, or combinations. The current-limited driver can include a sense resistor, comparator, foldback, fuse, thermal cutoff, or dedicated driver limit.

[0117] Protected safety limits are stored in read-only memory, one-time programmable memory, signed firmware verified by the safety controller, or a region not writable through the ordinary communication interface. An action-policy update cannot widen a protected hard limit.

[0118] FIG. 17 shows a head node 800, wrist node 802, paired task device 804, coordinator 806, and global dose manager 130. Authenticated clock synchronization allows the system to join signals and outputs across nodes and to distinguish simultaneous from sequential actions.

[0119] A global dose token can be a lease having a holder identifier, modality, maximum activation, expiration time, and sequence number. Loss of communication causes the lease to expire. Each output-capable node also enforces a local envelope, so a coordinator failure cannot remove local limits.

[0120] A rolling burden budget, enforced across nodes as a global rolling output budget, can combine actions of different modalities using modality-specific weights. The coordinator can prevent two nodes from issuing redundant cues or can deliberately coordinate a pattern when that coordinated pattern was enrolled as one action. The event record identifies all participating nodes and their executed doses.

[0121] The optional stabilization mode is distinct from ordinary discrete cueing, and its maintained output is referred to herein as a stabilization output. A psychometric controller estimates a probability of conscious notice as a function of measured executed dose and context. Calibration uses ascending and descending trials, adaptive Bayesian or staircase selection, catch trials with no output, and signal-detection analysis of hits and false alarms.

[0122] The controller also estimates a probability or magnitude of a predefined beneficial proximal response. Stabilization is permitted only if a dose region satisfies both P(notice given dose,context) less than alpha and P(benefit given dose,context) greater than beta, with confidence or credible bounds. Example alpha values include 0.05, 0.10, and 0.20. Example beta values include 0.55, 0.60, 0.70, and 0.80.

[0123] In-use catch trials and periodic perceptual questions monitor the notice model. Randomized no-output periods monitor benefit. If no dose region satisfies both criteria,

stabilization is unavailable and the system uses discrete cues or monitor-only operation. The specification does not assume that an imperceptible output improves cognition.

[0124] The system can process raw neural and physiological samples locally and retain them only in a volatile analysis buffer. Nonvolatile logs can store derived features, validity bits, model versions, action probabilities, executed-dose records, outcomes, and user-approved labels. Raw data export requires affirmative permission in a preferred privacy mode.

[0125] Model and firmware updates can be signed and versioned. The safety controller verifies its own image and protected limit table. A rollback partition stores a last-known-good version. Update failure leaves outputs disabled.

[0126] The system provides a manual pause, modality control, visible operating-mode indicator, and deletion controls. It can explain a cue using a concise reason code such as state-confidence pass, transition-hazard pass, user schedule, or diagnostic probe. An explanation does not expose raw neural data unless the user requests it.

[0127] The device is not a substitute for medical evaluation, and consumer-wellness embodiments avoid representing a cognitive-condition estimate as a diagnosis. Embodiments intended to diagnose, treat, or prevent disease require applicable clinical validation, risk management, human-factors work, and regulatory authorization.

[0128] FIG. 18 summarizes operation. The system defines a target state, acquires qualified features, infers state probability, uncertainty, and transition hazard, checks context and safety, randomizes action or no action, authorizes and executes the action, records executed dose, waits for artifact validity, measures a proximal outcome, estimates causal action effect, updates the constrained policy, and monitors drift and burden.

[0129] A reference implementation can use a four-channel frontal and temporal EEG front end at 256 samples per second, a two-channel EOG derivation, a wrist PPG and EDA node, and a six-axis IMU. Windows of two seconds with a 250-millisecond hop can produce spectral, ocular, autonomic, motion, and task features. These values are an enabling example, not a required configuration.

[0130] In that implementation, an eligible task-attention event requires a calibrated lapse probability of at least 0.75 for three consecutive hops, uncertainty below a validated limit, OOD score below a validated limit, acceptable contact, an active enrolled task, no protected context, and no cue within the preceding sixty seconds. The values can be modified after validation but are not silently learned from ordinary treatment response.

[0131] At each eligible event, the system can assign no output with probability 0.20, a single orienting haptic pattern with probability 0.40, and a paced haptic pattern with probability 0.40. The exact probabilities depend on study stage and burden. All probabilities are written before actuation.

[0132] A housing accelerometer and driver current monitor reconstruct the haptic dose. EEG outcome measurement begins only after the actuator-coupling residual stays below its limit

for three consecutive windows. The proximal outcome combines task resumption within fifteen seconds and change in validated attention features over a later ten-second window.

[0133] The foregoing implementation is a prophetic example unless and until it is actually built and tested. Numerical performance, detection accuracy, benefit, and user experience must be measured. No unperformed result is represented as an actual experimental result.

[0134] Alternative embodiments can omit EEG and use ocular, autonomic, behavioral, and contextual sensing when those signals meet validation criteria for the selected state species. Other embodiments can use EEG without a wrist node. The inventive separation of state evidence, causal effect evidence, executed-dose evidence, and safety authorization does not depend on a particular sensor count.

[0135] Alternative embodiments can operate in monitor-only mode, diagnostic-probe-only mode, discrete-cue mode, stabilization mode, or hybrid mode. A mode transition is logged and subject to the same protected safety and burden controls.

[0136] The functions described as models or modules can be implemented by software instructions, digital logic, analog circuitry, a field-programmable gate array, an application-specific integrated circuit, or combinations. A model can be a table, equation, statistical model, machine-learning model, or state machine.

[0137] The described embodiments provide multiple independently useful subcombinations, including executed-dose-linked causal learning, artifact-qualified outcome timing, validated diagnostic-probe state updating, uncertainty and unknown gating, independent safety authorization, and multi-device global burden control. The applicant intends to preserve support for claims directed to any disclosed subcombination that is consistent with the contribution of the human inventor.

[0138] As used herein, an induction mode is an operating mode in which the system delivers or maintains a bounded sensory output to promote a validated transition toward a predefined target state, or to prevent entry into an enrolled cognitive-noise condition, rather than to interrupt a condition already detected. Induction-mode operation uses the same eligibility, randomization, executed-dose, artifact-validity, safety, and burden machinery described above; the difference lies in the triggering evidence and in the sign convention of the predefined proximal outcome.

[0139] In a preemptive induction embodiment, an induction-mode eligibility event is created when the transition hazard described with reference to FIG. 7 and paragraphs [0077] and [0078] exceeds a hazard threshold and the expected action benefit is positive for the current context. The action randomizer 300 continues to include the no-output action 304 so that the causal-effect model 500 estimates the preemptive benefit relative to no output, and the lead-time distribution is evaluated on held-out labeled transitions as described in paragraph [0078].

[0140] In a sleep-induction embodiment, the user initiates the sleep-attempt context of paragraph [0063], and a paced sensory output of the kind described in paragraphs [0082] and [0083] is delivered or maintained to promote progression toward a sleep-transition state without

an awakening marker, the predefined proximal outcome being defined as in paragraph [0094]. Task-attention and recovery-period induction embodiments analogously pace task re-engagement or autonomic recovery using the disclosed paced cues.

[0141] In a maintained induction embodiment, the induced output is a stabilization output operated under the psychometric bounds of paragraphs [0121] to [0123], including the upper bound on notice probability, the lower bound on beneficial-response probability, in-use catch trials, and randomized no-output periods. Failure of either bound disables maintained induction, and the system reverts to discrete cues or monitor-only operation.

[0142] Induction-mode operation remains subject to the independent safety controller 116, the protected output envelope, the executed-dose monitor 402, artifact-qualified outcome windows, the dose, cooldown, and habituation gates 232, protected-context suppression as described in paragraph [0079], and mode-transition logging as described in paragraph [0135]. An induction-mode outcome is not used to relabel the state detector 502 except through the validated diagnostic-probe pathway.

CLAIMS

What is claimed is:

1. A wearable system comprising: at least one physiological sensor configured to acquire a physiological signal from a wearer; a sensory output transducer; a main processor; an executed-dose monitor configured to measure an output produced by the sensory output transducer; and a safety controller operationally independent of the main processor; wherein the main processor is configured to generate, from the physiological signal, a state probability and an uncertainty measure for a user-defined cognitive condition; create an eligible event only when the state probability satisfies a state criterion, the uncertainty measure satisfies an uncertainty criterion, and a signal-quality criterion and a context criterion are satisfied; select, for the eligible event, an assigned action from an available action set that includes a no-output action and a sensory-output action according to stored assignment probabilities, and record the assignment probabilities; cause the safety controller to authorize, limit, or reject the assigned action according to a protected output envelope; associate with the eligible event an executed-dose record produced from a measurement of the output actually produced; determine an opening of a response window according to a measured residual of output-induced artifact; estimate, using eligible events having recorded assignment probabilities and artifact-qualified outcomes, a causal effect of the sensory-output action relative to the no-output action; and update an action-selection policy according to the causal effect; wherein a state detector that generates the state probability is updated only from an independently obtained state label or from a response to a diagnostic probe having a validated state-conditional response model, and is not updated from an estimated causal effect of an ordinary sensory-output action.

2. The wearable system of claim 1, wherein the at least one physiological sensor comprises two or more of an electroencephalography sensor, an electrooculography sensor, a photoplethysmography sensor, an electrodermal-activity sensor, an inertial sensor, a temperature sensor, and a contact-quality sensor, and wherein the main processor further uses a behavioral feature received from a paired task device.

3. The wearable system of claim 1, wherein the user-defined cognitive condition is an operationally defined cognitive-noise condition selected from separately trained state species that include a task-attention lapse, perseverative cognitive arousal, working-memory overload, and pre-sleep cognitive arousal, each state species having a respective label source and proximal outcome.

4. The wearable system of claim 1, wherein the main processor is further configured to generate an out-of-distribution score and to produce an unknown-state output that prevents creation of the eligible event when the out-of-distribution score fails an out-of-distribution criterion.

5. The wearable system of claim 1, wherein the main processor is further configured to estimate, from independently labeled state-transition times, a hazard of entry into the user-defined cognitive condition within a future horizon, and wherein the state criterion is satisfied when

the hazard exceeds a hazard threshold, whereby the eligible event is created before the state probability alone satisfies the state criterion.

6. The wearable system of claim 1, wherein each assignment probability is written to an event record before the assigned action is supplied to the safety controller and is bounded away from zero for each action used by a causal-effect estimator.

7. The wearable system of claim 1, wherein the executed-dose record includes a requested dose, an authorized dose, a measured onset, a measured offset, a measured amplitude proxy, and a clip or reject status, and wherein the eligible event is excluded from estimation of the causal effect when the executed-dose record fails a dose-validity criterion.

8. The wearable system of claim 1, wherein the response window opens only after the measured residual of output-induced artifact remains below an artifact limit for a required plurality of analysis windows, and wherein an outcome is marked invalid if the artifact limit is not satisfied before a maximum delay.

9. The wearable system of claim 1, wherein the causal effect is estimated using a propensity-weighted or doubly robust estimator that receives the recorded assignment probabilities, a context vector, the assigned action, the executed-dose record, and the artifact-qualified outcome.

10. The wearable system of claim 1, wherein the diagnostic probe is a randomized low-burden sensory pattern, and wherein the main processor is configured to compute a state posterior proportional to a product of a pre-probe state probability and a state-conditional likelihood of an artifact-qualified probe response.

11. The wearable system of claim 1, wherein the safety controller asserts a fail-low enable signal, independently limits amplitude and duration, terminates output in response to loss of a heartbeat from the main processor, and controls a hardware limiter whose protected limit is not writable by the action-selection policy.

12. The wearable system of claim 1, wherein the action-selection policy is constrained by a rolling output-dose budget, a cooldown interval, an intervention-burden measure, a habituation measure, and a minimum probability of selecting the no-output action.

13. The wearable system of claim 1, further comprising an additional output-capable wearable node and an authenticated coordinator configured to grant a time-limited global output token, wherein each output-capable wearable node of the wearable system is configured to enforce a local protected output envelope even while holding the global output token.

14. A computer-implemented method performed by a wearable system, the method comprising: obtaining a pre-probe feature vector from at least one physiological signal; generating a prior probability for each of a plurality of candidate cognitive states; assigning, according to a recorded probability, a diagnostic probe or no diagnostic probe; when the diagnostic probe is assigned, verifying an executed dose of the diagnostic probe and delaying a probe-response window until a measured artifact residual satisfies an artifact criterion; obtaining a

probe response from the probe-response window; retrieving a state-conditional response likelihood that was calibrated using cognitive-state labels obtained independently of a treatment response; computing a posterior probability for at least one candidate cognitive state from the prior probability and the state-conditional response likelihood; updating a cognitive-state detector according to the posterior probability only when a probe-model validation criterion, an executed-dose criterion, and the artifact criterion are satisfied; separately assigning a treatment action or a no-treatment action at a treatment-eligible event; and updating an action-selection policy from a causal effect estimated from randomized treatment assignments without using the causal effect as a state label.

15. The computer-implemented method of claim 14, further comprising recording, before each treatment assignment, a probability of the treatment action and a probability of the no-treatment action, and maintaining each recorded probability above a minimum exploration probability while the respective action remains available for causal estimation.

16. The computer-implemented method of claim 14, further comprising coding a safety-rejected treatment assignment separately from a randomized no-treatment assignment and preventing the safety-rejected assignment from being automatically used as a no-treatment control.

17. The computer-implemented method of claim 14, further comprising enabling a stabilization output only when a psychometric model establishes an upper confidence bound on notice probability below a first limit and an effect model establishes a lower confidence bound on beneficial-response probability above a second limit, and disabling the stabilization output when either condition fails during catch-trial monitoring.

18. The computer-implemented method of claim 14, further comprising computing a drift score from at least one of feature-distribution shift, calibration error, unknown-state rate, diagnostic-probe separability, executed-dose inconsistency, and user-reported burden, and responsive to the drift score satisfying a suspension criterion, rolling back to a validated model or requesting re-enrollment.

19. A non-transitory computer-readable medium storing instructions that, when executed by one or more processors of a system comprising a plurality of wearable nodes, cause the system to: synchronize physiological signals, task events, and sensory-output events to a common time base; generate an eligible cognitive-state event only after signal-quality, uncertainty, out-of-distribution, context, cooldown, and burden gates pass; select an action including a no-output action under logged randomization probabilities; obtain from an output-capable node an executed-dose record and a local-safety authorization; open an outcome window according to a measured output-artifact residual; update a causal action-effect model from an artifact-qualified outcome without using the outcome to label a state detector; and enforce, across the plurality of wearable nodes, a global rolling output budget in addition to a protected local output envelope enforced at each output-capable node.

20. The non-transitory computer-readable medium of claim 19, wherein the instructions further cause the system to retain raw physiological samples in a volatile analysis buffer, store

derived validity and event records in nonvolatile memory, and prevent transmission of the raw physiological samples absent an affirmative user authorization.

21. A wearable system operable in an induction mode, the wearable system comprising: at least one physiological sensor configured to acquire a physiological signal from a wearer; a sensory output transducer; an executed-dose monitor configured to measure an output produced by the sensory output transducer; a safety controller; and a main processor; wherein the main processor is configured to estimate, from the physiological signal, a transition hazard of entry of the wearer into a predefined cognitive state within a future horizon, the transition hazard being trained from independently labeled state-transition times; create, in the induction mode, an eligibility event when the transition hazard satisfies a hazard threshold and a signal-quality criterion and a context criterion are satisfied; select, for the eligibility event, an assigned action from an available action set that includes a no-output action and at least one sensory-output action according to recorded assignment probabilities; and estimate, using eligibility events having recorded assignment probabilities, executed-dose records produced by the executed-dose monitor, and outcomes measured in response windows opened according to a measured residual of output-induced artifact, a causal effect of the at least one sensory-output action relative to the no-output action on a predefined proximal outcome representing promotion of a validated transition toward a predefined target state or prevention of entry into an enrolled cognitive-noise condition; and wherein the safety controller is configured to authorize, limit, or reject the assigned action according to a protected output envelope.

22. The wearable system of claim 21, wherein the predefined target state is a sleep-transition state, the induction mode is entered when the wearer initiates a sleep-attempt context, and the predefined proximal outcome includes progression toward the sleep-transition state without an awakening marker.

23. The wearable system of claim 21, wherein the at least one sensory-output action comprises a paced sensory output providing repeated pulses corresponding to inhalation, exhalation, task pacing, or a brief reset.

24. The wearable system of claim 21, wherein the at least one sensory-output action comprises a stabilization output maintained below a user-specific perception criterion only while a psychometric controller establishes an upper bound on notice probability and a lower bound on beneficial-response probability, the stabilization output being disabled upon failure of either bound during catch-trial monitoring.

25. The wearable system of claim 21, wherein induction-mode operation is subject to a cooldown interval, a rolling output-dose budget, an intervention-burden measure, and suppression during a configured protected context, and wherein an outcome measured in the induction mode is not used to update a state detector except through a response to a diagnostic probe having a validated state-conditional response model.

26. The wearable system of claim 1, wherein the executed-dose record, the recorded assignment probabilities, and an artifact-qualified outcome are joined in an event record protected by at

least one of a cryptographic message authentication code, a hash chain, a secure counter, or a write-once event sequence, and wherein a monotonic sequence check prevents an outcome from being associated with a wrong action.

27. The wearable system of claim 1, wherein creation of the eligible event is further conditioned on a receptivity estimate based on at least one of task phase, recent dismissal, motion, conversation, or notification burden, and wherein automatic sensory-output actions are disabled during a configured protected context comprising at least one of vehicle operation, machinery use, or another user-configured protected activity.

ABSTRACT

A wearable system detects a user-specific cognitive condition, including an operationally defined cognitive-noise condition, and selects a bounded sensory action, including a no-output action, under logged randomization probabilities. An independent safety controller authorizes the requested action, and an executed-dose monitor records the output actually produced. A response window opens only after a measured actuation-artifact residual satisfies a validity criterion. A causal-effect model estimates an action effect from randomized, dose-valid, artifact-valid observations and updates a constrained action-selection policy. A state detector with uncertainty and out-of-distribution gates is updated through a separate path using an independently obtained state label or a validated diagnostic-probe response; an ordinary treatment response is not used to relabel the detected state. Task-attention, cognitive-arousal, pre-sleep, and hazard-triggered induction-mode embodiments preserve causal attribution and fail-low output control.

APPENDIX A - REFERENCE NUMERALS

Reference	Component or function
10	wearer
100	wearable housing
104	EEG or physiological contact
105	EOG or auxiliary contact
106	sensory output transducer
110	main processor or local controller
112	memory
114	sensor front end
116	independent safety controller
118	fixed hardware limiter
120	output driver
124	power monitor
128	thermal monitor
130	global dose manager
150	wrist sensor
152	paired task device
154	neural and ocular sensors
156	autonomic and motion sensors
160	synchronized acquisition
162	raw EEG and EOG signals
164	PPG and EDA signals
166	IMU and contact signals
200	signal synchronization and conditioning
202	artifact estimator
204	quality record
206	state classifier
208	uncertainty estimator
210	out-of-distribution estimator
214	eligible state event
220	feature vector
222	state probability
224	unknown-state gate
230	task and receptivity gates
232	dose, cooldown, and habituation gates
300	action randomizer
302	assignment-probability log
304	no-output action
306	orienting haptic action
308	paced haptic action
310	audio or visual action
312	execute-or-withhold stage
320	time axis
322	no-output eligibility annotation
400	requested action
402	executed-dose sensor or monitor
404	executed-dose reconstruction or record
406	signed event record
407	event-record field set
408	artifact settling controller
410	artifact-qualified response window
412	proximal outcome
420	residual artifact ratio
422	artifact limit
424	output offset marker

426	response-window opening marker
500	causal-effect model
502	state detector
504	constrained action policy
506	independent state evidence
508	randomized action and valid outcome evidence
510	benefit, burden, and dose limits
512	detector-relabeling prohibition annotation
518	validated causal event tuple
520	baseline outcome models
522	propensity-weighted residual
524	invalid-observation exclusion
530	candidate latent states
532	diagnostic probe
534	artifact-qualified probe response
536	state-conditional response likelihood
538	probe-informed state posterior
540	pre-probe state prior
542	probe validation gate
550	bounded detector update
552	calibration monitor
554	drift score
556	bounded policy update
558	burden and habituation monitor
560	suspension state
562	rollback or re-enrollment
600	target-state definition
602	labeled task block collection
604	personal state model fitting
606	probability calibrator
608	probe-response calibration
610	transition-hazard model
612	context-specific baseline setting
614	perception and benefit limit setting
616	hold-out validation and arming
700	off state
702	armed state
704	enabled state
706	latched fault state
708	enable-deassertion annotation
720	fail-low enable gate
722	waveform source
724	dose-feedback converter
800	head-worn node
802	wrist node
804	task device
806	authenticated coordinator

FIG. 1. Wearable and paired-device embodiment

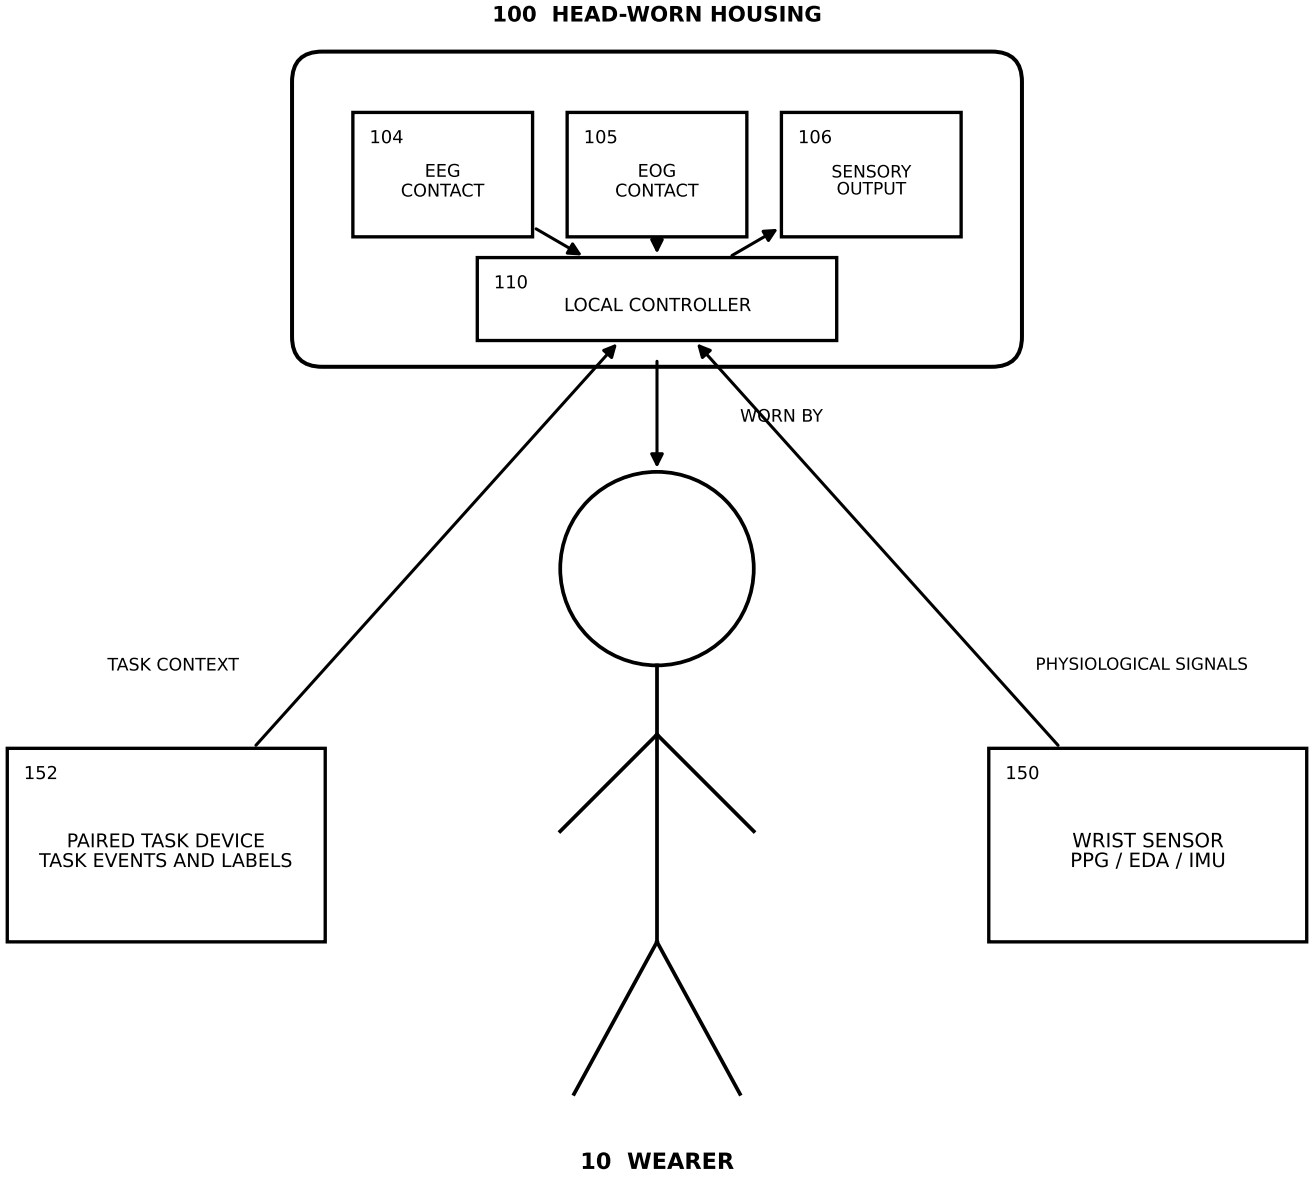

FIG. 1

FIG. 2. Hardware architecture with independent safety execution

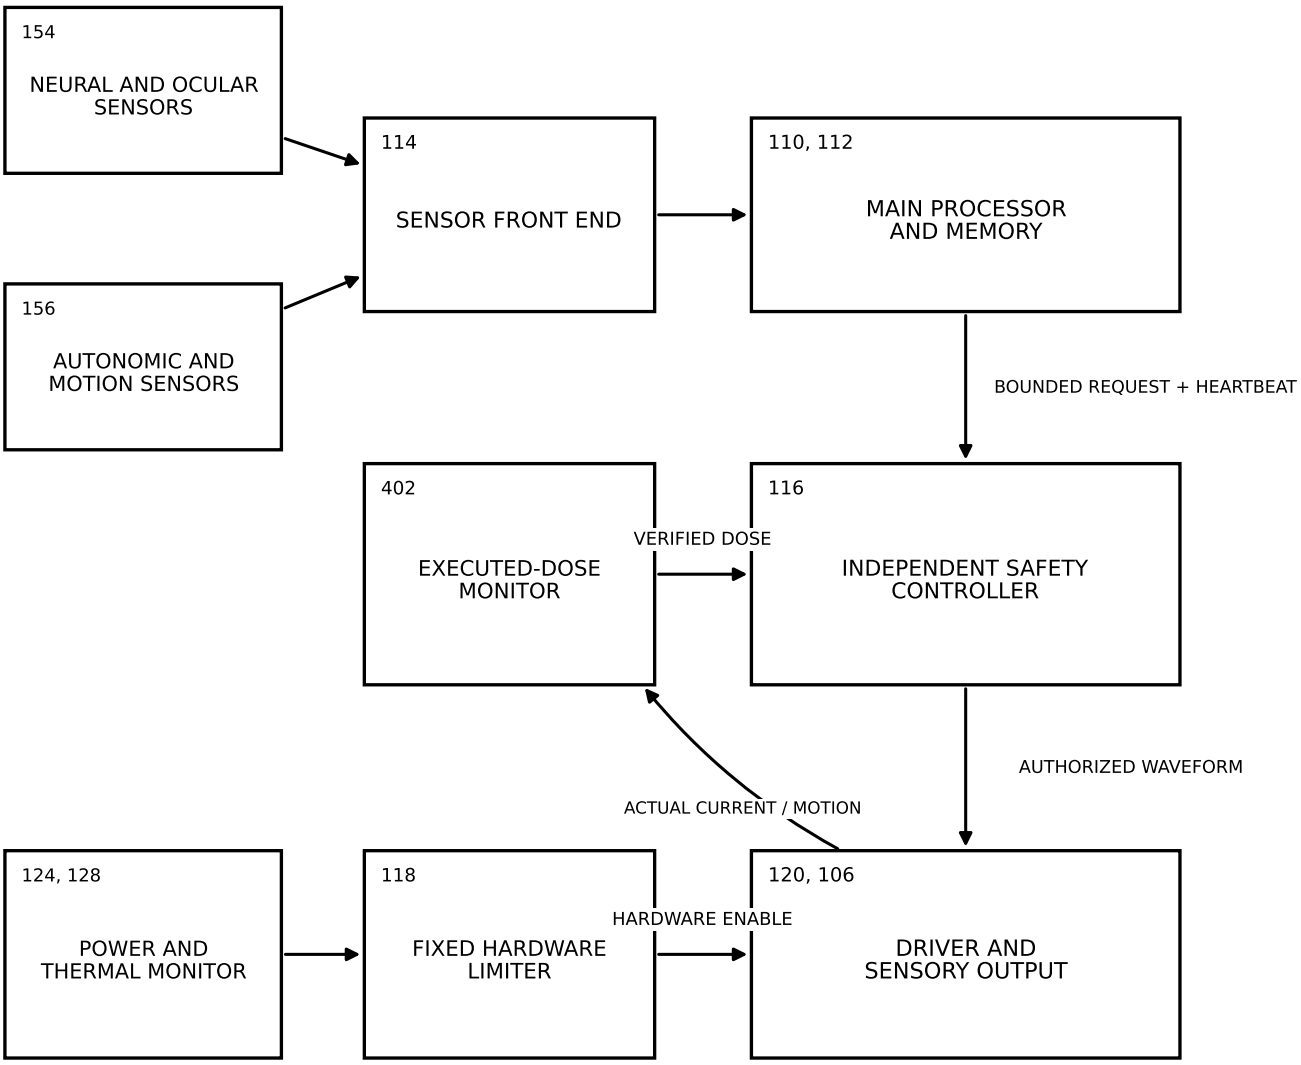

FIG. 2

FIG. 3. End-to-end closed-loop data flow

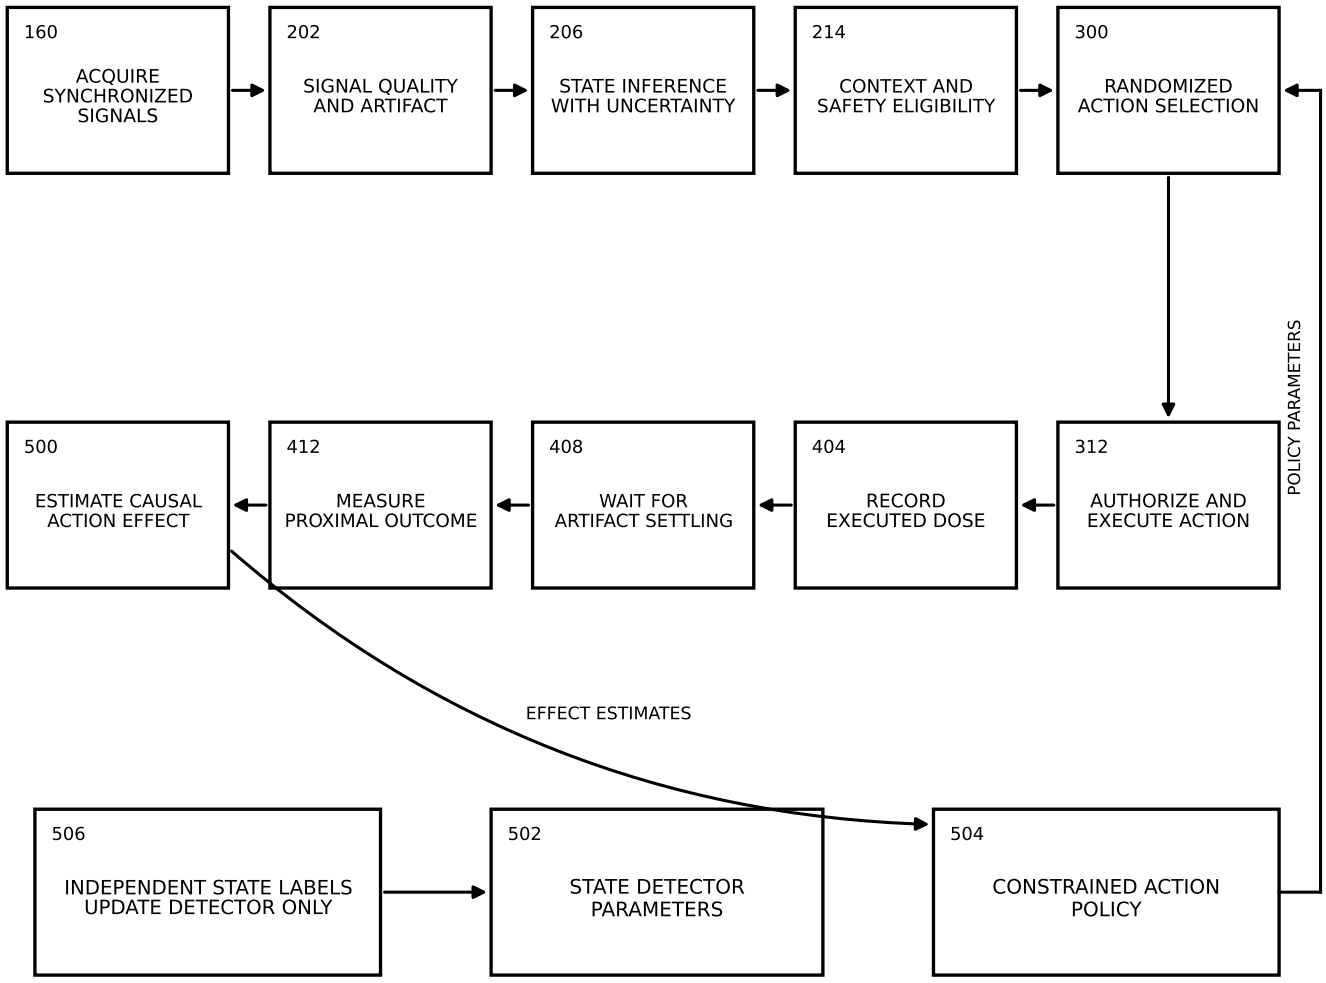

ORDINARY TREATMENT OUTCOMES DO NOT RELABEL THE STATE DETECTOR

FIG. 3

FIG. 4. Separation of state, causal-effect, and action-policy models

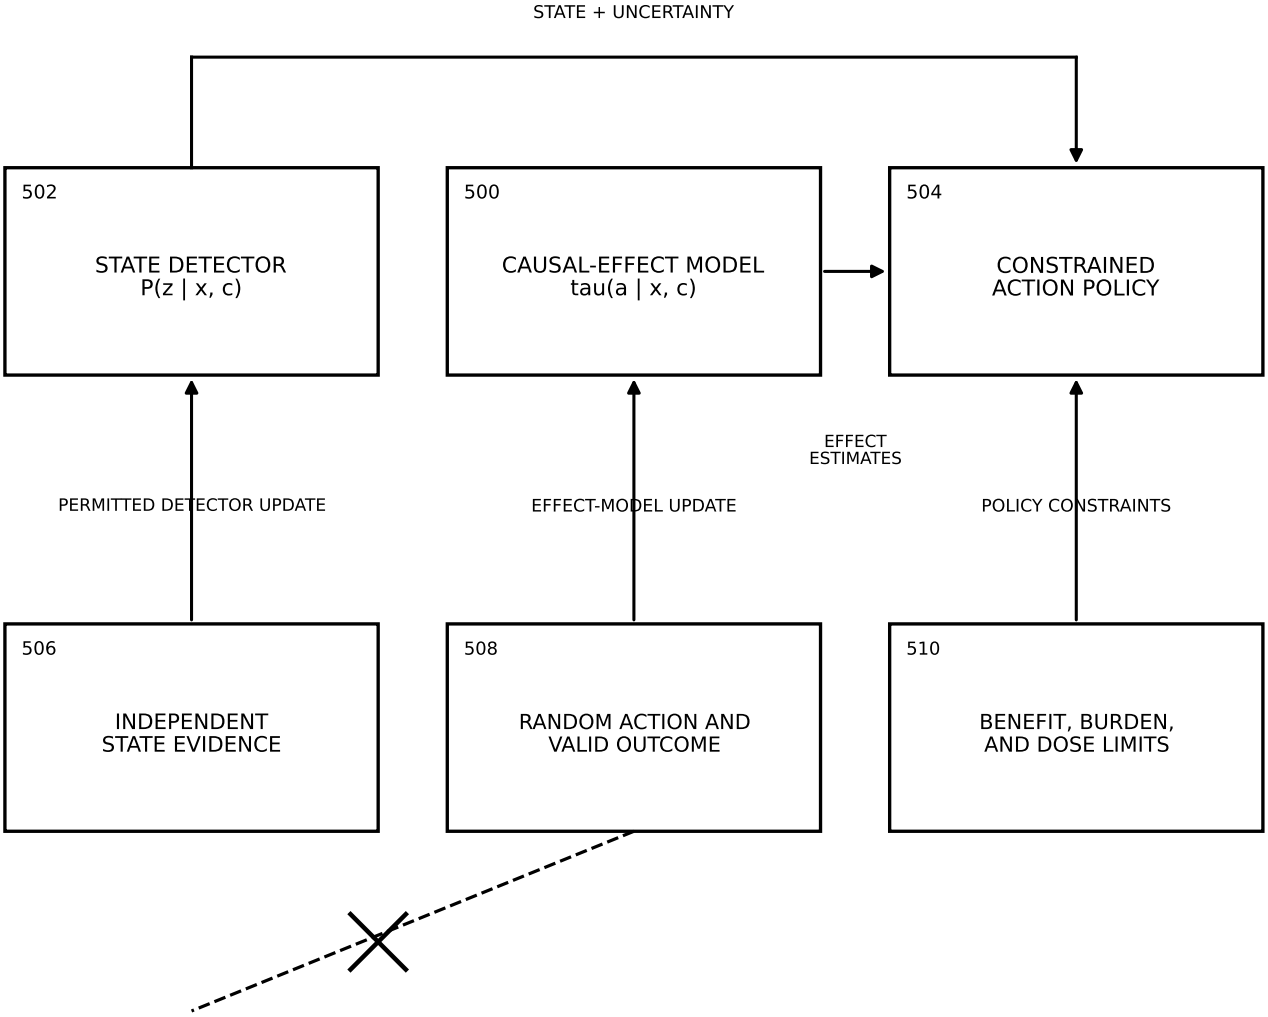

512 RESPONSE TO AN ORDINARY TREATMENT DOES NOT RELABEL THE DETECTED STATE

FIG. 4

FIG. 5. Enrollment and validation workflow

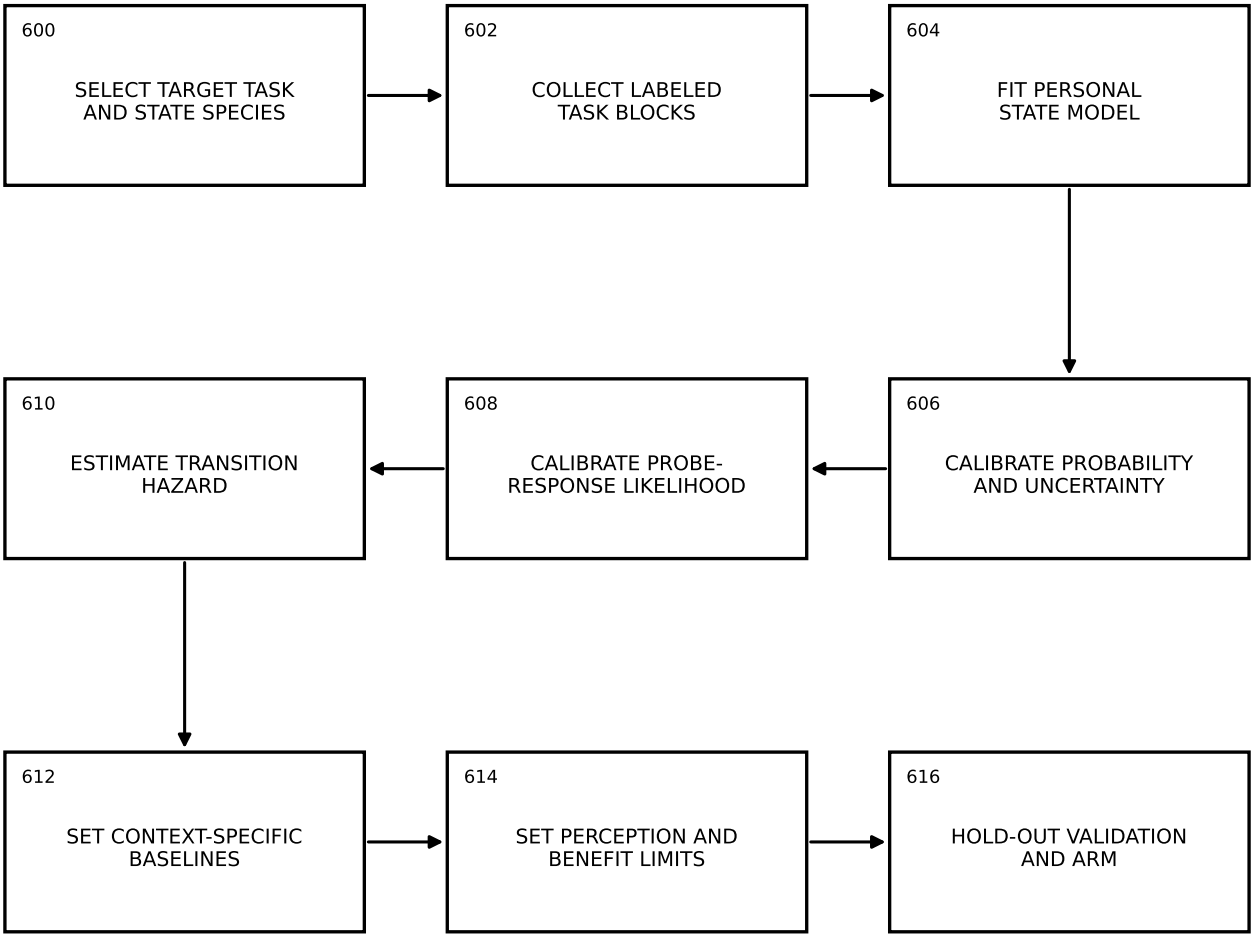

FIG. 5

FIG. 6. Synchronized sensing and artifact qualification

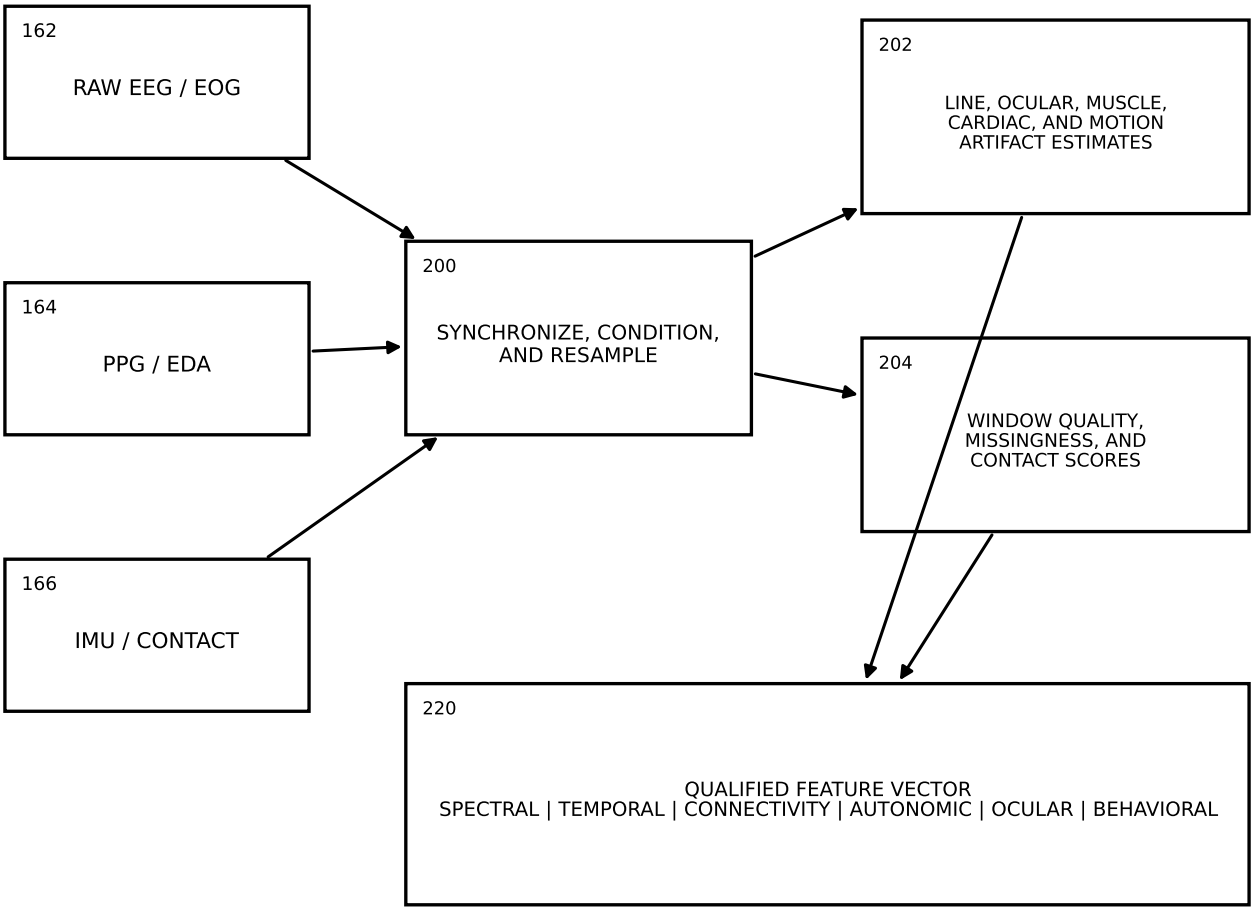

FIG. 6

FIG. 7. Uncertainty-aware state and transition-hazard inference

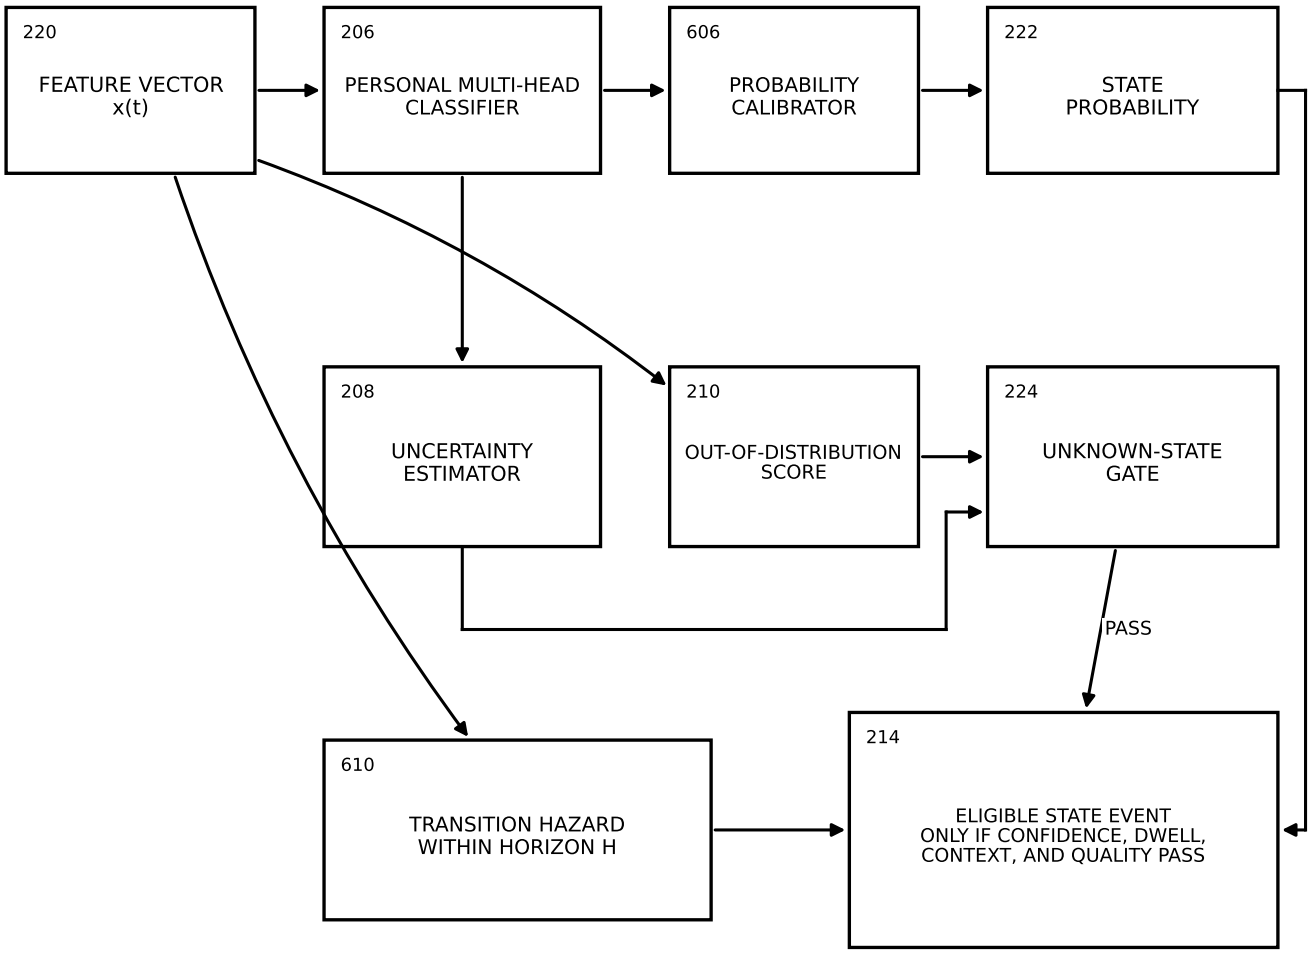

FIG. 7

FIG. 8. Eligibility gating and randomized action selection

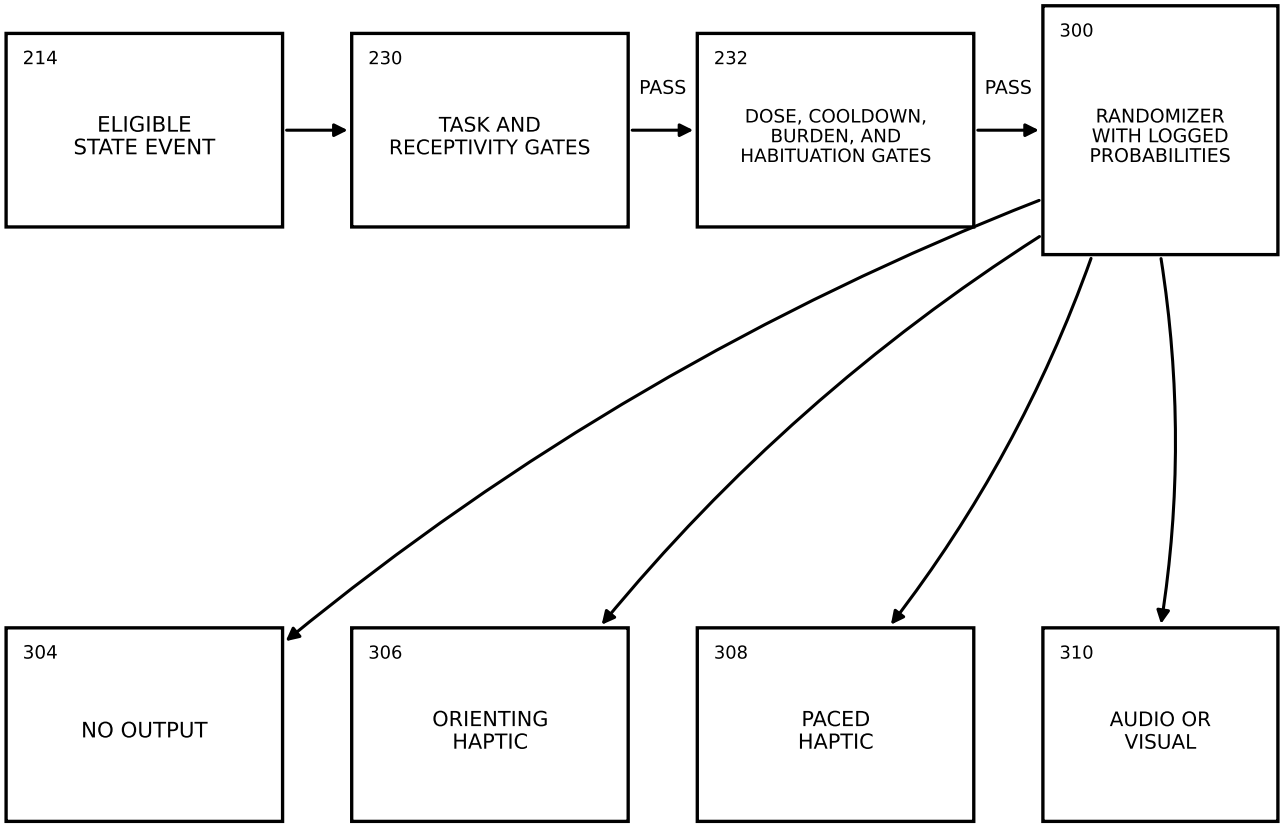

ONLY AVAILABLE, SAFE ACTIONS ARE RANDOMIZED; NO OUTPUT REMAINS AN ACTION

FIG. 8

FIG. 9. Logged randomized assignments including no output

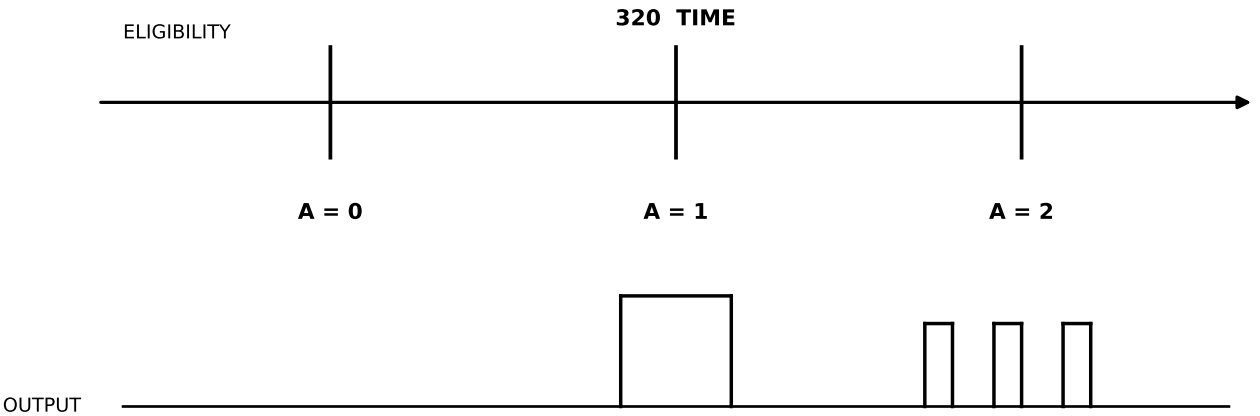

302

EVENT ID | CONTEXT | P(A=0) | P(A=1) | P(A=2) | SELECTED ACTION | MODEL VERSION

322 RANDOMIZED NO-OUTPUT EVENTS REMAIN ELIGIBLE OBSERVATIONS

SAFETY-REJECTED OR UNSAFE EVENTS ARE CODED SEPARATELY AND ARE NOT CONTROLS

FIG. 9

FIG. 10. Requested versus executed dose provenance

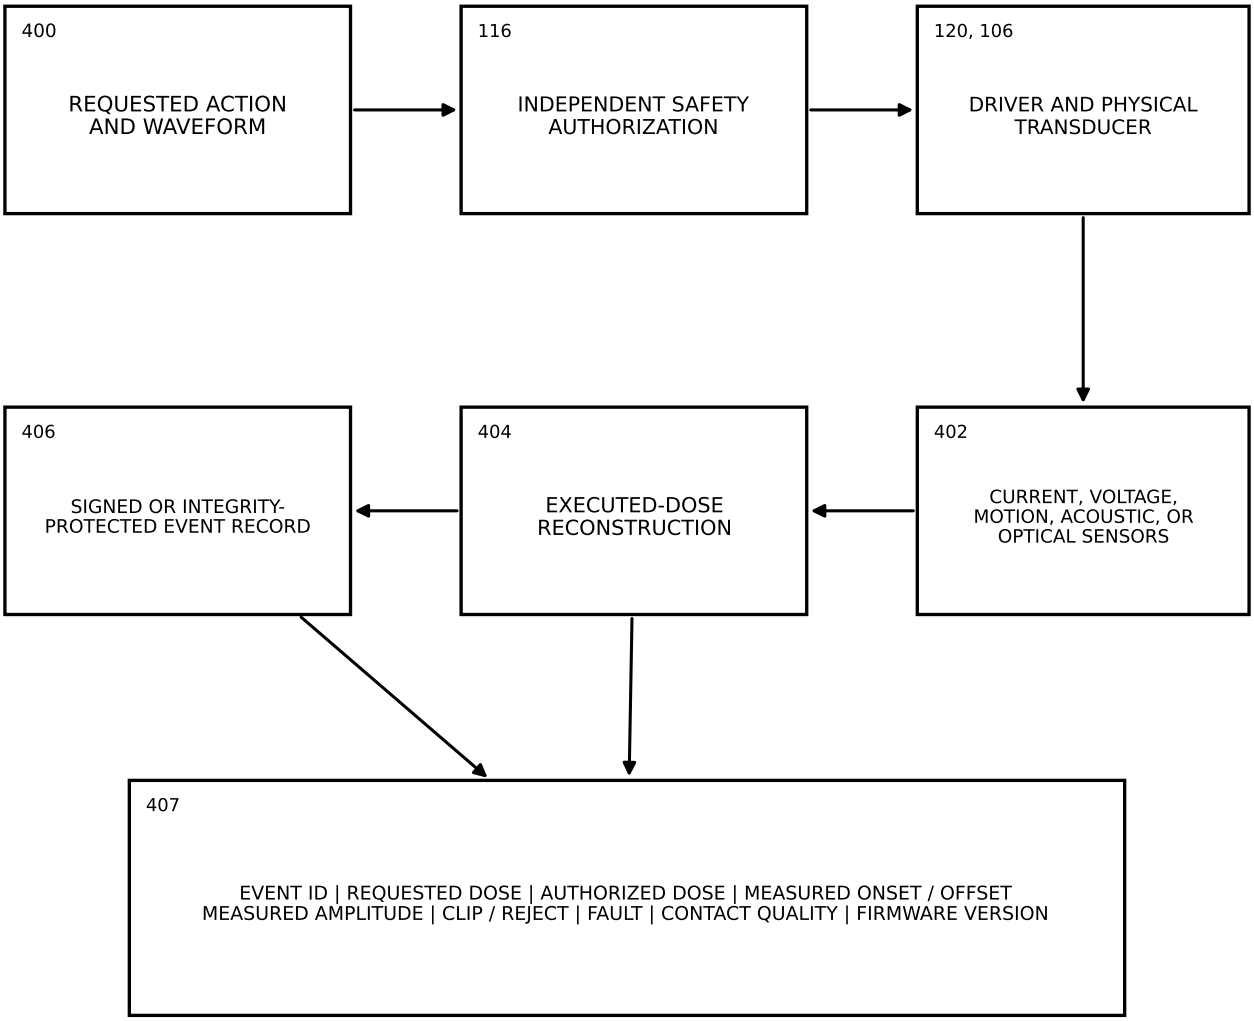

FIG. 10

FIG. 11. Adaptive artifact settling and valid outcome window

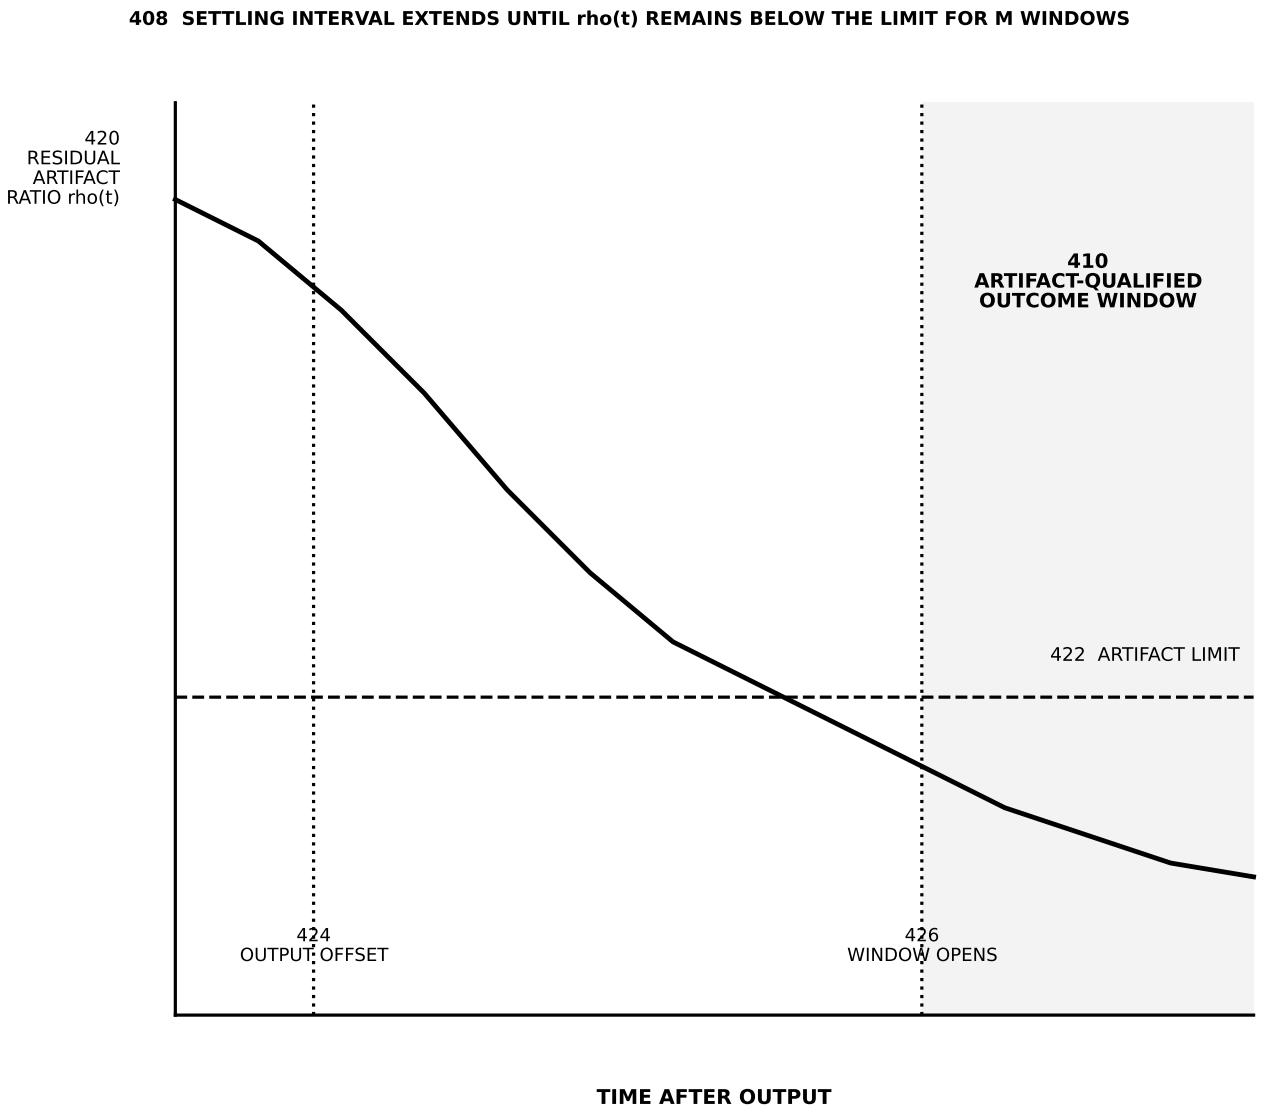

FIG. 11

FIG. 12. Causal-effect estimation from valid randomized events

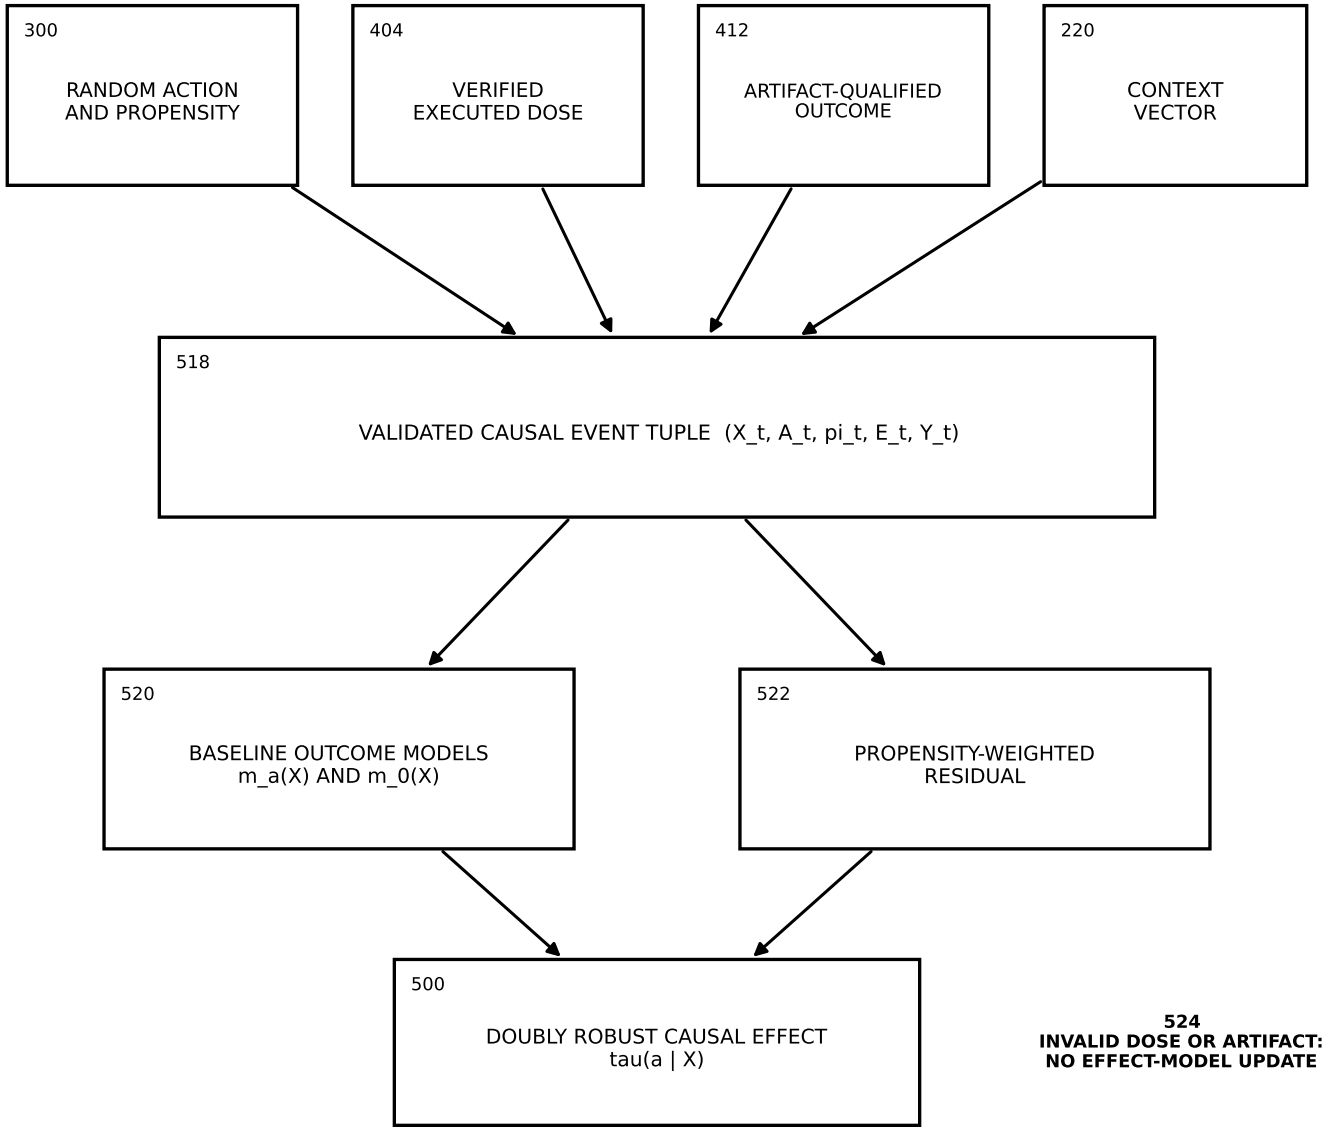

FIG. 12

FIG. 13. Validated active diagnostic-probe pathway

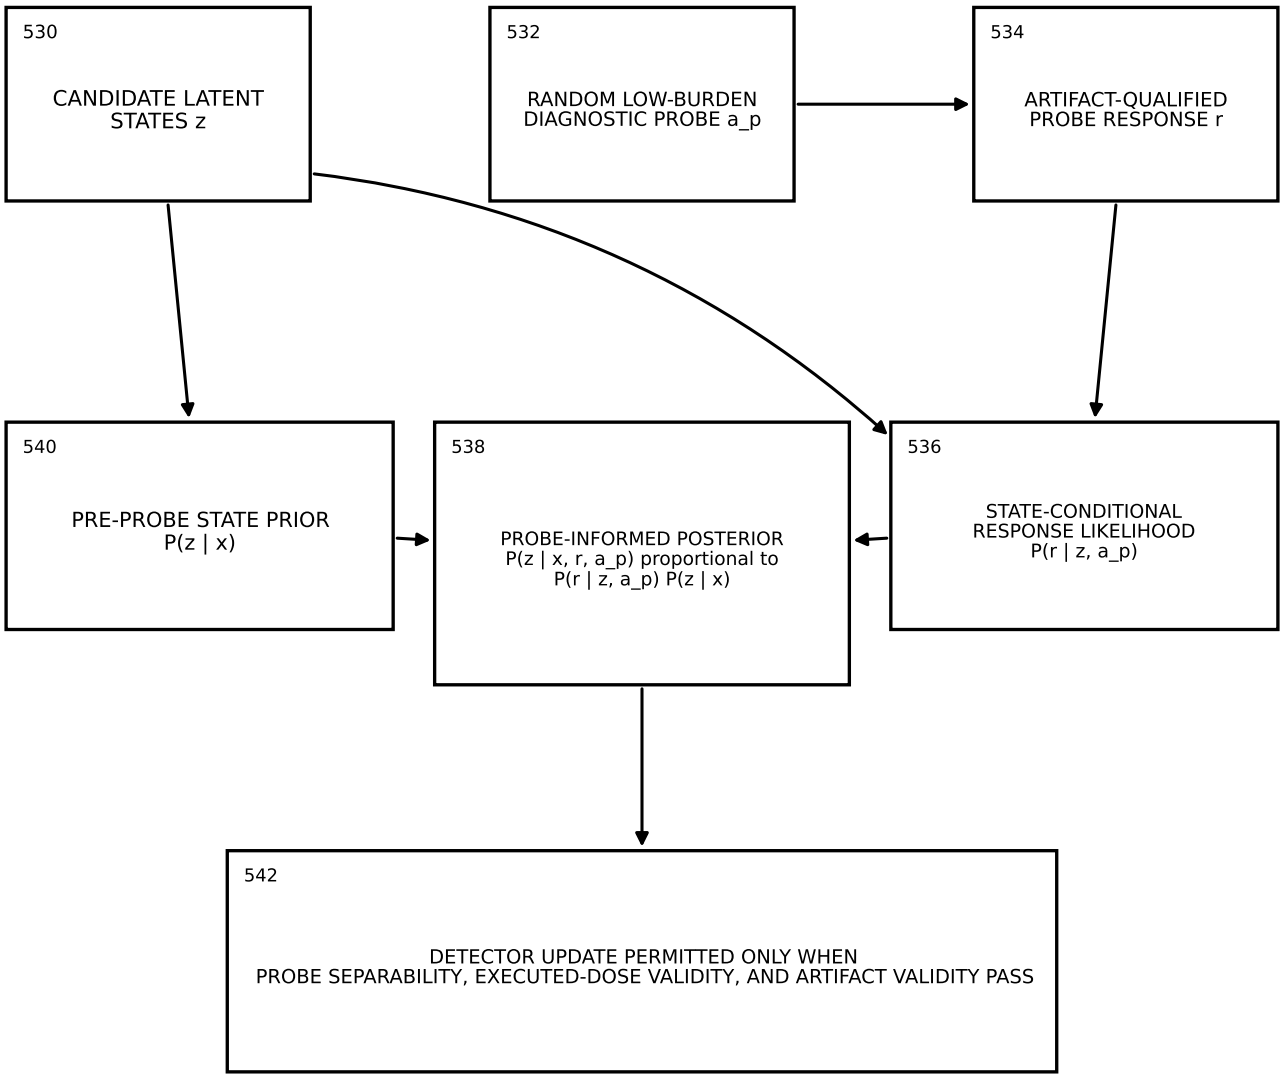

FIG. 13

FIG. 14. Bounded learning, drift detection, and rollback

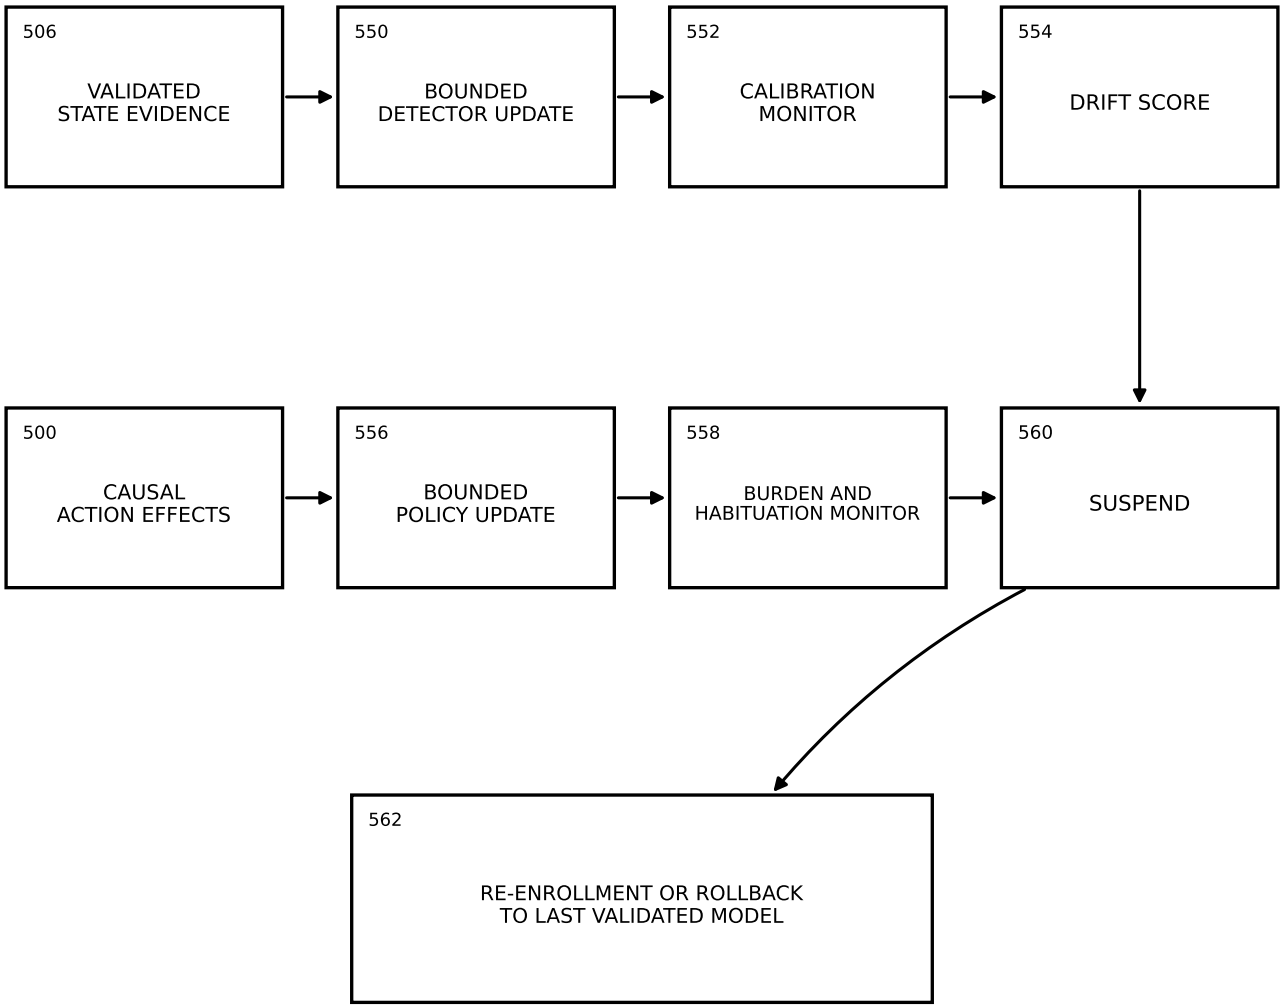

FIG. 14

FIG. 15. Independent safety-controller state machine

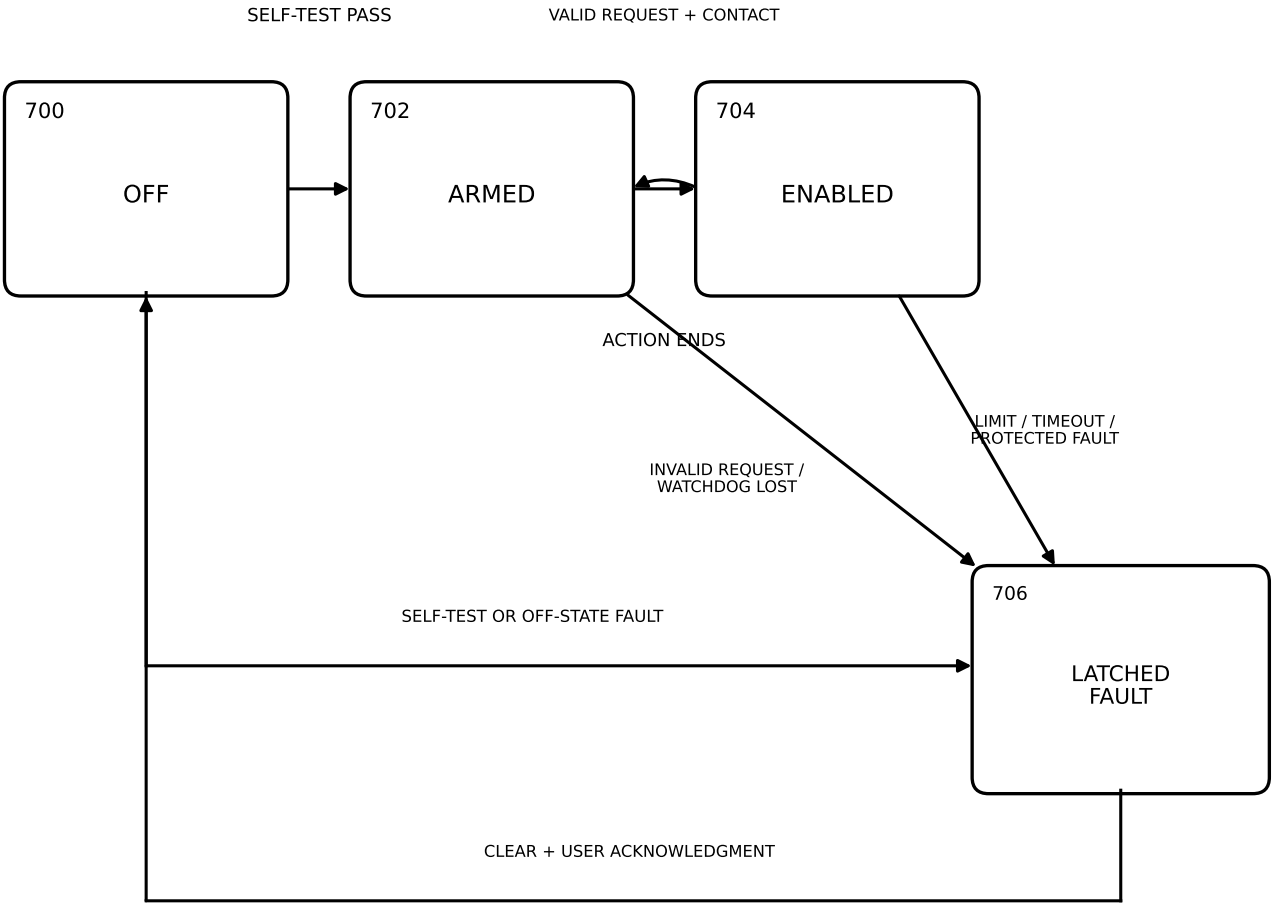

708 OUTPUT ENABLE IS DEASSERTED IN OFF AND LATCHED FAULT, AND UPON LOSS OF HEARTBEAT OR ANY PROTECTED SAFETY CONDITION

FIG. 15

FIG. 16. Fail-low output circuit and executed-dose feedback

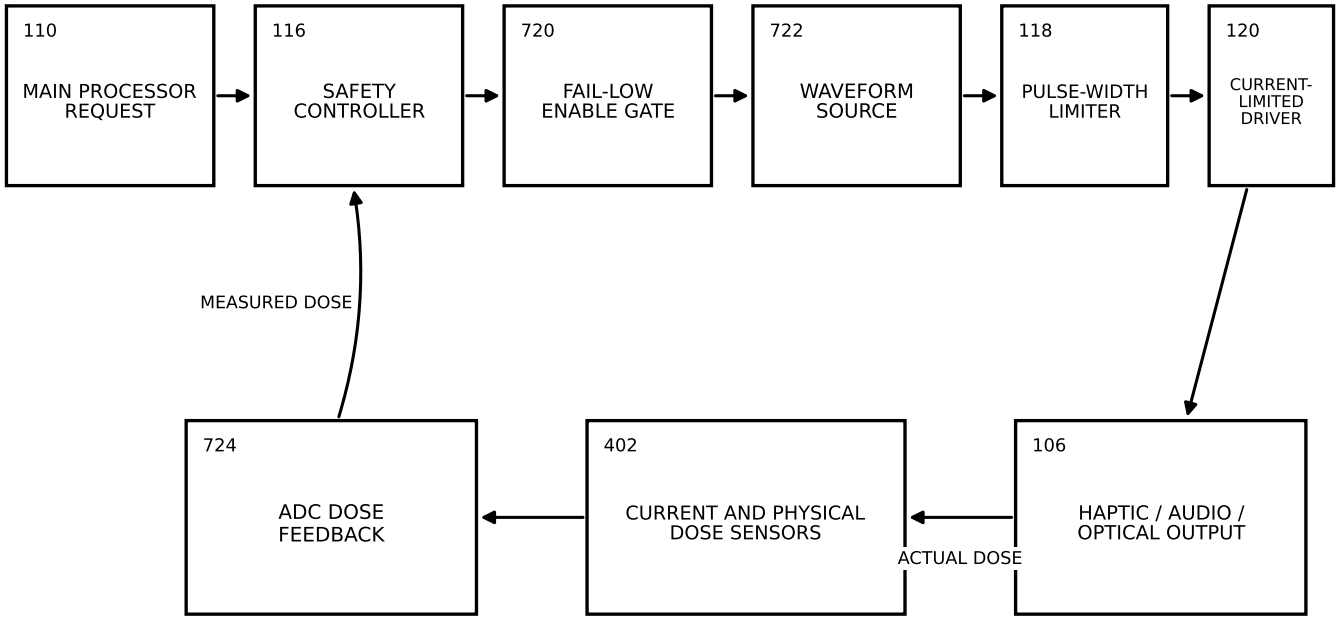

THE OUTPUT REQUIRES SIMULTANEOUS VALID REQUEST, WATCHDOG, SAFETY, AND HARDWARE-LIMITER CONDITIONS; ANY FAILURE FORCES OUTPUT OFF

FIG. 16

FIG. 17. Multi-device global dose coordination

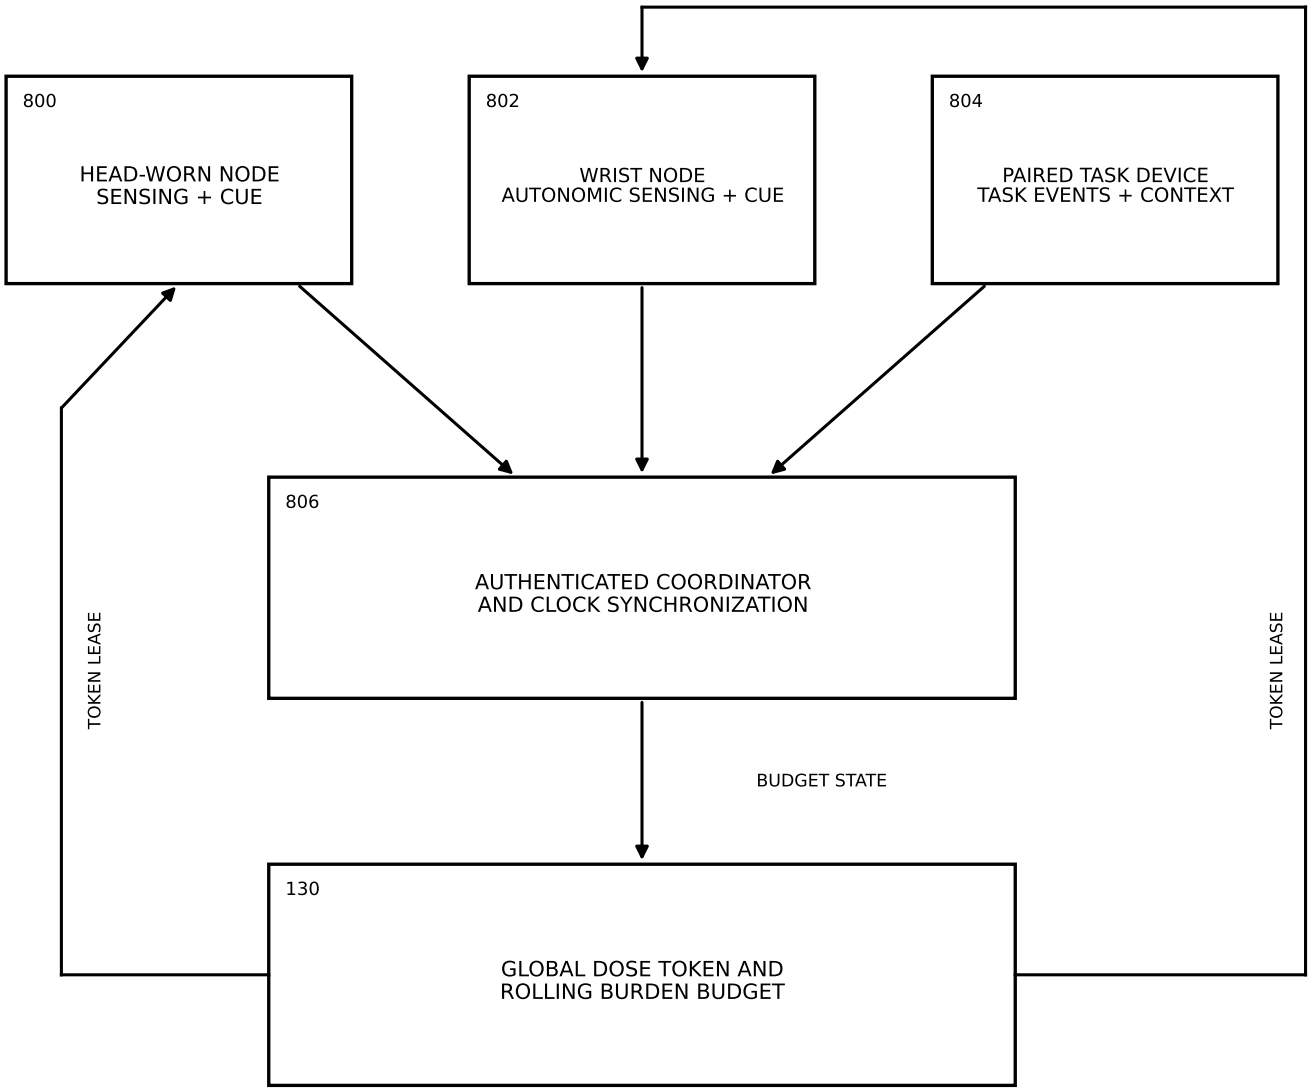

EACH OUTPUT-CAPABLE NODE ENFORCES ITS OWN LOCAL PROTECTED SAFETY ENVELOPE

FIG. 17

FIG. 18. Complete operating method

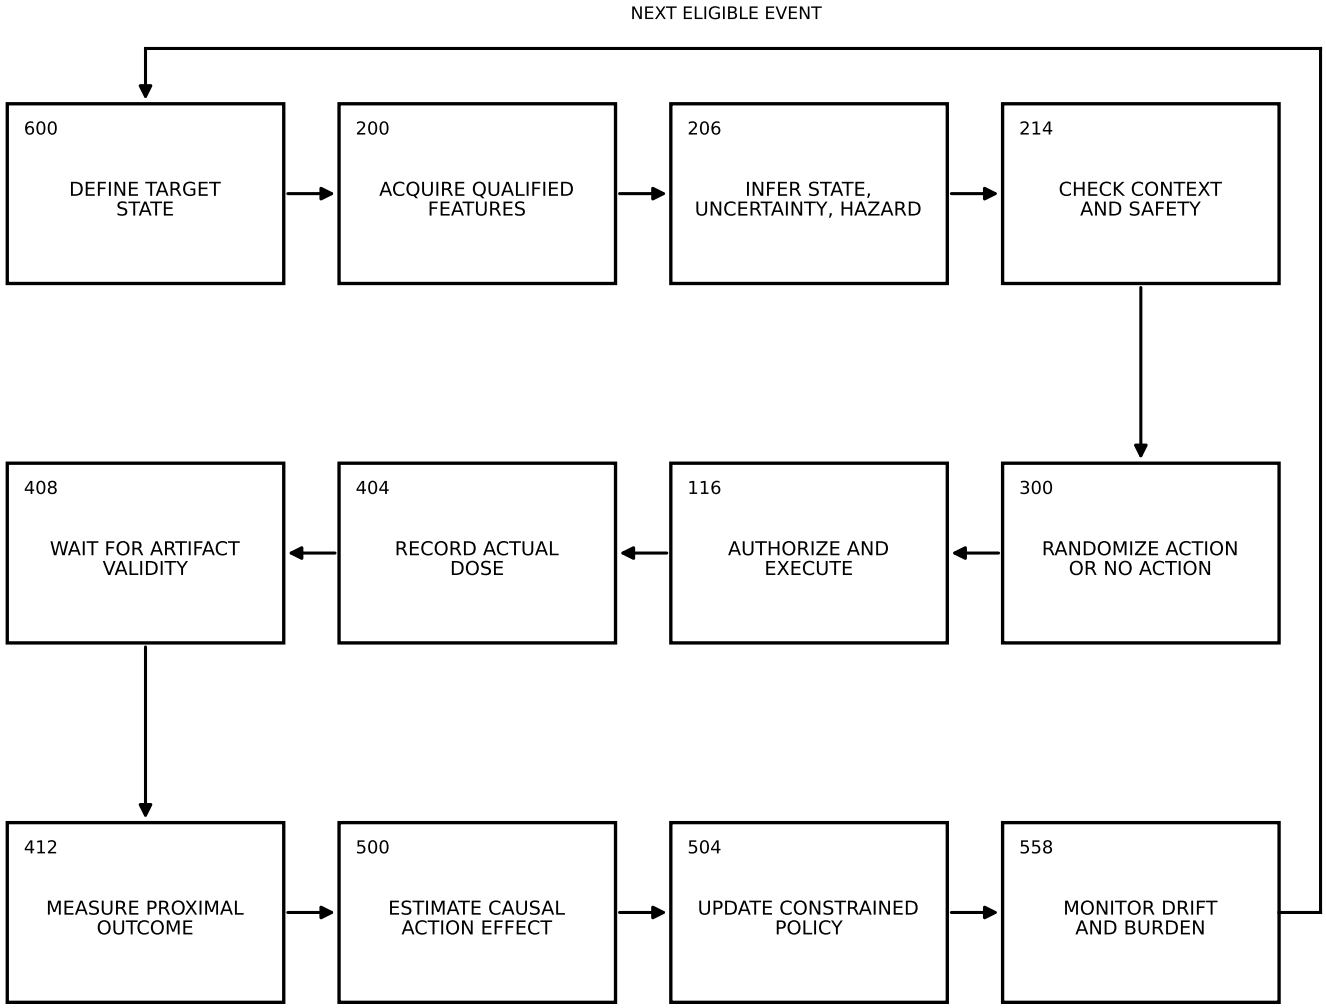

FIG. 18