Real-time Deployment of Unsupervised Changepoint Detection for Pediatric Cardiac Monitoring

Denys Herasymuk [1], Dmytro Lopushanskyy [1], Danny Eytan [2], Mjaye L. Mazwi [1,3]

[1] Seattle Children's Hospital; [2] Technion, Israel Institute of Technology; [3] University of Washington

Correspondence:
Dmytro.Lopushanskyy@seattlechildrens.org

Seattle Children's®
HOSPITAL • RESEARCH • FOUNDATION

Why this matters

ECG is one of the most widely used signals for continuous physiologic monitoring, yet clinically meaningful rhythm and morphology changes can be buried in hours of waveform data. In the ICU, these changes may precede deterioration or mark transitions in patient state. Threshold alarms dominate bedside monitoring, yet more than 90% of ECG monitor alarms are reported as false or clinically insignificant. We propose moving from threshold monitoring to state monitoring by detecting sustained ECG structure changes that indicate meaningful physiologic transitions.

What we developed

We built and deployed a real-time ECG changepoint detection system for sustained changes in rhythm, frequency, and morphology.

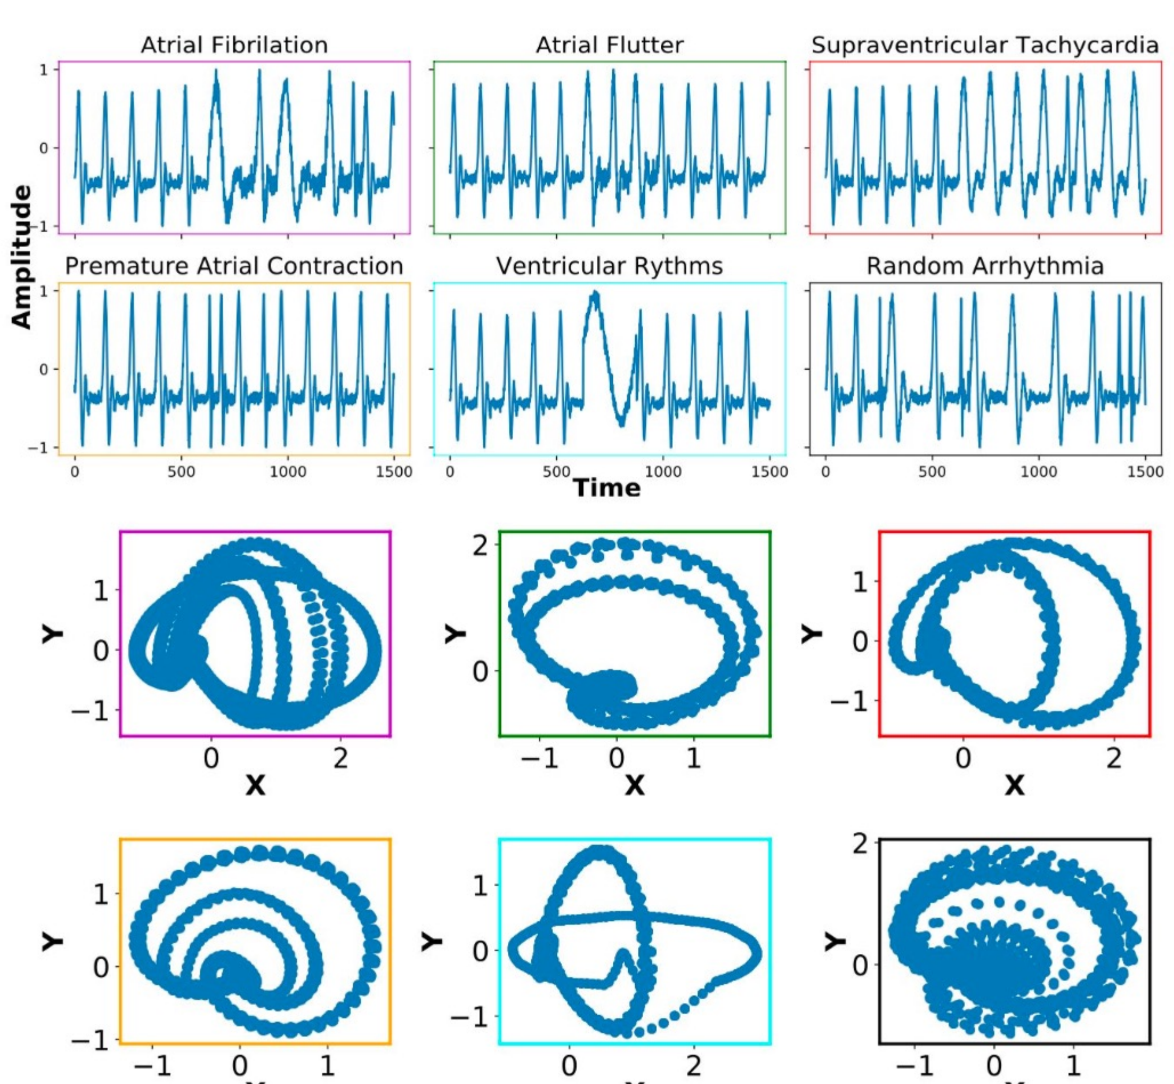

ECG segments are embedded as point clouds, where distinct waveform patterns produce distinct geometries [1].

Real-time Deployment Architecture

Live bedside ECG waveforms are streamed from Philips monitors into a Google Cloud-based inference pipeline. Pub/Sub routes the waveform stream, Dataflow jobs process incoming signals, BigTable stores signal and outputs, and Cloud Run executes the model.

Expert Validation

Candidate ECG segments were annotated in Label Studio by expert cardiologists. A valid changepoint is defined as a clear, consistent transition lasting ≥4 beats. Each segment was reviewed by three clinicians. Their annotations were merged into conservative consensus windows for evaluation.

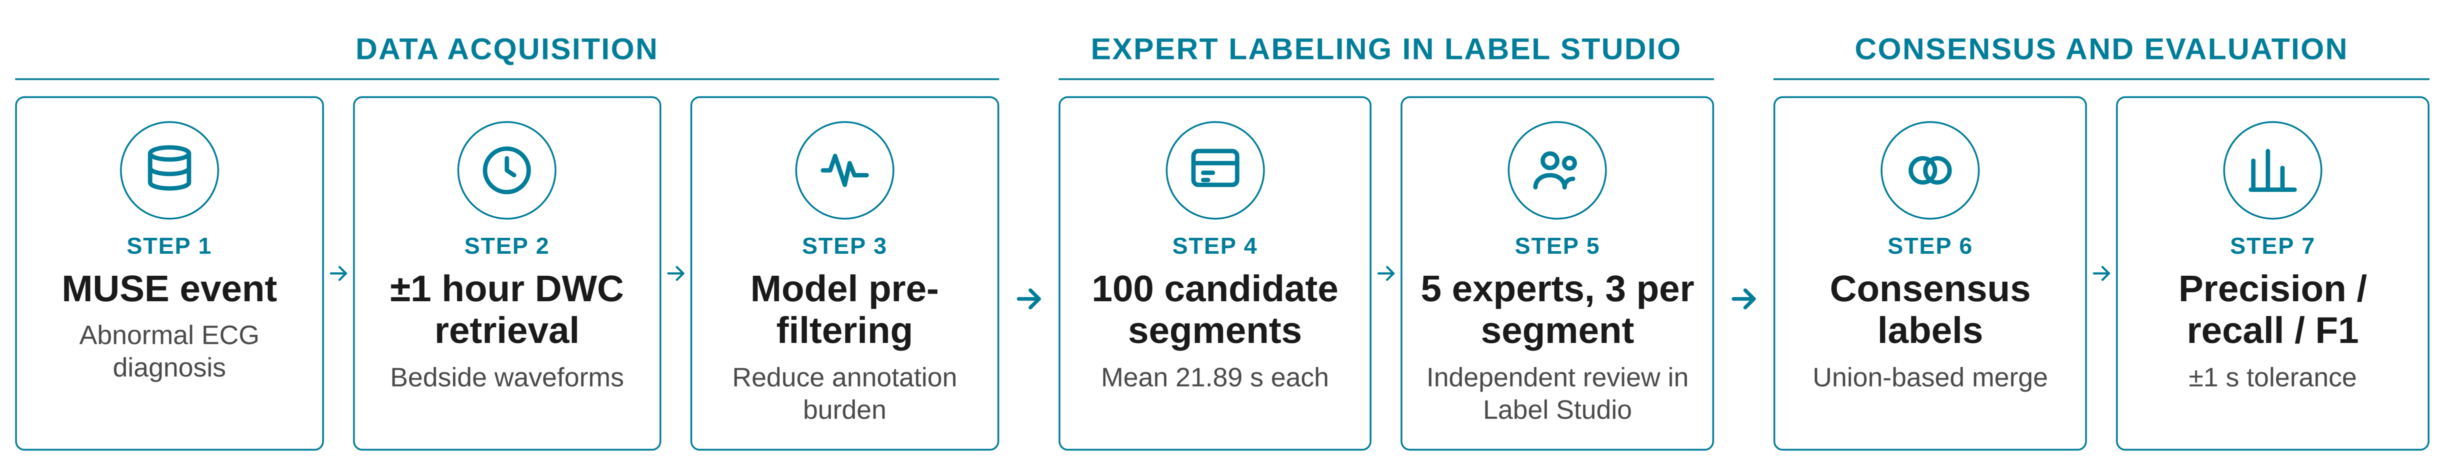

Method

ECG windowing over local beat structure

The ECG stream is divided into overlapping windows, each containing several beats. At each time point τ, the surrounding window is split into left and right empirical distributions, μ_l and μ_r, representing ECG structure before and after that point:

$$\mu_l(\tau) = \frac{1}{h}\sum_{t=\tau-h}^{\tau-1}\delta_{X_t}, \quad \mu_r(\tau) = \frac{1}{h}\sum_{i=\tau}^{\tau+h-1}\delta_{X_t}$$

Artifact handling

Artifact, dropout, baseline wander, and electrode noise are removed to normalize the signal.

Point cloud representation

Each ECG window is embedded as a point cloud [1]:

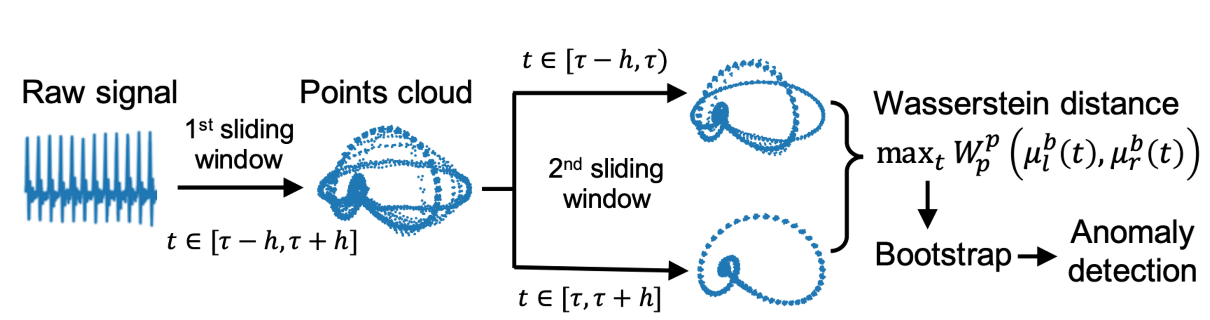

Wasserstein change score

The model compares the before and after point clouds using Wasserstein distance, which measures how much one distribution must be "moved" to match the other. A larger distance indicates a greater change in ECG structure at that time point.

$$W_p^p(\mu,\nu) = \left(\inf_{\gamma\in\Pi(\mu,\nu)}\int_{M*M}\|x-y\|^p d\gamma(x,y)\right)^{1/p}$$

Bootstrap detection threshold

A moving block bootstrap estimates the expected score under no meaningful change. Scores above this threshold are surfaced as changepoints.

Patient-specific baseline

The system estimates patient-specific beat structure and adjusts window parameters for pediatric ECG variability.

Results

We retrieved 45 bedside ECG waveforms from Philips DWC, each 2 hours at 250 Hz, totaling 90 hours, 81 million samples, and approximately 324,000 beats.

Model pre-filtered segments were reviewed by five expert cardiologists, with each segment independently labeled by three clinicians.

Against consensus labels, the model achieved **91.43% precision, 72.22% recall, and 80.70% F1 score.**

Statistic	Value
Total segments for CPD evaluation	100
Avg. task duration	21.89 s
Avg. task length (data points)	5,472
Avg. beats per task (≈ 250 pts/beat)	≈22 beats
Total beats across evaluation set	≈2,200
Total true change points (ground truth)	**126**

Prospective Evaluation

The model now runs live on real-world waveform streams from ICU environments in silent trial mode, generating research-only alerts. Alerts open the signal viewer at the detected changepoint for ECG review, signal quality assessment, and chart-based comparison. Outputs are stored for evaluation but are not visible to clinical staff and do not affect care.

Scope: 24-month silent trial, ~250 ICU patients per month, with up to 6,000 unique inpatient admissions per month, hospital-wide.

References

1. Shvetsov N, Buzun N, Dylov DV. Unsupervised non parametric change point detection in quasi periodic signals. arXiv:2002.02717. 2020.
2. Truong C, Oudre L, Vayatis N. Selective review of offline change point detection methods. Signal Processing. 2018.
3. Shao X, Zhang X. Testing for change points in time series. Journal of the American Statistical Association, 2010.
4. Li F, Runger GC, Tuv E. Supervised learning for change-point detection. International Journal of Production Research, 2006.