A PCLC-Based Coherence Diagnosis–Intervention Device

Design Specification for Build, and Integration Point into SWARP

J. Konstapel, Leiden, 20-7-2026.

A working design document — intended to be handed to an engineer or builder — that translates the PCLC Failure-Mode Taxonomy and the Kuramoto order parameter r(t) into a two-stage device architecture: a diagnosis layer that classifies which failure Type is present, and an intervention layer that delivers a corrective audio signal via headphones. This document is deliberately conservative about what is buildable today versus what remains a design placeholder.

Part 0 — What This Document Establishes, and What It Does Not

Consistent with the accounting convention in the PCLC case-studies paper, this section states the limits up front rather than at the end.

- **This is a concept design, not a certified medical device.** Nothing here has regulatory clearance (CE/MDR, FDA). A device that senses physiological signals and delivers a corrective stimulus to a person is a medical device under EU and US law, and would require clinical validation and regulatory approval before use on patients — including healthy volunteers in a trial.
- **The diagnosis layer is the weak link.** Of the eleven case studies in the taxonomy, only a few use signals that a consumer-grade sensor (phone, wearable) can plausibly capture today. Most (phase-singularity mapping in AFib, basal ganglia LFP in Parkinson's, receptor antibody titers in myasthenia gravis) require clinical-grade or invasive instrumentation. This document flags, case by case, which are realistic and which are not.
- **The intervention layer is unvalidated.** No claim is made that audio entrainment corrects any of the five failure types. This is a hypothesis to be tested, not a known mechanism.
- **Type II carries a specific, non-hypothetical risk that must gate the design.** Musicogenic and audiogenic epilepsy — seizures triggered by sound, including rhythmic or patterned auditory stimuli — is clinically documented. A device intended to *suppress* pathological hypersynchrony in a Type II condition could, if the entrainment logic is wrong, instead *induce* it. This is not a theoretical edge case to mention in a footnote; it is a reason Type II should not be the first condition this device is built or tested for, and any Type II application should not proceed without a clinical neurologist involved in the design from the start.

Part I — System Architecture

Two-stage closed loop, matching the sensor–controller–effector structure the taxonomy itself uses to classify disease:

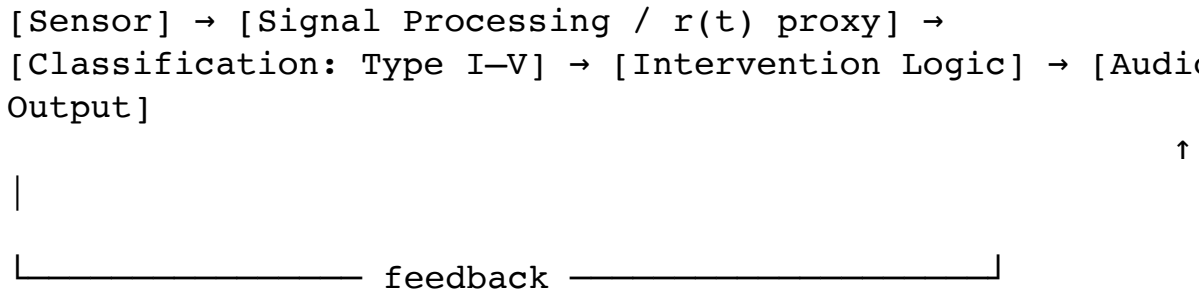

This is a **closed loop**, not open-loop playback: the intervention should be continuously informed by the ongoing sensor read, so the system can detect whether the coherence measure is moving toward or away from target and adjust — the same principle Part IV of "Coherence Across Scales" attributes to adaptive DBS ("steerability" as a controllability property, restored in closed loop rather than by a fixed suppression signal).

Part II — Diagnosis Layer

PCLC Type	Signal proxy used in the case-studies	Consumer-sensor feasibility today	Notes
I — cohere nce	HRV (sepsis); phase-singularity mapping (AFib)	**HRV: yes**, via PPG (phone camera or wrist wearable). AFib mapping: no, requires ECG-grade or catheter data.	HRV is the realistic starting point for this device.
II — patholo gical hyper-	Beta-band (13–30 Hz) power (Parkinson's); EEG hypersynchronizatio	Consumer EEG headbands exist but signal quality is far below clinical EEG; beta-band basal ganglia coupling specifically requires implanted electrodes (DBS lead) — not	Diagnosis alone is hard here; see Part 0 risk note before intervention is even
III — referen ce	Baroreceptor pressure-threshold shift; beta-cell	No direct consumer proxy. Blood pressure trend over time is a distant, indirect signal.	Not a realistic near-term target for this device.
IV — node destruc	Beta-cell mass / autoantibody status	Not measurable by any sensor in this class — this is a lab diagnostic.	Out of scope by definition: there is no loop left for a
V — compo nent-	Nerve conduction (sensor); ADH/ osmolality	Not measurable by phone/wearable sensors.	Out of scope for this device class.

Conclusion of Part II: Type I via HRV is the only cell in this table where diagnosis is both realistic today and low-risk. This is the recommended starting point for a first buildable prototype — not because the other Types are uninteresting, but because they either require instrumentation this device class doesn't have, or carry the Type II risk flagged in Part 0.

Part III — Classification Logic (Type I, HRV-based)

1. Continuous PPG capture → beat-to-beat interval series (RR intervals).
2. Compute standard HRV metrics (RMSSD, SDNN) over a rolling window as the empirical proxy for coupling/autonomic coherence — consistent with how the case-studies paper treats HRV as a Type I proxy for sepsis, not a literal Kuramoto r(t).

3. Compare against the individual's own baseline (not a population norm — inter-individual HRV variance is large), flagging sustained departure below baseline as a coherence-deficit state.
4. This is a **within-person trend detector**, not a diagnostic instrument. It should not be framed to a user as detecting or diagnosing a named disease.

Part IV — Intervention Layer (Type I only, for the first build)

- **Goal, per the taxonomy's own treatment logic (Part III of the case-studies paper):** Type I calls for restoring lost coupling, not for forcing a fixed rhythm.
- **Mechanism proposed:** binaural or isochronic tones at a frequency within the range associated with parasympathetic/vagal engagement (commonly studied: slow-paced breathing-entrainment around 0.1 Hz breath rate, translated into a paced audio cue rather than a raw binaural carrier frequency — paced-breathing audio has substantially more supporting evidence than binaural-beat carrier frequencies specifically).
- **Closed-loop behavior:** audio pacing continues only while HRV trend remains below the individual's baseline; as HRV recovers toward baseline, intervention intensity/duration should taper — this operationalizes "steerability" as an actual stop condition, not a fixed-duration session.
- **Explicitly out of scope for this document:** Type II intervention logic. See Part 0.

Part V — Integration Point into SWARP

- The **classification step (Part III)** is the natural home for SWARP's KAYS coherence engine — it is, functionally, a coherence-scoring function, which is what KAYS already exists to do.
- **AYYA360 profiling** is the natural home for the "individual baseline" requirement in Part III.3 — the device needs a per-person reference, not a population norm, and that's an individual-profiling function.
- **QUO?** is a plausible interface layer for surfacing the trend to the person, if the device is meant to inform rather than only auto-intervene.
- This mapping is a **proposed integration point**, not yet a specification of how KAYS/AYYA360/QUO? would need to be extended to actually do this — that would be a follow-on design task.

Part VI — What a Builder Needs Before Starting

Minimum open items this document does not resolve, and that block an actual build:

1. **Baseline-setting protocol** — how many days of data establish "this person's normal HRV," and how is that re-validated over time.
2. **Audio intervention parameters** — exact frequency, waveform, and volume envelope are not specified here; this requires either literature review of the paced-breathing-audio evidence base specifically, or a pilot.
3. **Hardware choice** — which PPG sensor (phone camera vs. dedicated wearable) meets the sampling-rate and noise requirements for RMSSD/SDNN computation; phone-camera PPG is noisier than wearable PPG and this tradeoff hasn't been decided here.
4. **Regulatory pathway** — even a "wellness" framing (avoiding disease-diagnosis claims) has boundaries under EU/US law once the device claims to detect and respond to a physiological state; this needs a regulatory consult before build, not after.

Falsifiable Prediction for This Design

If built as specified in Parts III–IV: HRV trend (RMSSD) during and after audio-paced sessions should show measurable improvement toward the individual's baseline, compared to a no-audio control period, within the same person. No systematic difference between session and control falsifies the design's core mechanism.

This document extends the PCLC Failure-Mode Taxonomy (Konstapel, 2026) and "Coherence Across Scales" (Konstapel, 2026) into a buildable first-prototype scope. It deliberately narrows to Type I / HRV as the only combination judged both technically realistic and safe to prototype without a clinical specialist embedded in the design team from the outset.