

---

# Position: Beyond Prediction: Toward Verifiable Physiological Waveform Reasoning with Foundation Models and Agentic LLMs

---

Anonymous Authors<sup>1</sup>

## Abstract

Physiological waveforms (e.g., ECG, PPG, EEG) encode clinically meaningful information in fine-grained morphology, precise timing, and cross-channel dynamics, yet most machine learning systems still treat them as generic time series and optimize end-to-end prediction. In this position paper, **we argue for verifiable physiological waveform reasoning: extracting localized, measurable signal evidence from raw signals, interpreting that evidence into physiological semantics, and supporting clinically grounded decisions.** Waveform reasoning is challenging due to acquisition heterogeneity, signal fidelity, complex semantics and cross-channel coupled dynamics. We analyze why existing model families remain insufficient: physiological foundation models learn strong perceptual representations but remain weak at verifiable reasoning, while LLM-based adaptations have limited waveform understanding. To bridge this gap, **we advocate verifiable, closed-loop systems that unify waveform semantics with language intelligence.** Concretely, we propose a dual-process architecture that System 1 aligns physiological waveforms with language, and System 2 provides agentic reasoning via a Plan–Act–Verify loop, together enabling verifiable physiological waveform reasoning. And we propose evaluations beyond accuracy, emphasizing traceability, replayability, counterfactual robustness, and calibrated abstention.

## 1. Introduction

Physiological waveforms such as electrocardiograms (ECG), photoplethysmograms (PPG), and electroencephalograms (EEG) are high-fidelity, time-resolved measurements of un-

---

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

derlying biological dynamics. Unlike many time-series forecasting benchmarks where performance can often be driven by coarse trends, seasonality, or global summary statistics (Chang et al., 2025a; Ye et al., 2025), the clinical meaning of physiological waveforms is concentrated in fine-grained morphology, precise timing, and structured dependencies across channels. Localized events (e.g., QRS onsets/offsets, dirotic notches) support standardized measurements and guideline-driven interpretation (Clifford et al., 2012; Orphanidou et al., 2014b; Wang et al., 2025a). In practice, clinicians do not “read” a waveform by a single global score; they delineate events, measure intervals and amplitudes, assess signal quality, and reconcile inconsistencies across leads or modalities before reaching guideline-constrained decisions (Shcherbina et al., 2017; Bent et al., 2020).

Recent years have seen rapid progress in applying deep learning to waveform modeling, achieving strong performance in arrhythmia detection and broader ECG interpretation (Rajpurkar et al., 2017; Hannun et al., 2019; Ribeiro et al., 2020; Attia et al., 2019; Ismail Fawaz et al., 2019). More recently, physiological foundation models (PhysioFMs) promise transferable representations and scalable pretraining regimes (Wiggins & Tejani, 2022; Mehari & Strodthoff, 2022; Li et al., 2025b). In parallel, large language model (LLM)-centric adaptations have emerged that use language as an interface for clinical tasks, including explanation generation, and guideline-aware summarization. Despite these advances, much of the literature still emphasizes end-to-end prediction. This clinician workflow also highlights a failure mode that end-to-end prediction can obscure: undetected wrongness under acquisition artifacts and heterogeneity. A motion-corrupted PPG segment can mimic irregular rhythm; a baseline wander can distort ST segments; missing or swapped ECG leads can yield plausible-looking signals that nevertheless invalidate downstream measurements. A system that outputs a prediction without exposing what evidence it relied on is difficult to audit.

**We argue that physiological signal AI should be framed as physiological waveform reasoning: extracting localized, verifiable evidence from raw signals, translating that evidence into physiological semantics, and producing clinically grounded decisions rather than only end-**

**to-end predictions.** Here, “verifiable” is an operational requirement that reasoning be entailed by measurable evidence objects. The decision should be creditable only to the extent that it can be reconstructed from recorded evidence objects and the procedures that produced them, yielding traceability and replayability rather than post-hoc narrative.

This reframing also clarifies why current model families remain insufficient. PhysioFMs excel as perceptual backbones, but their typical outputs do not constitute evidence objects with provenance; they provide limited mechanisms for unit-consistent measurements, explicit physiological semantics, or guideline-constrained decision logic (Rudin, 2019; Li et al., 2025b). Conversely, LLM-centric adaptations can improve the linguistic form of reasoning (e.g., producing fluent explanations, summaries, or differential diagnoses) but often rely on lossy interfaces to high-frequency morphology, weak grounding to raw signals, and rationales that are not guaranteed to be logically supported by the waveform (Schick et al., 2023; Shinn et al., 2023; Cao et al., 2024; Goswami et al., 2024; Jin et al., 2024b).

**We therefore advocate verifiable, closed-loop systems that unify waveform semantics with language intelligence.** Concretely, we propose a dual-process that *System 1* aligns the model with physiological waveforms and language, enabling faithful waveform understanding, while *System 2* performs agentic reasoning via a Plan–Act–Verify loop that decomposes tasks, requests missing evidence, invokes deterministic measurement and validation tools, checks cross-view consistency, and abstains or escalates when evidence quality or constraints do not support a safe conclusion (Guo et al., 2017; Ovadia et al., 2019; Geifman & El-Yaniv, 2019). Accordingly, evaluations should move beyond accuracy to score whether outputs are traceable to localized evidence, replayable under logged procedures, robust to nuisance perturbations yet sensitive to clinically meaningful counterfactual changes, and uncertainty-aware through calibrated abstention or escalation.

**Contributions.** The main contributions are as follows: ❶ We reframe physiological signal AI as *verifiable physiological waveform reasoning*: extracting evidence from raw signals, mapping evidence to physiological semantics, and supporting clinical decisions. ❷ We analyze the key challenges and why current model families remain insufficient. PhysioFMs rarely expose auditable verification objects, while LLM-centric adaptations often rely on lossy waveform grounding. ❸ We propose a dual-process, closed-loop blueprint in which *System 1* aligns physiological waveforms with language, while *System 2* performs agentic Plan–Act–Verify reasoning for meaningful physiological waveform reasoning. ❹ We advocate evaluations beyond end-to-end prediction, emphasizing evidence traceability, replayability, counterfactual robustness, and uncertainty-aware decisions.

## 2. Beyond Prediction: Toward Verifiable Physiological Waveform Reasoning

### 2.1. What are Physiological Waveforms

Physiological waveforms encompass the direct measurement of the body’s electrical activity or hemodynamic responses. We define physiological waveforms as high-fidelity temporal projections of continuous biological processes. Unlike generic time-series data (e.g., financial stocks or weather metrics) where analysis often centers on global trends, seasonality, or statistical distribution (Chang et al., 2025a; Ye et al., 2025; Chang et al., 2025b), physiological waveforms are characterized by precise morphological semantics and mechanistic coupling. The information is dense and encoded not just in the value at time  $t$ , but in the morphology, phase relationships, and quasi-periodic structure. Formally, we represent a physiological waveform recording as a multivariate tensor  $\mathbf{X} \in \mathbb{R}^{C \times T}$ , where  $T$  represents the discrete sampling of a continuous biological state and  $C$  denotes the spatial or modal channels (e.g., 12 leads of ECG). We use physiological waveforms to refer to widely used sensing modalities such as ECG, PPG, EEG, EMG, and PCG. For completeness, we summarize their sensing principles and typical characteristics in Appendix A.

### 2.2. Verifiable Physiological Waveform Reasoning

We define *physiological waveform reasoning* as the capability to transform raw biosignals into actionable clinical logic through a structured inference process grounded in waveform evidence. Unlike standard end-to-end prediction, which may rely on opaque correlations (Ismail Fawaz et al., 2019; Rudin, 2019), reasoning requires explicit intermediate products that connect low-level signal dynamics to high-level physiological concepts under established clinical principles (Rajpurkar et al., 2022; Wagner & Strauss, 2013).

We treat *verifiability* as an operational contract. Specifically, a system should emit *verification objects*—signal-quality summaries, localized events with explicit time/lead indices, and unit-consistent measurements defined by reproducible windows and procedures—so that intermediate claims and final decisions are acceptable only if they can be re-derived by replayable checks. This yields an auditable interface where conclusions remain grounded in localized, measurable waveform evidence rather than post-hoc narrative. We organize physiological waveform reasoning along two axes (Fig. 1): *Reasoning Structure* and *Reasoning Objective*

**Reasoning Structure.** We define *Reasoning Structure* as the topology by which a system composes intermediate states into a conclusion (Wei et al., 2022). We summarize three topologies: *Direct* reasoning, *Linear-chain* reasoning, and *Branch-structured* reasoning. In practice, many systems either adopt a single topology or are hybrids that combine

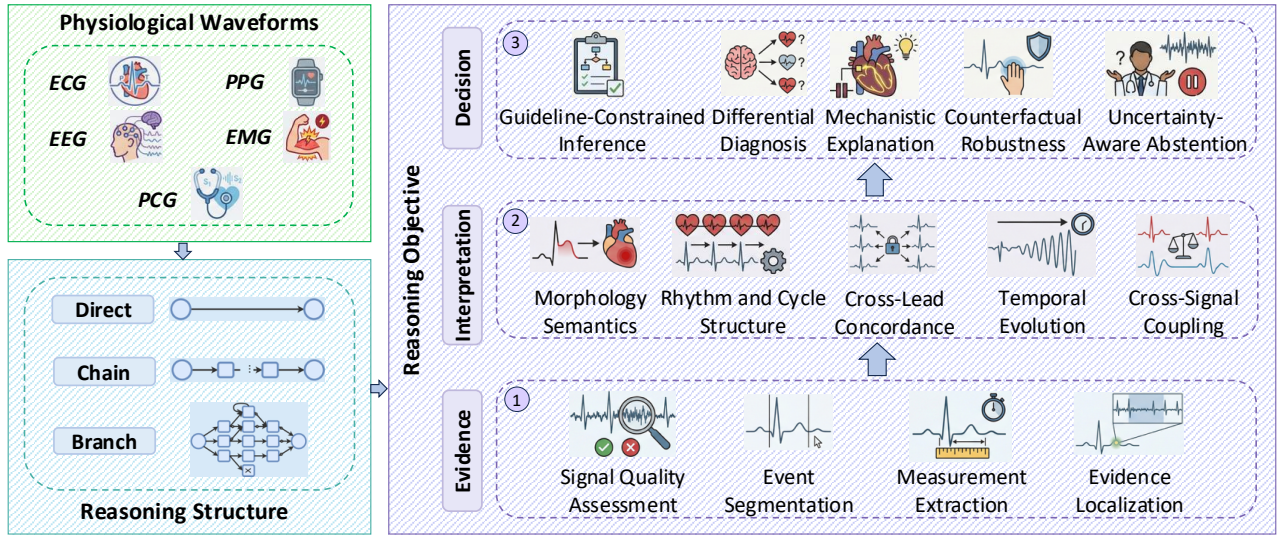


Figure 1. **Verifiable physiological waveform reasoning.** We organize reasoning along two axes: (1) **Left: Reasoning Structure** describes the inference topology: *Direct*, *Linear-Chain*, or *Branch-Structured*. (2) **Right: Reasoning Objective** specifies the target level: Level 1 *Evidence* (e.g., segmentation/measurement), Level 2 *Interpretation* (e.g., semantics), and Level 3 *Decision* (e.g., diagnosis).

multiple topologies.

(1) **Direct Reasoning.** A single-pass mapping from raw waveforms directly to a clinical output without explicit intermediate states. This is the dominant paradigm in current physiological deep learning (e.g., end-to-end arrhythmia detection (Hannun et al., 2019; Attia et al., 2019)). While effective for pattern matching, this “black-box” approach lacks the mechanisms to isolate specific morphological evidence (e.g., P-wave absence), making it difficult to distinguish between true pathology and artifacts.

(2) **Linear-Chain Reasoning.** A sequential inference process where conclusions are built through intermediate, checkable steps (e.g., *Event Delineation* → *Interval Measurement* → *Rule Application* → *Diagnosis*). Adapting the principles of Chain-of-Thought prompting (Wei et al., 2022) to waveforms, this structure mimics the standard clinical workflow of first quantifying indices before interpreting them, thereby enforcing logical traceability and allowing for error localization at specific steps (Koh et al., 2020).

(3) **Branch-Structured Reasoning.** A non-linear process that explores multiple concurrent hypotheses or verification paths before aggregating a final decision. Inspired by Tree-of-Thoughts frameworks (Yao et al., 2023), this approach is uniquely suited for complex differential diagnosis. For example, a system might spawn separate reasoning branches to evaluate competing explanations for a wide-complex tachycardia (e.g., *Branch A: VTach vs. Branch B: SVT with Aberrancy*), weighing the evidence for each hypothesis against clinical guidelines before converging on a conclusion.

**Reasoning Objective.** We define *Reasoning Objective* as

the level of capability a system is expected to achieve. We propose a hierarchy of increasing complexity: *Level 1: Evidence*, *Level 2: Interpretation*, and *Level 3: Decision*.

(1) **Level 1: Evidence.** This level is to extract verifiable observations from raw waveforms. It prioritizes auditable outputs (timestamps and unit-bearing measurements) that can be independently re-computed from the signal.

① **Signal Quality Assessment.** This objective characterizes whether a segment/lead is physiologically trustworthy by detecting acquisition failures and artifacts and reporting explicit quality indicators (e.g., SQIs) used for weighting or exclusion (Clifford et al., 2012; Orphanidou et al., 2014a).

② **Event Segmentation.** This objective anchors clinically defined fiducial points with precise time indices (e.g., P/QRS/T onsets/offsets for ECG) under standardized conventions, providing the coordinates required for reproducible measurement (Party et al., 1981).

③ **Physiological Measurement Extraction.** This objective derives clinical indices as scalar values with physical units and well-defined windows (e.g., QT/QTc, QRS duration, HRV, PTT), ensuring results are recomputable from fiducials rather than implicit latent states (Malik, 1996; Ding & Zhang, 2019).

④ **Evidence Localization.** This objective links every claim to the exact time ranges and channels/leads (and intermediate fiducials/measurements) that support it, enabling audibility and discouraging unsupported post-hoc narratives (Rudin, 2019).

(2) **Level 2: Interpretation.** This level synthesizes Level-1 evidence into physiological semantics, turning timestamps

and measurements into clinically meaningful concepts under biological constraints.

① *Morphology Semantics*. This objective maps waveform geometry to physiological and pathological concepts (e.g., sawtooth atrial activity → atrial flutter; ST elevation → acute myocardial injury/ischemia). The emphasis is concept grounding in standardized ECG interpretation conventions rather than pattern matching alone (Kligfield et al., 2007).

② *Rhythm and Cycle Structure*. This objective infers the organizing logic of cycles from event sequences (e.g., regularity, bigeminy, compensatory pauses), using interval patterns and beat-to-beat dependencies. It is distinct from morphology because the same beat shape can appear under different rhythm regimes (Moody & Mark, 2001).

③ *Cross-Lead Concordance*. This objective enforces spatial consistency across leads as multiple views of the same cardiac source (e.g., inferior MI patterns coherently appearing in II/III/aVF but not aVL). It also uses multi-lead redundancy to separate global physiology from lead-specific corruption (Surawicz et al., 2009).

④ *Temporal Evolution*. This objective reasons over trajectories rather than snapshots, capturing clinically meaningful state transitions (e.g., progressive QRS widening, evolving repolarization changes). The key output is a time-ordered narrative of change supported by repeated evidence over windows (Goldberger et al., 2000).

⑤ *Physiological Cross-Signal Coupling*. This objective enforces mechanistic coherence across signals (e.g., an ECG electrical event should be followed by a mechanical pulse in PPG/ABP within a plausible latency such as PTT/PAT). Violations are treated as evidence of misalignment, or sensor failure rather than physiological conclusions (Allen, 2007).

(3) **Level 3: Decision**. This level converts interpreted evidence into actionable clinical logic, including rule-based conclusions, hypothesis management, and risk-aware decision behavior.

① *Guideline-Constrained Inference*. This objective applies explicit clinical criteria to Level-1 measurements with transparent rule invocation, especially for borderline cases. Outputs should explicitly cite which guideline conditions were satisfied or violated (Surawicz et al., 2009).

② *Differential Diagnosis*. This objective generates and ranks competing hypotheses that can explain the same evidence (e.g., distinguishing VT from SVT with aberrancy). Ranking should reflect supporting versus contradicting evidence for each hypothesis (Reiter, 1987).

③ *Mechanistic Explanation*. This objective provides a physiology-grounded causal rationale for the conclusion (e.g., PR prolongation indicating delayed AV nodal conduc-

tion). The goal is a plausible causal account that connects observed measurements to underlying mechanisms (Neuberg, 2003).

④ *Counterfactual Robustness*. This objective stress-tests decision stability under realistic perturbations, identifying which evidence is necessary versus incidental. Robustness is treated as a verification step rather than a post-hoc justification (Wachter et al., 2017).

⑤ *Uncertainty-Aware Abstention*. This objective calibrates confidence and supports safe deferral when evidence quality or ambiguity is high (e.g., requesting re-recording instead of forcing a label). Abstention should be principled (risk-coverage tradeoff), not ad hoc (Kompa et al., 2021).

### 3. Why Current Model Families Fall Short

#### 3.1. Key Challenges

**Challenge 1: Acquisition Heterogeneity and Signal Fidelity.** Acquisition heterogeneity directly degrades signal fidelity: intermittent wear (battery/adherence gaps), motion/perfusion artifacts, and device-specific sampling/lead configurations can systematically warp waveform morphology and timing rather than producing rare “outliers” (Kligfield et al., 2007; Shcherbina et al., 2017; Bent et al., 2020; Orphanidou et al., 2014b; Clifford et al., 2012; Fine et al., 2021). Reasoning models must therefore infer and condition on acquisition state (quality, continuity, configuration) as an explicit latent variable before clinical interpretation (Orphanidou et al., 2014b; Clifford et al., 2012).

**Challenge 2: Complex Physiological Semantics.** Waveform “language” is encoded in localized, high-frequency morphology (e.g., J-point/ST-T shape, dicrotic notch) that must be preserved and mapped to physiological concepts (Wagner & Strauss, 2013; Kligfield et al., 2007). Bridging pattern to mechanism requires grounding measurements in guideline definitions (e.g., ST elevation and acute myocardial injury/ischemia) rather than relying on geometric similarity alone (Thygesen et al., 2018).

**Challenge 3: Multivariate and Cross-Channel Dynamics.** Physiological signals are coupled views of shared biology (e.g., ECG electrical activation preceding the hemodynamic PPG pulse), imposing tight phase/latency constraints across channels (Allen, 2007; Orphanidou et al., 2014b). Effective reasoning must test cross-channel concordance to separate true pathology from single-sensor corruption under motion-related noise (Fine et al., 2021; Clifford et al., 2012).

**Challenge 4: The Data-Reasoning Mismatch.** Most public datasets provide coarse labels without intermediate evidence (e.g., measurements, rule traces), limiting supervision for verifiable inference (Wagner et al., 2020; Johnson et al., 2023; Gow et al., 2023; Oh et al., 2023). In contrast to

NLP where explicit rationales (e.g., chain-of-thought) are common, waveform reasoning needs evidence-centric representations that are both inspectable and intervention-friendly (Wei et al., 2022; Koh et al., 2020; Rudin, 2019).

### 3.2. Model Families and Limitations

We categorize prior work by model family because architecture determines the evidence interface and the resulting justification, spanning physiological foundation models (Section 3.2.1) and LLM-centric pipelines (Section 3.2.2).

#### 3.2.1. PHYSIOLOGICAL FOUNDATION MODELS

PhysioFMs are shifting from task-specific supervised pipelines to reusable biosignal backbones pretrained with self-supervision (Yang et al., 2023; Jiang et al., 2025b; Kataria et al., 2025). One route centers on scaling data and external validation, producing open or broadly accessible backbones and cross-domain evaluations (McKeen et al., 2025; Xu et al., 2025b; Li et al., 2025b; Abbaspourazad et al., 2023; Luo et al., 2024; Xu et al., 2025a; Saha et al., 2025). Second route is to better encode physiological structure: self-supervised designs that preserve spatio-temporal dependencies and yield broad downstream utility (Coppola et al., 2024; Na et al., 2024b;a; Wang et al., 2025c). A third route pushes representations toward richer clinical semantics and unified multi-task interfaces, sometimes via diagnosis or disease-centric objectives, and via language-aligned or multi-task waveform modeling (Tian et al., 2024; Jiang et al., 2024; 2025b; Cui et al., 2024). Finally, multimodal and time–frequency pretraining is emerging (notably in sleep/PSG), while cross-modal guidance and new sensing modalities broaden the scope of what “physiological foundation models” cover (Thapa et al., 2026; Huang et al., 2026; Kjaer et al., 2025; Pillai et al., 2025; Nie et al., 2025; Chen et al., 2025; Zhang et al., 2024; 2023).

**Limitations Analysis.** PhysioFMs primarily strengthen Level-1 (Evidence) capability by learning robust waveform representations under noise and domain shift. However, they rarely support complex physiological waveform reasoning because their dominant interface remains prediction-centric rather than evidence-centric. So they do not reliably expose localized evidence, translate morphology into intermediate interpretations, or implement explicit decision procedures with verification and uncertainty-aware abstention.

#### 3.2.2. LLM-CENTRIC ADAPTATIONS

LLM-centric methods are increasingly organized around a language-grounded interface. Rather than producing predictions alone, recent systems align waveforms with text and leverage multimodal instruction tuning to generate reports, answer questions, and support interactive interpretation (Zhao et al., 2025b; Wan et al., 2025; Yang et al.,

2025). Building on this interface, a growing line of work strengthens ECG–text alignment with clinically informed supervision, improving semantic grounding and generalization (Yu et al., 2024; Liu et al., 2025a; Li et al., 2025a). In parallel, “ECG understanding” and “clinical reasoning” are made more measurable by reframing evaluation as clinically oriented question answering, supported by knowledge-informed multimodal QA protocols (Oh et al., 2023; Wang et al., 2025b; Pham et al., 2025). To further improve factuality and traceability at inference time, many pipelines attach external clinical knowledge and retrieval (RAG-style) to report generation and diagnosis/QA (Tang et al., 2025; Yu et al., 2023). Finally, several efforts move toward more LLM-native signal interfaces by mapping waveforms into language-compatible units and bridging EEG–language for open-vocabulary decoding and assisted documentation (Jiang et al., 2025b; Chan et al., 2025; Jiang et al., 2025a).

**Limitations Analysis.** Many LLM-centric works still frame the task as question answering or answer prediction and thus rarely perform complex reasoning. Moreover, their waveform understanding is often shallow and indirect: upstream feature extraction can be lossy, the mapping from representations back to measurable evidence is frequently opaque, and generated rationales may be fluent but not entailed by the waveform under artifacts, missing channels, or distribution shift (Huang et al., 2025). Without explicit verification loops, these systems rarely guarantee guideline-consistent decisions or uncertainty-aware abstention.

## 4. Framework Design: Unifying Waveform Semantics and Language Intelligence

### 4.1. Design Goals

To bridge physiological signals and symbolic logic, we distill three design goals for waveform reasoning systems. (1) *Joint waveform understanding and language-level reasoning.* The system must align physiological waveform with language to have strong reasoning ability and waveform understanding ability (Nie et al., 2023). (2) *Agentic reasoning for complex reasoning objectives.* The system must iteratively test hypotheses via targeted measurements, cross-lead consistency checks, and tool-grounded computation (Yao et al., 2023; Shinn et al., 2023). (3) *Human-centered closed-loop evaluation and oversight.* The system must provide auditable verification objects, abstain or escalate under uncertainty, and improve via feedback-driven evaluation (Shinn et al., 2023; Yao et al., 2023).

### 4.2. Dual-Process Framework Architecture

Inspired by the dual-process theory in cognitive science (Kahneman, 2011) and its recent adaptations in machine learning (Bengio, 2017; Goyal & Bengio, 2022), we

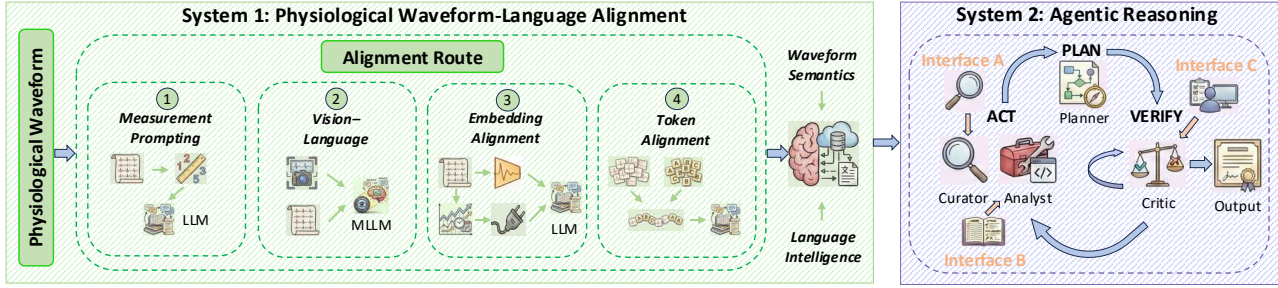


Figure 2. Framework Design: Unifying Waveform Semantics and Language Intelligence. System 1 aligns physiological waveforms and language via four alignment routes; System 2 performs plan-act-verify reasoning with tool grounding and human oversight.

advocate a *Dual-Process Architecture* with two systems to satisfy these goals simultaneously. System 1 and System 2 are **functionally distinct but operationally coupled**: (1) *System 1 (Physiological Waveform-Language Alignment)*: this system defines the alignment mechanism between physiological waveforms and the language model, so the system attains both strong reasoning ability and strong waveform understanding. (2) *System 2 (Agentic Reasoning: Plan-Act-Verify)*: this system is the controller that performs agentic reasoning. It plans, acquires missing verification objects via System 1 and tools, verifies claims with deterministic measurements, checks physiological/guideline constraints, and abstains or escalates under uncertainty.

#### 4.2.1. SYSTEM 1: PHYSIOLOGICAL WAVEFORM-LANGUAGE ALIGNMENT

The central question for System 1 is: *how should waveforms be aligned with a language model so that the system achieves strong reasoning ability and waveform understanding?* We propose **four alignment routes**, distinguished by the interface granularity and the degree of coupling:

**(i) Measurement Prompting Alignment.** The most simple route is to not align raw waveforms at all, but instead align waveform measurements into structured prompts that an LLM can reliably read. Recent works perform few-shot prompting on physiological time series by directly serializing measured sequences into the prompt (Liu et al., 2023), and propose retrieval-augmented, measurement-driven prompting for ECG diagnosis (Yu et al., 2023). In the general time-series domain, prompt-based reprogramming of LLMs (treating time series as a foreign language via prompt design) (Liu et al., 2024a; Kong et al., 2025) also fits this route when the interface is primarily textual verification objects rather than learned waveform embeddings.

**(ii) Vision-Language Alignment.** Vision-Language alignment converts waveforms into a visual surrogate and then leverages the well-developed MLLM stack for reasoning. This matches how humans consume waveforms (“look at the tracing”), enabling direct reuse of vision-language training recipes and instruction tuning. Recent work shows strong potential for ECG-image instruction tuning and benchmark-

ing (Liu et al., 2024b), and for grounded ECG understanding by combining plots with additional modalities or verification objects (Lan et al., 2025; Seki et al., 2025).

**(iii) Embedding Alignment.** Embedding alignment feeds the LLM continuous vectors computed from the raw waveform, retaining more signal detail. There are two common sub-routes. (1) *Specialized waveform encoders*: train an encoder jointly with an LLM-facing interface under reasoning supervision (instruction tuning, or contrastive alignment to clinical text), so the representation is optimized for downstream reasoning rather than generic reconstruction (Chow et al., 2024; Jin et al., 2024a; Langer et al., 2025). (2) *Adapter-based embedding alignment*: freeze a strong pre-trained PhysioFM encoder and learn lightweight adapters to map its representations into the LLM’s embedding space (Yu et al., 2025). The advantage is compute efficiency, as well as better retention of raw morphology than rendering.

**(iv) Discrete Token Alignment.** Token alignment makes waveform a *native language* by discretizing signal patches into tokens that are processed by an autoregressive LLM. This enables the cleanest conceptual integration—signals and text become a single sequence model—and opens the door to true native multimodal chain-of-thought reasoning, where the model can attend directly to signal tokens. Examples include large-scale tokenization-based pretraining that frames forecasting as language modeling (Ansari et al., 2024), and wavelet-based tokenization that discretizes time-localized frequency coefficients for autoregressive forecasting (Masserano et al., 2025).

#### 4.2.2. SYSTEM 2: AGENTIC REASONING: PLAN-ACT-VERIFY

System 2 treats waveform interpretation and decision-making as an iterative Plan-Act-Verify procedure rather than a single forward pass from inputs to labels. The reasoner plans, acts by requesting additional verification objects or invoking external procedures, and verifies intermediate claims before committing (Yao et al., 2023; Shinn et al., 2023). This agentic loop is especially important for complex waveform reasoning, where correctness depends on iterative evidence acquisition, cross-checking, and principled back-

tracking under uncertainty (Wei et al., 2026; Zhao et al., 2025a; Liu et al., 2025b).

Crucially, reasoning quality is measured by the trajectory: which verification objects are requested, which tools are used, how contradictions are handled, and when the system stops. System 2 maintains an explicit working state and selects actions that reduce uncertainty with minimal cost, terminating only when the decision is supported by auditable verification objects and physiological or guideline constraints are satisfied; otherwise it triggers targeted re-measurement or evidence expansion (Shinn et al., 2023). To make this loop auditable, we propose a minimal **Role-Based Agentic Architecture** with a Curator (evidence seeking), an Analyst (deterministic execution), and a Critic (verification).

**(i) The Curator: hypothesis-driven active perception (Act).** Most models passively accept a fixed input window, failing when decisive verification object lies outside the receptive field or is obscured by artifacts. The Curator addresses this by steering System 1 toward *targeted verification object acquisition*: it requests additional temporal context, alternative lead subsets, or focused re-representations conditioned on the current hypothesis (e.g., “Retrieve the preceding 30 seconds to test sudden vs. gradual onset”) (Zhao et al., 2025a). This active perception mirrors clinical workflows in which clinicians zoom, scroll telemetry, or examine specific leads before declaring a rhythm or morphology-based diagnosis.

**(ii) The Analyst: deterministic tool execution with provenance (Act → Verify).** Waveform reasoning often hinges on precise measurements, which LLMs are not reliable at producing directly. The Analyst therefore implements *code-as-reasoning*: instead of emitting numbers as free-form tokens, it calls deterministic tools or executable code for delineation, peak detection, interval computation, and statistical checks, ensuring each quantitative claim is backed by a reproducible execution trace (provenance) rather than a stochastic guess (Liu et al., 2025b). This also enables principled re-measurement when the Critic requests alternative windows, preprocessing assumptions, or robustness checks.

**(iii) The Critic: reflexion, constraint checking, and guideline adherence (Verify).** A major failure mode of end-to-end systems is producing fluent but physiologically impossible claims. The Critic implements a reflexion loop that stress-tests intermediate conclusions against (i) physiological constraints (e.g., rhythm regularity vs. RR variability, plausible interval ranges, cross-lead consistency) and (ii) guideline logic when applicable. When inconsistencies are detected (e.g., “irregular rhythm” declared but RR intervals are constant), the Critic triggers *backtracking*: it revises the plan, asks the Curator for additional verification object, and instructs the Analyst to re-measure or run alternative checks (Shinn et al., 2023; Liu et al., 2025b). Importantly,

the Critic should also control *stop conditions*: if verification object quality remains low or tool outputs conflict, it should abstain or escalate rather than force a brittle decision.

**Closing the Loop: Verification, Provenance, and Oversight Interfaces.** While the Curator–Analyst–Critic decomposition specifies who acts in the Plan–Act–Verify loop, verifiable waveform reasoning requires explicit *closed-loop interfaces* that specify *what* is exchanged and recorded. **Interface A: Verification object querying from System 1.** System 2 can call System 1 to obtain alternative waveform-grounded views under uncertainty. Concretely, System 2 may request different temporal context, lead subsets, preprocessing assumptions, or specific verification objects (e.g., delineation- or measurement-oriented outputs). Importantly, this is not model updating; it is inference-time, hypothesis-driven verification object acquisition over a fixed alignment mechanism. **Interface B: Auditable ledger of verification object for tool-grounded claims.** Rather than reiterating tool use, we require a protocol-level verification object ledger: any quantitative assertion produced by System 2 can be accompanied by a replayable record of how it was obtained (tool name, parameters, software version, input segment/lead identifiers, and failure modes when applicable). This ledger makes outputs reproducible and supports debugging and downstream auditing, while leaving tool selection and execution to the Analyst behavior. **Interface C: Human oversight and feedback-to-memory.** For safety-critical deployment, the system should expose auditable artifacts so clinicians can validate what the model relied on and intervene when verification object quality is low. Crucially, human signals should be captured in structured form and stored as workflow memory. Over time, these feedback traces can be retrieved to guide future Plan–Act–Verify trajectories or internalized through post-training optimization, enabling continual improvement beyond one-off reflexion.

## 5. Evaluation: From Prediction to Verifiable Physiological Waveform Reasoning

Physiological waveform reasoning should be evaluated as **verifiable, episode-level decision making**, not single-shot prediction. The goal is not only to be correct, but to be *checkable*: intermediate claims and final decisions are credited only if they can be re-derived from reproducible **verification objects** and replayable procedures, with uncertainty handled via abstention or escalation when support is insufficient.

### 5.1. Episode-Level Evaluation via a Verification Ledger

**Unit (reasoning episode).** We evaluate a *reasoning episode*: the complete trajectory from waveform input and task query to final output, including intermediate requests for additional views, verification-object extraction, measurements, tool calls, and any backtracking. The episode is the atomic unit.

**Interface (Verification Ledger).** To make verifiability testable, each episode must expose a *Verification Ledger*: a minimal structured record that links (i) *verification objects*: signal-quality summaries, localized events with explicit time/lead indices, and unit-consistent measurements defined by reproducible windows and procedures; (ii) *tool provenance*: tool names, parameters, software versions/ hashes, input segment/lead identifiers, and failure modes; and (iii) *the final decision*: including abstention or escalation when warranted. A decision is credited only if it can be reconstructed from the ledger by replayable checks.

**Criteria and metrics.** Given the ledger interface, we score task-agnostic episode-level criteria: (i) *Traceability*: major claims and decision labels are linked to sufficient verification objects or tool outputs for audit; (ii) *Replayability*: quantitative statements are reproducible by rerunning the logged procedures under recorded parameters; (iii) *Robustness vs. sensitivity*: outputs remain stable under nuisance perturbations (noise, baseline wander, resampling, preprocessing variants, missing leads) yet respond appropriately to clinically meaningful counterfactual changes; (iv) *Uncertainty-aware safety*: the system surfaces low-quality signals, conflicts, or violated constraints and abstains/escalates when a safe conclusion is not supported. And *budgeted verifiable success* should be additionally reported: performance as a function of query/tool-call budget and audit effort, rather than a single unconstrained score.

## 5.2. Component and Closed-Loop Evaluation

**Component evaluation.** In the dual-process framework, System 1 is evaluated by *verification-object fidelity*: correctness of extracted signal-quality summaries, localized events, and unit-consistent measurements, consistency across views, and calibration of quality indicators under heterogeneity. System 2 is evaluated by *trajectory validity*: whether it requests appropriate verification objects, applies coherent checks and constraints, ensures intermediate claims and final decisions are entailed by the ledger, detects and resolves contradictions via principled backtracking, and invokes abstention/escalation when warranted.

**Closed-loop Interfaces.** System-level evaluation tests whether closed-loop integration improves verifiable outcomes under realistic constraints. **Interface A: Verification-object querying** compares single-shot episodes to budgeted multi-turn episodes where System 2 can query System 1 for alternative views, reporting improvement in budgeted verifiable success and reduction in unresolved conflicts per additional query. **Interface B: Auditable Verification Ledger** measures replay consistency, mismatch rate, and conflict handling. **Interface C: Audit and correction** quantifies audit cost and correction yield, with special focus on abstention under low-quality signals, where safe deferral is

preferable to confident but uncheckable conclusions.

## 6. Alternative Views

**Predicting measurements is enough, interpretation and decision are unnecessary.** A plausible view is that models should focus on predicting measurements, and that interpretation and decision-making are unnecessary extras. Our response is that interpretation and decision are essential for clinical meaning and safety: medicine is not just producing numbers, but determining what they imply in context, resolving cross-lead conflicts, and recognizing when evidence is insufficient. So a system that can handle interpretation and decision-making is what makes waveform modeling clinically meaningful and safe.

**End-to-end prediction is sufficient.** A credible alternative is that end-to-end prediction is enough: prioritize accuracy, calibration, and robustness under shift, and avoid explicitly staged reasoning that can compound errors. Our response is not to replace end-to-end learning, but to constrain it where clinical stakes demand verifiability. The central risk is undetected wrongness under artifacts, missing leads, and device/site heterogeneity, so high average AUROC is not a safety guarantee. Reasoning is therefore warranted only when it is checkable: decisions must be tied to localized time-lead evidence, supported by replayable measurements, challenged by consistency/counterfactual tests, and paired with explicit abstention when evidence is insufficient.

**Data and deployment, not reasoning, are the bottleneck.**

Some researchers argue that progress is constrained by practicalities rather than new framings: better datasets matter more than reasoning. In waveforms, failures are dominated by acquisition noise, device/protocol heterogeneity; without realistic deployment tests, reasoning claims are hard to verify. Our response is that our position is necessary because turning these practical needs into implementable requirements demands a system that couples *waveform semantics* with *language intelligence*: System 1 aligns physiological waveform-language, and System 2 supports agentic Plan-Act-Verify reasoning for complex waveform reasoning.

## 7. Conclusion

We frame physiological waveform reasoning as verifiable, episode-level inference that links raw biosignals to clinical decisions beyond end-to-end prediction. Accordingly, we propose a dual-process blueprint: System 1 provides a stable waveform-language alignment interface that exposes queryable signal evidence, while System 2 performs agentic Plan-Act-Verify reasoning, acquiring missing information and applying deterministic consistency checks under uncertainty. To make progress comparable and deployable, we advocate evaluation through a Verification Ledger that prioritizes traceability and robustness beyond accuracy.

## References

- 440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494
- Abbaspourazad, S., Elachqar, O., Miller, A. C., Emrani, S., Nallasamy, U., and Shapiro, I. Large-scale training of foundation models for wearable biosignals. *arXiv preprint arXiv:2312.05409*, 2023.
- Allen, J. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1, 2007.
- Ansari, A. F., Stella, L., Turkmen, C., Zhang, X., Mercado, P., Shen, H., Shchur, O., Rangapuram, S. S., Arango, S. P., Kapoor, S., et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.
- Attia, Z. I., Kapa, S., Lopez-Jimenez, F., McKie, P. M., Ladewig, D. J., Satam, G., Pellikka, P. A., Enriquez-Sarano, M., Noseworthy, P. A., Munger, T. M., et al. Screening for cardiac contractile dysfunction using an artificial intelligence-enabled electrocardiogram. *Nature medicine*, 25(1):70–74, 2019.
- Bengio, Y. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.
- Bent, B., Goldstein, B. A., Kibbe, W. A., and Dunn, J. P. Investigating sources of inaccuracy in wearable optical heart rate sensors. *NPJ digital medicine*, 3(1):18, 2020.
- Buzsáki, G. *Rhythms of the Brain*. Oxford university press, 2006.
- Cao, D., Jia, F., Arik, S. O., Pfister, T., Zheng, Y., Ye, W., and Liu, Y. Tempo: Prompt-based generative pre-trained transformer for time series forecasting. In *The Twelfth International Conference on Learning Representations*, 2024.
- Chan, N., Parker, F., Zhang, C., Bennett, W., Jia, M. Y., Fackler, J., and Ghobadi, K. MedtSLLM: Medical time series analysis using multimodal llms. *IEEE Journal of Biomedical and Health Informatics*, 2025.
- Chang, C., Hwang, J., Shi, Y., Wang, H., Peng, W.-C., Chen, T.-F., and Wang, W. Time-imm: A dataset and benchmark for irregular multimodal multivariate time series. 2025a.
- Chang, C., Lo, M.-C., Peng, W.-C., and Chen, T.-F. PromptSS: A prompting-based approach for interactive multi-granularity time series segmentation. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pp. 209–220, 2025b.
- Chen, Z., Ding, C., Kataria, S., Yan, R., Wang, M., Lee, R., and Hu, X. Gpt-ppg: a gpt-based foundation model for photoplethysmography signals. *Physiological Measurement*, 46(5):055004, 2025.
- Chow, W., Gardiner, L., Hallgrímsson, H. T., Xu, M. A., and Ren, S. Y. Towards time series reasoning with llms. *arXiv preprint arXiv:2409.11376*, 2024.
- Clifford, G., Behar, J., Li, Q., and Rezek, I. Signal quality indices and data fusion for determining clinical acceptability of electrocardiograms. *Physiological measurement*, 33(9):1419, 2012.
- Coppola, E., Savardi, M., Massussi, M., Adamo, M., Metra, M., and Signoroni, A. Hubert-ecg as a self-supervised foundation model for broad and scalable cardiac applications. *medRxiv*, pp. 2024–11, 2024.
- Craik, A., He, Y., and Contreras-Vidal, J. L. Deep learning for electroencephalogram (eeg) classification tasks: a review. *Journal of Neural Engineering*, 16(3):031001, 2019.
- Cui, W., Jeong, W., Thölke, P., Medani, T., Jerbi, K., Joshi, A. A., and Leahy, R. M. Neuro-gpt: Towards a foundation model for eeg. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5. IEEE, 2024.
- De Luca, C. J. The use of surface electromyography in biomechanics. *Journal of Applied Biomechanics*, 13(2):135–163, 1997.
- Ding, X. and Zhang, Y.-T. Pulse transit time technique for cuffless unobtrusive blood pressure measurement: from theory to algorithm. *Biomedical engineering letters*, 9(1):37–52, 2019.
- Fine, J., Branan, K. L., Rodriguez, A. J., Boonya-Ananta, T., Ajmal, Ramella-Roman, J. C., McShane, M. J., and Cote, G. L. Sources of inaccuracy in photoplethysmography for continuous cardiovascular monitoring. *Biosensors*, 11(4):126, 2021.
- Geifman, Y. and El-Yaniv, R. Selectivenet: A deep neural network with an integrated reject option. In *International conference on machine learning*, pp. 2151–2159. PMLR, 2019.
- Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- Goswami, M., Szafer, K., Choudhry, A., Cai, Y., Li, S., and Dubrawski, A. Moment: A family of open time-series foundation models. In *Forty-first International Conference on Machine Learning*, 2024.
- Gow, B., Pollard, T., Nathanson, L. A., Johnson, A., Moody, B., Fernandes, C., Greenbaum, N., Waks, J. W., Eslami,

- 495 P., Carbonati, T., et al. Mimic-iv-ecg: Diagnostic elec-  
496 trocardiogram matched subset. *Type: dataset*, 6:13–14,  
497 2023.
- 498
- 499 Goyal, A. and Bengio, Y. Inductive biases for deep learn-  
500 ing of higher-level cognition. *Proceedings of the Royal*  
501 *Society A*, 478(2266):20210068, 2022.
- 502
- 503 Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On  
504 calibration of modern neural networks. In *International*  
505 *conference on machine learning*, pp. 1321–1330. PMLR,  
506 2017.
- 507
- 508 Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H.,  
509 Bourn, C., Turakhia, M. P., and Ng, A. Y. Cardiologist-  
510 level arrhythmia detection and classification in ambu-  
511 latory electrocardiograms using a deep neural network.  
512 *Nature medicine*, 25(1):65–69, 2019.
- 513
- 514 Hermens, H. J., Freriks, B., Disselhorst-Klug, C., and Rau,  
515 G. Development of recommendations for semg sensors  
516 and sensor placement procedures. *Journal of Electromyo-*  
517 *graphy and Kinesiology*, 10(5):361–374, 2000.
- 518
- 519 Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H.,  
520 Chen, Q., Peng, W., Feng, X., Qin, B., et al. A survey on  
521 hallucination in large language models: Principles, taxon-  
522 omy, challenges, and open questions. *ACM Transactions*  
523 *on Information Systems*, 43(2):1–55, 2025.
- 524
- 525 Huang, W., Wang, Y., Cheng, H., Xu, W., Li, T., Wu, X.,  
526 Xu, H., Liao, P., Cui, Z., Zou, Q., et al. A unified time-  
527 frequency foundation model for sleep decoding. *Nature*  
528 *Communications*, 2026.
- 529
- 530 Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L.,  
531 and Muller, P.-A. Deep learning for time series classifica-  
532 tion: a review. *Data mining and knowledge discovery*, 33  
533 (4):917–963, 2019.
- 534
- 535 Jiang, M., Zhang, S., Yang, Z., Wu, M., Jiang, W., Guo,  
536 Z., Zhang, W., Liu, R., Zhang, S., Li, Y., et al. Elastiq:  
537 Eeg-language alignment with semantic task instruction  
538 and querying. *arXiv preprint arXiv:2509.24302*, 2025a.
- 539
- 540 Jiang, W., Zhao, L., and Lu, B.-l. Large brain model for  
541 learning generic representations with tremendous eeg  
542 data in bci. In *The Twelfth International Conference on*  
543 *Learning Representations*, 2024.
- 544
- 545 Jiang, W., Wang, Y., Lu, B.-l., and Li, D. Neurolm: A  
546 universal multi-task foundation model for bridging the  
547 gap between language and eeg signals. In *The Thirteenth*  
548 *International Conference on Learning Representations*,  
549 2025b.
- Jin, M., Wang, S., Ma, L., Chu, Z., Zhang, J., Shi, X., Chen,  
P.-Y., Liang, Y., Li, Y.-f., Pan, S., et al. Time-llm: Time se-  
ries forecasting by reprogramming large language models.  
In *International Conference on Learning Representations*,  
2024a.
- Jin, M., Zhang, Y., Chen, W., Zhang, K., Liang, Y., Yang,  
B., Wang, J., Pan, S., and Wen, Q. Position: What can  
large language models tell us about time series analysis.  
In *41st International Conference on Machine Learning*.  
MLResearchPress, 2024b.
- Johnson, A. E., Bulgarelli, L., Shen, L., Gayles, A., Sham-  
mout, A., Horng, S., Pollard, T. J., Hao, S., Moody, B.,  
Gow, B., et al. Mimic-iv, a freely accessible electronic  
health record dataset. *Scientific data*, 10(1):1, 2023.
- Kahneman, D. *Thinking, fast and slow*. macmillan, 2011.
- Kataria, S., Wu, Y., Chen, Z., Kwak, H. G., Xu, Y., Panchu-  
marthi, L. Y., Xiao, R., Lu, J., Ermis, A., Zhao, A.,  
et al. Generalist vs specialist time series foundation  
models: Investigating potential emergent behaviors in  
assessing human health using ppg signals. *arXiv preprint*  
*arXiv:2510.14254*, 2025.
- Kjaer, M. R., Thapa, R., Ganjoo, G., Moore IV, H., Jen-  
num, P. J., Westover, B. M., Zou, J., Mignot, E., He, B.,  
and Brink-Kjaer, A. Stanford sleep bench: Evaluating  
polysomnography pre-training methods for sleep founda-  
tion models. *arXiv preprint arXiv:2512.09591*, 2025.
- Kligfield, P., Gettes, L. S., Bailey, J. J., Childers, R., Deal,  
B. J., Hancock, E. W., Van Herpen, G., Kors, J. A., Mac-  
farlane, P., Mirvis, D. M., et al. Recommendations for  
the standardization and interpretation of the electrocar-  
diogram: part i: the electrocardiogram and its technology  
a scientific statement from the american heart association  
electrocardiography and arrhythmias committee, council  
on clinical cardiology; the american college of cardiology  
foundation; and the heart rhythm society endorsed by the  
international society for computerized electrocardiology.  
*Journal of the American College of Cardiology*, 49(10):  
1109–1127, 2007.
- Koh, P. W., Nguyen, T., Tang, Y. S., Mussmann, S., Pierson,  
E., Kim, B., and Liang, P. Concept bottleneck models.  
In *International Conference on Machine Learning*, pp.  
5338–5348. PMLR, 2020.
- Kompa, B., Snoek, J., and Beam, A. L. Second opinion  
needed: communicating uncertainty in medical machine  
learning. *NPJ Digital Medicine*, 4(1):4, 2021.
- Kong, Y., Yang, Y., Hwang, Y., Du, W., Zohren, S., Wang,  
Z., Jin, M., and Wen, Q. Time-mqa: Time series multi-  
task question answering with context enhancement. *arXiv*  
*preprint arXiv:2503.01875*, 2025.

- 550 Lan, X., Wu, F., He, K., Zhao, Q., Hong, S., and Feng, M.  
551 Gem: Empowering mllm for grounded ecg understand-  
552 ing with time series and images. *Advances in Neural*  
553 *Information Processing Systems*, 2025.
- 554 Langer, P., Kaar, T., Rosenblattl, M., Xu, M. A., Chow, W.,  
555 Maritsch, M., Verma, A., Han, B., Kim, D. S., Chubb,  
556 H., et al. Opentslm: Time-series language models for  
557 reasoning over multivariate medical text-and time-series  
558 data. *arXiv preprint arXiv:2510.02410*, 2025.
- 559 Li, H., Liu, C., Ding, Z., Liu, Z., and Huang, Z. Fine-grained  
560 ecg-text contrastive learning via waveform understanding  
561 enhancement. *arXiv preprint arXiv:2505.11939*, 2025a.
- 562 Li, J., Aguirre, A. D., Junior, V. M., Jin, J., Liu, C., Zhong,  
563 L., Sun, C., Clifford, G., Brandon Westover, M., and  
564 Hong, S. An electrocardiogram foundation model built on  
565 over 10 million recordings. *NEJM AI*, 2(7):AIoa2401033,  
566 2025b.
- 567 Liu, C., Ouyang, C., Wan, Z., Wang, H., Bai, W., and Ar-  
568 cucci, R. Knowledge-enhanced multimodal ecg represen-  
569 tation learning with arbitrary-lead inputs. *arXiv preprint*  
570 *arXiv:2502.17900*, 2025a.
- 571 Liu, H., Zhao, Z., Wang, J., Kamarthi, H., and Prakash,  
572 B. A. Lstprompt: Large language models as zero-shot  
573 time series forecasters by long-short-term prompting. In  
574 *Findings of the Association for Computational Linguistics*  
575 *ACL 2024*, pp. 7832–7840, 2024a.
- 576 Liu, P., Fons, E., Vyetrenko, S., Borrajo, D., Potluru, V.,  
577 and Veloso, M. Ts-agent: A time series reasoning agent  
578 with iterative statistical insight gathering. *arXiv preprint*  
579 *arXiv:2510.07432*, 2025b.
- 580 Liu, R., Bai, Y., Yue, X., and Zhang, P. Teach multimodal  
581 llms to comprehend electrocardiographic images. *arXiv*  
582 *preprint arXiv:2410.19008*, 2024b.
- 583 Liu, X., McDuff, D., Kovacs, G., Galatzer-Levy, I., Sun-  
584 shine, J., Zhan, J., Poh, M.-Z., Liao, S., Di Achille, P.,  
585 and Patel, S. Large language models are few-shot health  
586 learners. *arXiv preprint arXiv:2305.15525*, 2023.
- 587 Luo, Y., Chen, Y., Salekin, A., and Rahman, T. Toward  
588 foundation model for multivariate wearable sensing of  
589 physiological signals. *arXiv preprint arXiv:2412.09758*,  
590 2024.
- 591 Malik, M. Heart rate variability: Standards of measurement,  
592 physiological interpretation, and clinical use: Task force  
593 of the european society of cardiology and the north amer-  
594 ican society for pacing and electrophysiology. *Annals of*  
595 *Noninvasive Electrocardiology*, 1(2):151–181, 1996.
- 596 Masserano, L., Ansari, A. F., Han, B., Zhang, X., Faloutsos,  
597 C., Mahoney, M. W., Wilson, A. G., Park, Y., Rangapu-  
598 ram, S. S., Maddix, D. C., et al. Enhancing foundation  
599 models for time series forecasting via wavelet-based tok-  
600 enization. In *Forty-second International Conference on*  
601 *Machine Learning*, 2025.
- 602 McKeen, K., Masood, S., Toma, A., Rubin, B., and Wang, B.  
603 Ecg-fm: An open electrocardiogram foundation model.  
604 *JAMIA open*, 8(5):ooaf122, 2025.
- Mehari, T. and Strodthoff, N. Self-supervised representation  
learning from 12-lead ecg data. *Computers in biology  
and medicine*, 141:105114, 2022.
- Moody, G. B. and Mark, R. G. The impact of the mit-bih  
arrhythmia database. *IEEE engineering in medicine and  
biology magazine*, 20(3):45–50, 2001.
- Na, Y., Park, M., Tae, Y., and Joo, S. Guiding masked repre-  
sentation learning to capture spatio-temporal relationship  
of electrocardiogram. In *International Conference on*  
*Learning Representations*, 2024a.
- Na, Y., Park, M., Tae, Y., and Joo, S. Guiding masked repre-  
sentation learning to capture spatio-temporal relationship  
of electrocardiogram. In *The Twelfth International Con-  
ference on Learning Representations*, 2024b.
- Neuberg, L. G. Causality: models, reasoning, and infer-  
ence, by judea pearl, cambridge university press, 2000.  
*Econometric Theory*, 19(4):675–685, 2003.
- Nie, G., Tang, G., Xiao, Y., Li, J., Huang, S., Zhang,  
D., Zhao, Q., and Hong, S. Anypgp: An ecg-guided  
ppg foundation model trained on over 100,000 hours of  
recordings for holistic health profiling. *arXiv preprint*  
*arXiv:2511.01747*, 2025.
- Nie, Y., Nguyen, N. H., Sinthong, P., and Jayaraman, K.  
A time series is worth 64 words: Long-term forecasting  
with transformers. *International Conference on Learning*  
*Representations (ICLR)*, 2023.
- Oh, J., Lee, G., Bae, S., Kwon, J.-m., and Choi, E. Ecg-qa:  
A comprehensive question answering dataset combined  
with electrocardiogram. *Advances in Neural Information*  
*Processing Systems*, 36:66277–66288, 2023.
- Orphanidou, C., Bonnici, T., Charlton, P., Clifton, D., Val-  
lance, D., and Tarassenko, L. Signal-quality indices for  
the electrocardiogram and photoplethysmogram: Deriva-  
tion and applications to wireless monitoring. *IEEE jour-  
nal of biomedical and health informatics*, 19(3):832–838,  
2014a.
- Orphanidou, C., Bonnici, T., Charlton, P., Clifton, D., Val-  
lance, D., and Tarassenko, L. Signal-quality indices for

- 605 the electrocardiogram and photoplethysmogram: Deriva-  
606 tion and applications to wireless monitoring. *IEEE jour-*  
607 *nal of biomedical and health informatics*, 19(3):832–838,  
608 2014b.
- 609 Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D.,  
610 Nowozin, S., Dillon, J., Lakshminarayanan, B., and  
611 Snoek, J. Can you trust your model’s uncertainty? evalu-  
612 ating predictive uncertainty under dataset shift. *Advances*  
613 *in neural information processing systems*, 32, 2019.
- 614 Party, C. E. W., Willems, J., Arnaud, P., van Bommel, J.,  
615 Bourdillon, P., Brohet, C., Dalla Volta, S., Degani, R.,  
616 Denis, B., Demeester, M., et al. Common standards for  
617 quantitative electrocardiography: The cse pilot study. In  
618 *Medical Informatics Europe 81: Third Congress of the*  
619 *European Federation of Medical Informatics Proceed-*  
620 *ings, Toulouse, France March 9–13, 1981*, pp. 319–326.  
621 Springer, 1981.
- 622 Pham, H. M., Tang, J., Saeed, A., and Ma, D. Q-heart: Ecg  
623 question answering via knowledge-informed multimodal  
624 llms. *arXiv preprint arXiv:2505.06296*, 2025.
- 625 Pillai, A., Spathis, D., Kawsar, F., and Malekzadeh, M. Pa-  
626 pagei: Open foundation models for optical physiological  
627 signals. In *The Thirteenth International Conference on*  
628 *Learning Representations*, 2025.
- 629 Rajpurkar, P., Hannun, A. Y., Haghpanahi, M., Bourn,  
630 C., and Ng, A. Y. Cardiologist-level arrhythmia detec-  
631 tion with convolutional neural networks. *arXiv preprint*  
632 *arXiv:1707.01836*, 2017.
- 633 Rajpurkar, P., Chen, E., Banerjee, O., and Topol, E. J. Ai  
634 in health and medicine. *Nature medicine*, 28(1):31–38,  
635 2022.
- 636 Reiter, R. A theory of diagnosis from first principles. *Artifi-*  
637 *cial Intelligence*, 32(1):57–95, 1987.
- 638 Ribeiro, A. H., Ribeiro, M. H., Paixão, G. M., Oliveira,  
639 D. M., Gomes, P. R., Canazart, J. A., Ferreira, M. P.,  
640 Andersson, C. R., Macfarlane, P. W., Meira Jr, W., et al.  
641 Automatic diagnosis of the 12-lead ecg using a deep neu-  
642 ral network. *Nature communications*, 11(1):1760, 2020.
- 643 Rudin, C. Stop explaining black box machine learning  
644 models for high stakes decisions and use interpretable  
645 models instead. *Nature machine intelligence*, 1(5):206–  
646 215, 2019.
- 647 Saha, M., Xu, M. A., Mao, W., Neupane, S., Rehg, J. M.,  
648 and Kumar, S. Pulse-ppg: An open-source field-trained  
649 ppg foundation model for wearable applications across  
650 lab and field settings. *Proceedings of the ACM on Inter-*  
651 *active, Mobile, Wearable and Ubiquitous Technologies*, 9  
652 (3):1–35, 2025.
- Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli,  
M., Hambro, E., Zettlemoyer, L., Cancedda, N., and  
Scialom, T. Toolformer: Language models can teach  
themselves to use tools. *Advances in Neural Information*  
*Processing Systems*, 36:68539–68551, 2023.
- Seki, T., Kawazoe, Y., Ito, H., Akagi, Y., Takiguchi, T.,  
and Ohe, K. Assessing the performance of zero-shot  
visual question answering in multimodal large language  
models for 12-lead ecg image interpretation. *Frontiers in*  
*cardiovascular medicine*, 12:1458289, 2025.
- Shcherbina, A., Mattsson, C. M., Waggott, D., Salisbury, H.,  
Christle, J. W., Hastie, T., Wheeler, M. T., and Ashley,  
E. A. Accuracy in wrist-worn, sensor-based measure-  
ments of heart rate and energy expenditure in a diverse  
cohort. *Journal of personalized medicine*, 7(2):3, 2017.
- Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., and  
Yao, S. Reflexion: Language agents with verbal rein-  
forcement learning. *Advances in Neural Information*  
*Processing Systems*, 36:8634–8652, 2023.
- Surawicz, B., Childers, R., Deal, B. J., and Gettes, L. S.  
Aha/accf/hrs recommendations for the standardization  
and interpretation of the electrocardiogram: part iii: intra-  
ventricular conduction disturbances a scientific statement  
from the american heart association electrocardiography  
and arrhythmias committee, council on clinical cardi-  
ology; the american college of cardiology foundation;  
and the heart rhythm society endorsed by the international  
society for computerized electrocardiology. *Journal of*  
*the American College of Cardiology*, 53(11):976–981,  
2009.
- Tang, J., Xia, T., Lu, Y., Mascolo, C., and Saeed, A. Elec-  
trocardiogram report generation and question answering  
via retrieval-augmented self-supervised modeling. In  
*ICASSP 2025-2025 IEEE International Conference on*  
*Acoustics, Speech and Signal Processing (ICASSP)*, pp.  
1–5. IEEE, 2025.
- Thapa, R., Kjaer, M. R., He, B., Covert, I., Moore IV,  
H., Hanif, U., Ganjoo, G., Westover, M. B., Jennum,  
P., Brink-Kjaer, A., et al. A multimodal sleep foundation  
model for disease prediction. *Nature Medicine*, pp. 1–11,  
2026.
- Thygesen, K., Alpert, J. S., Jaffe, A. S., Chaitman,  
B. R., Bax, J. J., Morrow, D. A., White, H. D., and  
on behalf of the Joint European Society of Cardiology  
(ESC)/American College of Cardiology (ACC)/American  
Heart Association (AHA)/World Heart Federation (WHF)  
Task Force for the Universal Definition of Myocardial In-  
farction, E. G. Fourth universal definition of myocardial  
infarction (2018). *Journal of the American college of*  
*cardiology*, 72(18):2231–2264, 2018.

- 660 Tian, Y., Li, Z., Jin, Y., Wang, M., Wei, X., Zhao, L., Liu, Y.,  
661 Liu, J., and Liu, C. Foundation model of ecg diagnosis:  
662 Diagnostics and explanations of any form and rhythm on  
663 ecg. *Cell Reports Medicine*, 5(12), 2024.
- 664  
665 Wachter, S., Mittelstadt, B., and Russell, C. Counterfactual  
666 explanations without opening the black box: Automated  
667 decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.
- 668  
669 Wagner, G. S. and Strauss, D. G. *Marriott's practical elec-*  
670 *trocardiography*. Lippincott Williams & Wilkins, 2013.
- 671  
672 Wagner, P., Strodthoff, N., Bousseljot, R.-D., Kreiseler, D.,  
673 Lunze, F. I., Samek, W., and Schaeffter, T. Ptb-xl, a large  
674 publicly available electrocardiography dataset. *Scientific*  
675 *data*, 7(1):1–15, 2020.
- 676  
677 Wan, Z., Liu, C., Wang, X., Tao, C., Shen, H., Xiong, J.,  
678 Arcucci, R., Yao, H., and Zhang, M. Meit: Multimodal  
679 electrocardiogram instruction tuning on large language  
680 models for report generation. In *Findings of the Asso-*  
681 *ciation for Computational Linguistics: ACL 2025*, pp.  
682 14510–14527, 2025.
- 683  
684 Wang, X., Han, K., Xu, Y., Luo, X., Sun, Y., Wang, W., and  
685 Yang, C. Simulator and experience enhanced diffusion  
686 model for comprehensive ecg generation. *International*  
687 *Conference on Learning Representations*, 2025a.
- 688  
689 Wang, X., Kang, J., Han, P., Zhao, Y., Liu, Q., He, L., Zhang,  
690 L., Dai, L., Wang, Y., and Tao, J. Ecg-expert-qa: A bench-  
691 mark for evaluating medical large language models in  
692 heart disease diagnosis. *arXiv preprint arXiv:2502.17475*,  
693 2025b.
- 694  
695 Wang, X., Liu, X., Liu, X., Si, Q., Xu, Z., Li, Y., and  
696 Zhen, X. Eeg-dino: Learning eeg foundation models via  
697 hierarchical self-distillation. In *International Conference*  
698 *on Medical Image Computing and Computer-Assisted*  
699 *Intervention*, pp. 196–205. Springer, 2025c.
- 700  
701 Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi,  
702 E., Le, Q. V., Zhou, D., et al. Chain-of-thought prompting  
703 elicits reasoning in large language models. *Advances in*  
704 *neural information processing systems*, 35:24824–24837,  
705 2022.
- 706  
707 Wei, T., Li, T.-W., Liu, Z., Ning, X., Yang, Z., Zou, J., Zeng,  
708 Z., Qiu, R., Lin, X., Fu, D., et al. Agentic reasoning for  
709 large language models. *arXiv preprint arXiv:2601.12538*,  
710 2026.
- 711  
712 Wiggins, W. F. and Tejani, A. S. On the opportunities and  
713 risks of foundation models for natural language process-  
714 ing in radiology. *Radiology: Artificial Intelligence*, 4(4):  
e220119, 2022.
- Xu, M. A., Narain, J., Darnell, G., Hallgrímsson, H. T.,  
Jeong, H., Forde, D., Fineman, R. A., Raghuram, K. J.,  
Rehg, J. M., and Ren, S. Y. Relcon: Relative contrastive  
learning for a motion foundation model for wearable data.  
In *The Thirteenth International Conference on Learning*  
*Representations*, 2025a.
- Xu, Y., Wang, X., Wu, Y., Jin, W., Hu, X., and Yang, C. Ecg-  
moe: Mixture-of-expert electrocardiogram foundation  
model. In *NeurIPS 2025 Workshop on Learning from*  
*Time Series for Health*, 2025b.
- Yang, C., Westover, M., and Sun, J. Biot: Biosignal trans-  
former for cross-data learning in the wild. *Advances*  
*in Neural Information Processing Systems*, 36:78240–  
78260, 2023.
- Yang, K., Hong, M., Zhang, J., Luo, Y., Zhao, S., Zhang,  
O., Yu, X., Zhou, J., Yang, L., Zhang, P., et al. Ecg-lm:  
Understanding electrocardiogram with a large language  
model. *Health Data Science*, 5:0221, 2025.
- Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T., Cao, Y.,  
and Narasimhan, K. Tree of thoughts: Deliberate problem  
solving with large language models. *Advances in neural*  
*information processing systems*, 36:11809–11822, 2023.
- Ye, W., Liu, J., Cao, D., Yang, W., and Liu, Y. When  
llm meets time series: Can llms perform multi-step  
time series reasoning and inference. *arXiv preprint*  
*arXiv:2509.01822*, 2025.
- Yu, F., Zhao, H., and Zhou, T. Ts-reasoner: Aligning time se-  
ries foundation models with llm reasoning. *arXiv preprint*  
*arXiv:2510.03519*, 2025.
- Yu, H., Guo, P., and Sano, A. Zero-shot ecg diagnosis  
with large language models and retrieval-augmented gen-  
eration. In *Machine learning for health (MLAH)*, pp.  
650–663. PMLR, 2023.
- Yu, H., Guo, P., and Sano, A. Ecg semantic integrator (esi):  
A foundation ecg model pretrained with llm-enhanced car-  
diological text. *arXiv preprint arXiv:2405.19366*, 2024.
- Zhang, D., Yuan, Z., Yang, Y., Chen, J., Wang, J., and Li,  
Y. Brant: Foundation model for intracranial neural signal.  
*Advances in Neural Information Processing Systems*, 36:  
26304–26321, 2023.
- Zhang, Y., Xia, T., Han, J., Wu, Y., Rizos, G., Liu, Y.,  
Mosuily, M., Ch, J., and Mascolo, C. Towards open  
respiratory acoustic foundation models: Pretraining and  
benchmarking. *Advances in Neural Information Process-*  
*ing Systems*, 37:27024–27055, 2024.
- Zhao, H., Zhang, X., Wei, J., Xu, Y., He, Y., Sun, S., and  
You, C. Timeseriesscientist: A general-purpose ai agent

715 for time series analysis. *arXiv preprint arXiv:2510.01538*,  
716 2025a.

717 Zhao, Y., Kang, J., Zhang, T., Han, P., and Chen, T. Ecg-  
718 chat: A large ecg-language model for cardiac disease  
719 diagnosis. In *2025 IEEE International Conference on*  
720 *Multimedia and Expo (ICME)*, pp. 1–6. IEEE, 2025b.  
721

722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769

## **A. Common Physiological Waveforms**

This appendix briefly recaps representative physiological waveform modalities referenced in the main text.

**Electrocardiogram (ECG).** Records cardiac electrical depolarization and repolarization. It is a quasi-periodic signal (typically 0.05–50Hz) where specific morphological components (P, QRS, T waves) map directly to conduction pathway integrity and myocardial state. Clinical datasets typically consist of 10-second, 12-lead recordings with structured diagnostic annotations (Wagner et al., 2020), often linked to auxiliary clinical data like comorbidities (Johnson et al., 2023).

**Photoplethysmogram (PPG).** Measures peripheral blood volume changes via optical absorption. It encodes hemodynamic parameters including heart rate and vascular tone (typically 0.5–2Hz) and is coupled to the ECG via Pulse Transit Time. PPG quality is highly sensitive to sensor placement, skin tone, and motion artifacts (Bent et al., 2020; Fine et al., 2021), often requiring task-specific window lengths (e.g., shorter for HR, longer for respiratory modulation) (Allen, 2007).

**Electroencephalogram (EEG).** Captures synaptic potentials aggregated across the scalp. It is characterized by spectral dominance across distinct bands—delta (~0.5–4Hz) to gamma (>30Hz)—and transient graphoelements linked to cognitive or pathological brain states (Buzsáki, 2006). Signal characteristics are heavily influenced by the electrode montage (e.g., 10-20 system) and neural dynamics (Craik et al., 2019).

**Electromyogram (EMG).** Measures muscle motor unit activation via electrical potentials. It reflects neuromuscular recruitment through burst patterns and high-frequency spectral signatures (20–500Hz) (De Luca, 1997). Unlike ECG, EMG lacks a universally standardized montage, often adapting sensor placement to specific muscle groups following guidelines like SENIAM (Hermens et al., 2000).

**Phonocardiogram (PCG).** Records mechanical heart sounds (acoustic vibrations). It validates the mechanical response to electrical excitation, emphasizing valve closure timing (S1/S2) and turbulent flow (murmurs).

**Other physiological waveforms.** Many additional physiological signals appear in modern monitoring, including arterial blood pressure (ABP), respiratory waveforms such as airflow or impedance-based respiration, capnography for end-tidal CO<sub>2</sub>, electrodermal activity (EDA), and peripheral temperature.

## **B. Call to Action for Verifiable Physiological Waveform Reasoning**

This appendix translates the paper’s position into concrete, stakeholder-specific steps. The goal is to move the community from end-to-end prediction toward verifiable, replayable, budget-aware waveform reasoning, where every decision is supported by localized evidence and reproducible checks.

### **B.1. Who should do what**

**Benchmark and dataset builders.** (i) Publish benchmarks that mirror deployment: artifacts, missing channels, device/site heterogeneity, and label uncertainty. (ii) Provide an official evaluation harness that consumes a minimal Verification Ledger; phase it in via (a) bonus scoring for ledger outputs, then (b) main-leaderboard requirement. At minimum, the ledger should record: claim type, time/lead span, tool-derived measurements (with units), checker/tool version, and provenance. (iii) Release perturbation generators with fixed seeds and documented ranges, covering noise, baseline wander, motion artifacts, resampling, channel drop, plus at least one deployment-realistic failure mode. (iv) Define tiered audit budgets and protocols that cap inspection time, tool calls, and optional human review, enabling realistic and comparable evaluation.

**Model builders.** (i) Expose structured evidence outputs as first-class predictions, including signal quality, localized events with time and channel indices, and unit-consistent measurements. (ii) Implement a Plan–Act–Verify loop that can request missing evidence, invoke deterministic tools, perform consistency checks, and abstain or escalate when evidence is insufficient. (iii) Report verifiability metrics together with task performance, including traceability to evidence, replay success, robustness under perturbations, and calibrated abstention under budgets.

**Tooling community.** (i) Maintain versioned, deterministic tool suites for quality assessment, delineation, interval measurement, morphology descriptors, and artifact detection. (ii) Enforce provenance logging as part of the tool interface, including tool identity, version hash, parameters, input window, channels, and failure status. (iii) Provide conformance tests and reference outputs so that ledger checks are replayable across platforms and environments.

**Clinical and deployment partners.** (i) Define operational policies for abstention and escalation, specifying what constitutes

sufficient evidence for each workflow setting. (ii) Specify audit budgets and oversight requirements, including what evidence must be shown for acceptance and what triggers human review. (iii) Provide structured feedback traces that identify which evidence was missing, unreliable, or misleading, enabling targeted improvement of tools and models.

**B.2. Benchmark roadmap as a community plan**

This roadmap is a staged plan that coordinates multiple stakeholders. Dataset builders provide the benchmark and perturbations, tool developers provide replayable checks, model builders produce ledger-backed decisions.

**Phase 1: Ledger-first leaderboards, months 0 to 6.** (i) Benchmark builders add a ledger requirement to existing tasks and publish a validator that checks schema compliance and replayability. (ii) Tooling community releases a reference tool suite and conformance tests so that reported measurements can be reproduced. (iii) Model builders submit systems that produce structured verification objects, explicit checks, and abstention decisions under stated budgets.

**Phase 2: Artifact and heterogeneity challenge sets, months 6 to 18.** (i) Benchmark builders release controlled perturbation suites and device or site shift splits with documented protocols. (ii) Community evaluates robustness together with sensitivity, focusing on unsafe high-confidence errors and whether abstention reduces such failures under budgets.

**Phase 3: End-to-end workflow episodes, months 18 and beyond.** (i) Benchmark builders publish multi-step episodes that mirror real workflows, including quality gating, evidence extraction, measurement, decision-making, and escalation. (ii) Clinical partners define workflow-specific evidence requirements and escalation policies so success reflects operational safety constraints. (iii) Model builders integrate System 1 evidence extraction with System 2 Plan–Act–Verify, demonstrating replayable reasoning under realistic audit budgets.

**C. Verification Ledger and Verification Objects: Schema and Case Studies**

This appendix specifies a minimal, auditable *Verification Ledger* and concrete *verification objects*. The ledger is the evaluation interface: an episode-level conclusion is credited only if it can be re-derived from (i) localized verification objects and (ii) replayable procedures recorded in the ledger.

**Minimal Verification Ledger fields**

<b>Inputs</b>	Record identifier; record hash; sampling rate; lead set; acquisition metadata; task query.
<b>Pointer</b>	Canonical scope pointer for all evidence: record_id; signal_space raw or preproc; leads; sample_start/sample_end; fs; preproc_id; view_id.
<b>Preprocessing</b>	preproc_id; ordered ops such as filtering, resampling, normalization; parameters; determinism settings; output hash.
<b>Views</b>	view_id; scope pointer lead subset and time window; preproc_id; render parameters.
<b>Verification Objects</b>	Signal-quality summaries; localized events; unit-consistent measurements with units and scope pointer; conflicts with scope and resolution pointers.
<b>Tool Run</b>	tool name; semantic version; code hash; runtime or container hash; parameters; seed; input pointers; status; failure mode; output hash.
<b>Deterministic Checks</b>	name; input object ids or pointers; rule; threshold; tolerance; result; tool-run provenance reference.
<b>Decision</b>	Final label or text; abstain flag; escalate flag; supporting object/check pointers; unresolved conflicts; constraints_checked.
<b>Budget</b>	Tier id; caps max views, max tool calls, max wall-clock, optional human review; realized usage views, tool calls, wall-clock, audit steps.

**Verification object and check templates**

<b>Pointer</b>	{id, record_id, signal_space, leads, s_start, s_end, fs, preproc_id, view_id}
<b>Preproc</b>	{id, input_pointer, ops, params, determinism, output_hash}
<b>View</b>	{id, scope_pointer, preproc_id, render_params}
<b>Signal Quality Summary</b>	{id, metric, value, unit, scope_pointer, threshold, flag, method_ref, provenance_ref}
<b>Localized Event</b>	{id, type, scope_pointer, confidence, attributes, notes, method_ref, provenance_ref}
<b>Unit Consistent Measurement</b>	{id, name, value, unit, scope_pointer, method_ref, uncertainty, tolerance, provenance_ref}
<b>Tool Run (Provenance)</b>	{id, tool, version, code_hash, runtime_env_hash, params, seed, input_pointers, status, failure_mode, output_hash}
<b>Conflict</b>	{id, scope_pointer, sources object_ids, description, severity, resolution_status, resolution_pointers}
<b>Deterministic Check</b>	{id, name, input_object_ids, rule, threshold, tolerance, result, provenance_ref}
<b>Decision</b>	{id, label, text, abstain, escalate, supports, constraints_checked, unresolved_conflicts}

**Audit rules: What makes an episode verifiable**

<b>Traceability</b>	Every major claim and the final decision reference explicit pointers to verification objects and deterministic checks. Free-form rationales without pointers are not creditable.
<b>Replayability</b>	Any quantitative statement is reproducible by re-running the logged tool run using the recorded code hash, runtime hash, parameters, and seed under the recorded scope pointer, matching within the stated tolerance.
<b>Unit and Scope Discipline</b>	Measurements carry units and an explicit scope pointer lead and time in sample indices plus fs and method_ref. Ambiguity in unit, scope, or method breaks replayability.
<b>Determinism Discipline</b>	All tools that can affect values must log code hash, runtime or container hash, and seed; checks must log tolerance. Missing determinism metadata downgrades the episode to non-verifiable.
<b>Quality Gating</b>	If a quality summary fails for a required scope pointer, the episode must query an alternative scope, abstain or escalate, or log a deterministic check that justifies validity under the limitation.
<b>Conflict Handling</b>	If conflicts are detected across tools, views, or scopes, the ledger logs scope, sources, and resolution pointers. Unresolved high-severity conflicts cannot support a definitive decision.
<b>Budget Compliance</b>	Credited decisions must satisfy the declared tier caps for views, tool calls, and wall-clock; budget overruns are logged and may be disqualified per protocol.

**C.1. Case Study**

This case study illustrates uncertainty-aware safety in a verifiable episode. The system first records objective signal-quality evidence (e.g., lead dropout and low SNR) as verification objects tied to the evaluated time/lead scope. It then runs two independent rhythm analyses whose outputs disagree at high severity (AF vs. sinus), and logs this disagreement as an unresolved conflict object. Two deterministic checks—(i) a quality-gating rule requiring required leads and SNR above a threshold, and (ii) a conflict policy that defers when high-severity disagreements remain unresolved—are applied to these objects. Because the quality gate fails and the conflict policy triggers deferral, the episode returns an explicit abstain+escalate decision (“unknown rhythm”), with a support list that points back to the quality metrics, the conflict, and the specific checks that forced deferral. This makes the safety behavior auditable: a reviewer can replay the checks and verify that abstention follows directly from recorded evidence under the stated policy and budget.

Case Study: Low-quality recording triggers abstention

<b>Inputs</b>	Record <b>rec_789</b> sampled at <b>500 Hz</b> . Leads: I, II, III, aVR, aVL, aVF, V1–V6. Task: determine rhythm and triage.
<b>View</b>	Single view over the full <b>10-second</b> segment (all leads), using the default preprocessing.
<b>Tools executed</b>	A signal-quality module, two independent rhythm analysis tools (A and B), and a ledger checker.
<b>Quality evidence</b>	<b>Lead I dropout detected</b> $\Rightarrow$ quality failure. Estimated <b>SNR = 6 dB</b> (required $\geq$ <b>10 dB</b> ) $\Rightarrow$ low-quality flag.
<b>Rhythm hypotheses</b>	Tool A indicates <b>atrial fibrillation (AF)</b> ; Tool B indicates <b>sinus rhythm</b> . These outputs form a <b>high-severity unresolved disagreement</b> .
<b>Deterministic checks</b>	<b>Quality gating</b> : required leads present and $\text{SNR} \geq 10 \text{ dB} \Rightarrow$ <b>FAIL</b> . <b>Conflict policy</b> : unresolved high-severity disagreement $\Rightarrow$ <b>DEFER</b> .
<b>Decision</b>	Return <b>unknown rhythm</b> ; <b>abstain</b> from definitive rhythm labeling and <b>escalate</b> for further review. Supports: lead dropout, low SNR, AF vs sinus disagreement, quality-gating failure, conflict-policy deferral.
<b>Budget</b>	Tier: standard. Limits: <b>1 view</b> , <b>4 tool calls</b> , <b>2 s</b> wall-clock. Usage: 1 view, 4 tool calls, 0.9 s wall-clock.

*Note:* Detailed object identifiers, hashes, and runtime metadata follow Appendix C and are omitted here for readability.

935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989