

A Wearable Device for Online and Long-Term ECG Monitoring

Marco Longoni¹, Diego Carrera¹, Beatrice Rossi²,
 Pasqualina Fragneto², Marco Pessione², Giacomo Boracchi¹

¹ Politecnico di Milano, Milan, Italy

² STMicroelectronics, Agrate Brianza, Italy

marco.longoni@mail-polimi.it, diego.carrera@polimi.it, beatrice.rossi@st.com,
 pasqualina.fragneto@st.com, marco.pessione@st.com, giacomo.boracchi@polimi.it

Abstract

We present a prototype wearable device able to perform online and long-term monitoring of ECG signals, and automatically detect anomalous heartbeats such as arrhythmias. Our solution is based on user-specific dictionaries which characterizes the morphology of normal heartbeats and are learned every time the device is positioned. Anomalies are detected via an optimized sparse coding procedure, which assesses the conformance of each heartbeat to the user-specific dictionary. The dictionaries are adapted during online monitoring, to track heart rate variations occurring during everyday activities. Perhaps surprisingly, dictionary adaptation can be successfully performed by transformations that are user-independent and learned from large datasets of ECG signals.

1 Introduction

In the last few years there has been a flourishing of smart-watches and bracelets able to monitor health parameters. Due to their limited sensing apparatus, these devices can not be used to ease the transitioning from hospital to home/mobile monitoring: as an example, the heart rate is often the only ECG measure provided. We present the *Bio2Bit-Dongle*, a prototype wearable device developed by STMicroelectronics that features a proper sensing apparatus and implements advanced machine learning algorithms to analyze the heartbeats shapes of each specific user, and detect anomalous heartbeats that can be due to potentially dangerous arrhythmias, acquisition errors and device displacement. These algorithms have been designed to run online on a low-power dongle in a long-term monitoring scenario.

The Bio2Bit-Dongle (Figure 1) is composed of the *Bio2Bit Move* [Gentile *et al.*, 2017], which is plugged on an adjustable chest strap and acquires and transmits via Bluetooth Low Energy (BLE) the ECG signals, and a low-power *dongle* (Section 3) that analyzes in real time the received signals. In particular, the dongle adopts a data-driven model describing the heartbeats' morphology of the current user, and detects

as anomalous any heartbeat that does not conform to that. This anomaly detection algorithm [Carrera *et al.*, 2016] has a very low computational and battery requirements, and as such can be steadily executed on the considered low-power dongle. Online monitoring allows to raise timely alerts at each anomalous heartbeat, or to store only anomalies for further inspection.

Since the heartbeat morphology depends on the electrodes locations [Hoekema *et al.*, 2001], the user-specific model has to be learned every time the Bio2Bit Move is positioned. To this purpose, a few minutes of ECG signals acquired in resting conditions are enough to learn the user-specific model on any smartphone or powerful smart-watch. Such a customization is not enough to enable long-term monitoring, since during everyday activities the heart rate changes and heartbeats get consequently transformed. Thus, the dongle estimates the current user heart rate, and adapts the learned model accordingly as in [Carrera *et al.*, 2017].

2 ECG Monitoring Algorithms

In what follows we describe the algorithms we designed to perform ECG monitoring on the Bio2Bit-Dongle.

2.1 Anomaly Detection

Pursuing a universal approach to anomaly detection [Chandola *et al.*, 2009], we learn a model [Bengio *et al.*, 2013] providing meaningful representations of the normal heartbeats for each specific user, and we detect as anomalous any heartbeat that does not conform to the learned model. In particular, we learn for each user u a dictionary $D_u \in \mathbb{R}^{p \times n}$ such that any normal heartbeat $\mathbf{s} \in \mathbb{R}^p$ can be well approximated by a sparse representation [Bruckstein *et al.*, 2009], i.e. a linear combination of few columns of D_u . We learn D_u from a training set of heartbeats acquired in resting conditions using the KSVD algorithm [Aharon *et al.*, 2006]. Thus, normal heartbeats admits a representation $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{s} \approx D\mathbf{x}$ and \mathbf{x} is sparse, namely the number of its nonzero coefficients is small. This modeling assumption corresponds to learning a union of low dimensional subspaces of \mathbb{R}^p where normal heartbeats live. To detect [Carrera *et al.*, 2016] whether a heartbeat \mathbf{s} is anomalous or not we solve the *sparse coding*

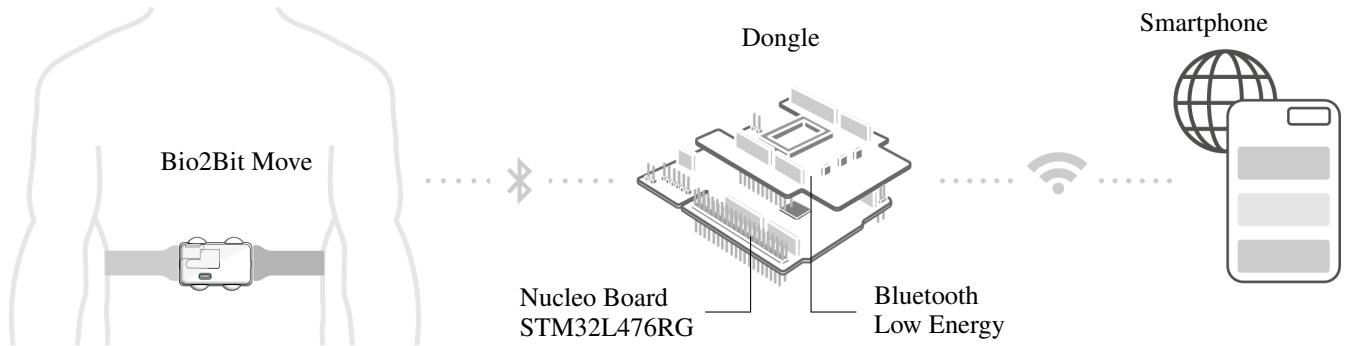


Figure 1: The user wears Bio2Bit Move (on a chest strap) and the dongle (in a bracelet, or hanged on a belt), which communicate over a Bluetooth connection. The *training phase* requires a few minutes of ECG signals acquired in resting conditions, which the dongle segments into heartbeats. These represent the training set that is sent to the user’s smartphone, where the dictionary D_u and the threshold γ_u are learned. These are sent back to the dongle together with adapted versions by built-in user-independent transformations. During *online monitoring*, the dongle receives the ECG signal from by the Bio2Bit Move, segments each heartbeats and estimates the heart rate. For each heartbeat, it performs the sparse coding w.r.t. the dictionary that matches the current heart rate. Finally the dongle computes the reconstruction error, and raises an alert when it is above the selected threshold.

problem

$$\mathbf{x} = \arg \min_{\tilde{\mathbf{x}}} \|\mathbf{s} - D\tilde{\mathbf{x}}\|_2, \quad \text{such that } \|\tilde{\mathbf{x}}\|_0 \leq \kappa, \quad (1)$$

where $\|\cdot\|_0$ denotes the number of nonzero coefficients. This problem is NP-Hard and has to be addressed by greedy algorithms, such as Orthogonal Matching Pursuit (OMP) [Pati *et al.*, 1993]. We assess the conformance of \mathbf{s} to D_u by analyzing the reconstruction error $e(\mathbf{s}) = \|\mathbf{s} - D\mathbf{x}\|_2$ and detect \mathbf{s} as anomalous when $e(\mathbf{s}) > \gamma_u$. The threshold γ_u is the minimum distance of anomalies to the learned subspaces of \mathbb{R}^p containing normal heartbeats and controls the amount of false positives and is estimated during configuration.

To allow online execution on the dongle, we designed a highly optimized variant of the OMP [Carrera *et al.*, 2017], meant for undercomplete dictionaries (having $n < p$), which are the most effective at modeling heartbeats. The experiments in [Carrera *et al.*, 2016] show that our anomaly-detection algorithm achieves excellent performance on the benchmark dataset MIT-BIH [Moody and Mark, 2001] and on ECG signals acquired using the Bio2Bit Move.

2.2 Domain adaptation

Heartbeats get transformed when the heart-rate increases, and the dictionary D_u learned in resting conditions might not be appropriate to monitor the user u in everyday activities. Similarly, γ_u might not maintain a constant false positive rate. An extended training set including heartbeats at different heart rates would allow learning multiple pairs (D_u, γ_u) , but for safety reasons, training data can be only acquired in resting conditions. Thus, we pursue a transfer-learning approach [Pan and Yang, 2010], and adapt D_u and γ_u to track heart rate variations.

We learn transformations [Carrera *et al.*, 2017] to adapt both D_u and γ_u from a large dataset of ECG signals from several users [Jager *et al.*, 2003]. To preserve the structure of subspaces where normal heartbeats lives, we learn a set of linear transformations for adapting dictionaries to different

heart rates. Perhaps surprisingly, our user-independent transformations can successfully adapt user-specific dictionaries and thresholds, even when the sensing apparatus changes. In [Carrera *et al.*, 2017], we show that our dictionary-adaptation algorithm outperforms other transfer-learning solutions [Shekhar *et al.*, 2013], and can be executed online on the dongle.

3 Practical Aspects

The Bio2Bit Move is a prototype device placed in a 55x40x14 mm envelope, even though the final product sizes can be significantly reduced to yield a non-invasive wearable device. Its battery allows to steadily sense and transmit via BLE the ECG signals for up to 20 hours. The dongle has been replaced by a development board NUCLEO-L476RG [ST, 2016] which embeds an ultra-low-power microcontroller unit (MCU) based on the high-performance ARM Cortex-M4 32-bit RISC core, operating at a frequency of 80 MHz. This MCU is a likely reference for low-power dongles. The board is also equipped by a BLE evaluation board (X-NUCLEO-IDB05A1) to communicate with the Bio2Bit Move.

To determine whether our algorithm can be executed online on the Bio2Bit-Dongle in a long term monitoring scenario, we measure the execution time and the power consumption. Our optimized OMP [Carrera *et al.*, 2017] requires on average 1.360 ms to compute the reconstruction error of an heartbeat, which is way below the acquisition time of each heartbeat, and is 48% less than the computing time of a standard OMP implementation. Dictionary adaptation has no overheads, since dictionaries can be pre-transformed during training and stored in memory. The average current absorbed by the dongle during 20 minutes of online monitoring is 10.01 mA. Thus, when the dongle is equipped with a 592 mWh battery such as the one adopted on the Bio2Bit Move, it has 16 hours monitoring autonomy.

References

- [Aharon *et al.*, 2006] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing over-complete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- [Bengio *et al.*, 2013] Yoshua Bengio, Aaron Courville, and Pierre Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [Bruckstein *et al.*, 2009] Alfred M Bruckstein, David L Donoho, and Michael Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM review*, 51(1):34–81, 2009.
- [Carrera *et al.*, 2016] Diego Carrera, Beatrice Rossi, Daniele Zambon, Pasqualina Fragneto, and Giacomo Boracchi. Ecg monitoring in wearable devices by sparse models. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD)*, pages 145–160, 2016.
- [Carrera *et al.*, 2017] Diego Carrera, Beatrice Rossi, Pasqualina Fragneto, and Giacomo Boracchi. Domain adaptation for online ecg monitoring. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, pages 775–780, 2017.
- [Chandola *et al.*, 2009] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):15, 2009.
- [Gentile *et al.*, 2017] Paolo Gentile, Marco Pessione, Antonio Suppa, Alessandro Zampogna, and Fernanda Irrera. Embedded wearable integrating real-time processing of electromyography signals. In *Proceedings of Eurosensor*, volume 1, page 600, 2017.
- [Hoekema *et al.*, 2001] Rudi Hoekema, Gérard JH Uijen, and Adriaan Van Oosterom. Geometrical aspects of the interindividual variability of multilead ecg recordings. *IEEE Transactions on Biomedical Engineering*, 48(5):551–559, 2001.
- [Jager *et al.*, 2003] Franc Jager, Alessandro Taddei, George B Moody, Michele Emdin, G Antolič, Roman Dorn, Ales Smrdel, Carlo Marchesi, and Roger G Mark. Long-term st database: a reference for the development and evaluation of automated ischaemia detectors and for the study of the dynamics of myocardial ischaemia. *Medical and Biological Engineering and Computing*, 41(2):172–182, 2003.
- [Moody and Mark, 2001] George B Moody and Roger G Mark. The impact of the mit-bih arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50, 2001.
- [Pan and Yang, 2010] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.
- [Pati *et al.*, 1993] Yagyensh Chandra Pati, Ramin Rezaiifar, and Perinkulam Sambamurthy Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of the Asilomar Conference on Signals, Systems and Computers*, pages 40–44, 1993.
- [Shekhar *et al.*, 2013] Sumit Shekhar, Vishal M Patel, Hien V Nguyen, and Rama Chellappa. Generalized domain-adaptive dictionaries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 361–368, 2013.
- [ST, 2016] ST. <http://www.st.com/en/evaluation-tools/nucleo-l476rg.html>, 2016.