Abstract (GTTI 2026)

Title: A Probabilistic Dynamic Bayesian Framework for Explainable Anomaly Detection in PPG Signals

Authors: Anna Lo Grasso[1], Antonio Affanni[1], Lucio Marcenaro[2], Carlo Regazzoni[2], Roberto Rinaldo[1], Pamela Zontone[2]

Affiliations: [1]University of Udine, [2]University of Genova

Photoplethysmography (PPG) signals are widely used in wearable cardiovascular monitoring, enabling continuous and non-invasive assessment of physiological states in everyday life. In safety-critical settings such as driving, early detection of cardiovascular anomalies, including arrhythmic events, sudden autonomic changes, or stress-induced alterations, is essential to prevent dangerous situations and ensure timely intervention. Reliable anomaly detection in PPG signals therefore plays a key role in supporting driver well-being, long-term health monitoring, and preventive healthcare applications. Achieving robust, real-time detection in wearable devices, however, requires methods that operate efficiently, generalize across individuals, and provide interpretable outputs that clinicians and system designers can readily understand and use.

In this work, we address these needs by proposing a fully probabilistic and explainable framework for anomaly detection in PPG signals. Our method models cardiovascular dynamics through a Dynamic Bayesian Network (DBN) and performs inference using a Markov Jump Particle Filter (MJPF) (see Fig.1), enabling joint modeling of discrete region changes and continuous waveform evolution. Normal PPG segments, acquired exclusively with a custom necklace-type wearable device, are mapped into a structured latent space using Growing Neural Gas, which identifies a set of superstates representing typical cardiovascular patterns. Within each superstate, locally linear dynamical models capture short-term temporal evolution. During online inference, the MJPF estimates both discrete mode transitions and continuous latent states, generating complementary anomaly indicators based on likelihood consistency, innovation, and Kullback–Leibler divergence between predicted and observation-driven superstate distributions (see Fig. 2).

Our probabilistic framework is evaluated on unseen normal data and on publicly available datasets reflecting stress or exercise (walking, cycling) as well as pathological rhythms (atrial fibrillation). Despite using far fewer samples than deep models that require extensive training data, the proposed method achieves strong discrimination across all scenarios, with AUC values reaching up to 1.00. Comparisons with a convolutional autoencoder and a recent subject-aware representation-learning approach show that it offers competitive or superior performance while maintaining interpretability and substantially lower computational demands. Its key characteristics, hierarchical interpretability, data efficiency, and lightweight inference, make our approach a practical and robust solution for real-time anomaly detection in wearable PPG monitoring and, more generally, for embedded health-sensing applications.

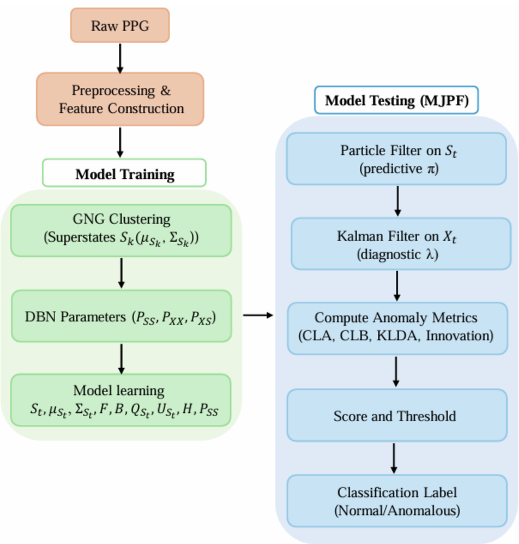

Fig.1 Block diagram of the proposed method.

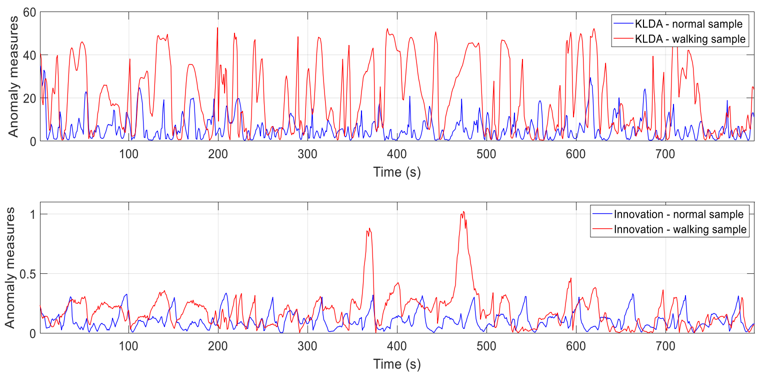

Fig.2 Temporal profiles of anomaly scores for a normal PPG signal (blue) and an anomalous PPG signal recorded during walking (red). (a) Kullback-Leibler Divergence Anomaly (KLDA); (b) Innovation.