ECGencode: Compact and computationally efficient DL feature encoder for ECG signals

Lennert Bontinck, Karel Fonteyn, Dirk Deschrijver, Tom Dhaene

lennert.bontinck@ugent.be, karel.fonteyn@ugent.be, dirk.deschrijver@ugent.be, tom.dhaene@ugent.be

IDLab - SUMO

Time series team

Expert Systems with Applications

Problem Statement

Deep learning (**DL**) models are becoming essential for **ECG analysis**, identifying patterns that are **undetectable by traditional methods**. These models enable tasks like **AFib detection during NSR** and **new-onset AFib risk prediction**. However, the absence of **ECG-specific models** leads current solutions to rely on general architectures, which are often **computationally expensive**, **difficult to configure**, lack **intrinsic visualisation**, and struggle to **generalize** to **out-of-distribution (OOD)** data. As interest in deploying these models on **medical edge devices** grows, the need for an **ECG-specific DL feature encoder** that is **easy to configure**, **computationally efficient**, **interpretable**, and **generalisable** while maintaining **state-of-the-art performance** becomes clear.

Proposed Solution

We present **ECGencode**, an **ECG-specific DL feature encoder** that overcomes key limitations in existing DL models:

- **ECG Specific, Compact, and Expert-Inspired Architecture**
  Tailored for ECG data, inspired by the FBCSP technique in EEG processing.
- **Computationally Efficient and Low-Parameter Design**
  Optimized for inference on resource-constrained medical edge devices.
- **Intuitive Model Configuration and Interpretable Architecture**
  Simplifies model configuration and offers clear latent space interpretability.
- **Versatile By Design**
  Easily adaptable to various tasks, including classification and regression.

Input Data

12 x 5,000 -> 12 x 5,000 x 1

ECG Input
Input shape: 12 x 5000
Output shape: 12 x 5,000 x 1

ECGencode

12 x 5,000 x 1 -> 4 x 40 x 4

Temporal Convolution
Input shape: 12 x 5,000 x 1
Output shape: 12 x 625 x 16

Spatial Convolution
Input shape: 12 x 625 x 16
Output shape: 4 x 625 x 16

Feature Convolutions
Input shape: 4 x 625 x 16
Output shape: 4 x 40 x 4

Extension(s) & Output

4 x 40 x 4 -> 640 -> 2

Classifier
Input shape: 4 x 40 x 4
Flatten shape: 640
Output shape: 2

Input Data Properties

The most commonly used input is a **12-lead, 10-second ECG**, widely available in clinical settings.

Recordings are often sampled at **500Hz**, resulting in a high-dimensional signal of **12 x 5000**.

Temporal Properties

2D convolution is applied to each lead independently, with varying kernel lengths to capture both **high** and **low-frequency** details.

Striding is used to downscale the signal, reducing **computational complexity** for subsequent layers.

Spatial Properties

Captures **spatial correlations** across all leads to create **augmented leads**.

Utilises **depthwise convolutions** to ensure information from different **frequency filters** remains distinct.

Feature Properties

Extracts **deep learning features** using efficient **depthwise separable convolutions** with configurable **complexity**.

ECGencode's output preserves the **time-ordered structure across augmented leads**.

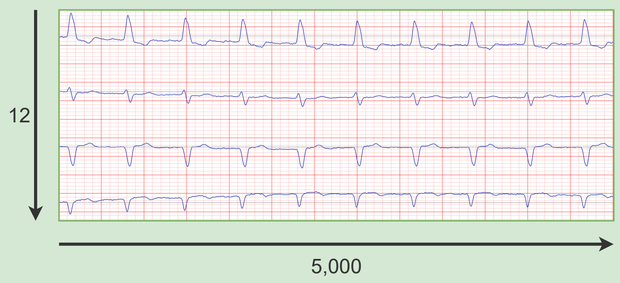
Extension Properties

The output of ECGencode can be **extended** with **LSTM** or other functionality for improved **temporal modelling**.

Applicable to a wide range of tasks, including **binary classification**, **multi-label classification**, and **regression**.

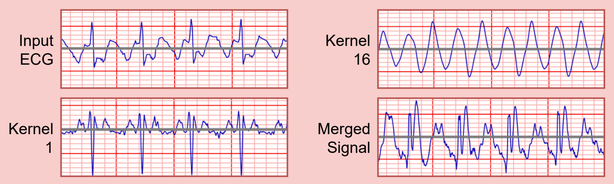
ECG signal

Standard 12-lead 10-second ECG



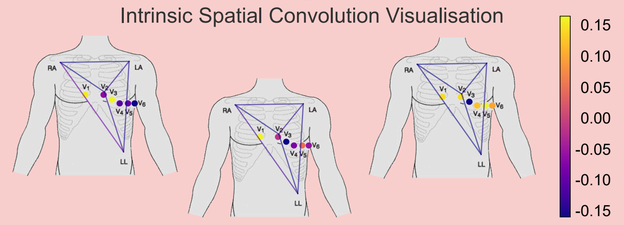
Frequency Filtering

Intrinsic Temporal Convolution Visualisation



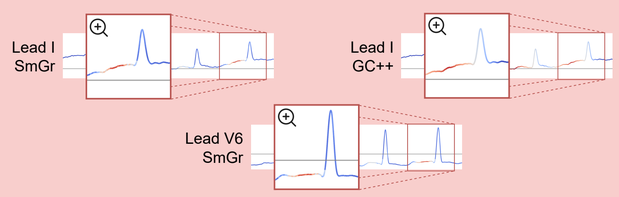
Lead Importance

Intrinsic Spatial Convolution Visualisation



Segment Importance

Post-Hoc Model Attention Visualisation



Easily extendable

ECGencode is made Versatile By Design

LSTM
Input shape: 4 x 40 x 4
Output shape: 40 x 16

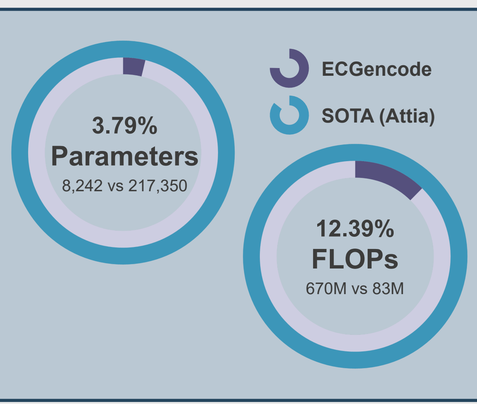
→

Multilabel
Input shape: 40 x 16
Output shape: 71

	# Params	# FLOPs
ECG signal	0	0M
Frequency Filtering	5,440	81M
Lead Importance	816	1M
Segment Importance	704	1M
Easily extendable	1,282	0M
ECGencode	8,242	83M

Results

- **AFib Detection during NSR and New-Onset AFib Prediction**
  ECGencode achieves significant reductions in computational cost while matching (test fold) or outperforming (OOD data) SOTA models.
- **Multi-label Arrhythmia Classification with LSTM Extension**
  With intuitive configuration, ECGencode matches benchmark performance and achieves substantial reductions in trainable parameters, demonstrating versatility across different tasks.



Future work

- **Medical Validation of Visualisations**
  Assess the clinical usefulness of the proposed visualisations through an easy UI.
- **AutoML Toolbox**
  Develop an AutoML toolbox that includes ECGencode and various extensions to optimise for a specific task while controlling computational cost.

imec embracing a better life

IDLab INTERNET & DATA LAB

GHENT UNIVERSITY

Artificial Intelligence Flanders

Flanders AI Research Program