

Robust Heartbeat Detection from Multimodal Data via CNN-based Generalizable Information Fusion

B. S. Chandra*, C. S. Sastry and S. Jana

Abstract—Objective: Heartbeat detection remains central to cardiac disease diagnosis and management, and is traditionally performed based on electrocardiogram (ECG). To improve robustness and accuracy of detection, especially, in certain critical-care scenarios, the use of additional physiological signals such as arterial blood pressure (BP) has recently been suggested. There, estimation of heartbeat location requires information fusion from multiple signals. However, reported efforts in this direction often obtain multimodal estimates somewhat indirectly, by voting among separately obtained signal-specific intermediate estimates. In contrast, we propose to directly fuse information from multiple signals without requiring intermediate estimates, and thence estimate heartbeat location in a robust manner. **Method:** We propose as a heartbeat detector, a convolutional neural network (CNN) that learns fused features from multiple physiological signals. This method eliminates the need for hand-picked signal-specific features and ad hoc fusion schemes. Further, being data-driven, the same algorithm learns suitable features from arbitrary set of signals. **Results:** Using ECG and BP signals of PhysioNet 2014 Challenge database, we obtained a score of 94%. Further, using two ECG channels of MIT-BIH arrhythmia database, we scored 99.92%. Both those scores compare favorably with previously reported database-specific results. Also, our detector achieved high accuracy in a variety of clinical conditions. **Conclusion:** The proposed CNN-based information fusion (CIF) algorithm is generalizable, robust and efficient in detecting heartbeat location from multiple signals. **Significance:** In medical signal monitoring systems, our technique would accurately estimate heartbeat locations even when only a subset of channels are reliable.

Index Terms—Convolutional neural networks, heartbeat detection, multimodal data fusion, electrocardiogram, blood pressure.

I. INTRODUCTION

CARDIOVASCULAR diseases (CVDs) are a leading cause of death worldwide [1], and their management has become a global imperative. In certain scenarios, including intensive care unit (ICU) monitoring, cardiac conditions are continuously assessed for possible worsening. There, automated anomaly detection and alert generation are anticipated to improve timely intervention. In this direction, the first step often involves accurate heartbeat detection based on related physiological signals such as electrocardiogram (ECG) and arterial blood pressure (BP) signals (Figure 1a). Detection based on multiple signals is generally more accurate than that based on individual signals [2], because quality of one signal

*B. S. Chandra is with the Department of Electrical Engineering, Indian Institute of Technology Hyderabad, India (e-mail: bschandra@iith.ac.in).

C. S. Sastry is with the Department of Mathematics, Indian Institute of Technology Hyderabad, India.

S. Jana is with the Department of Electrical Engineering, Indian Institute of Technology Hyderabad, India.

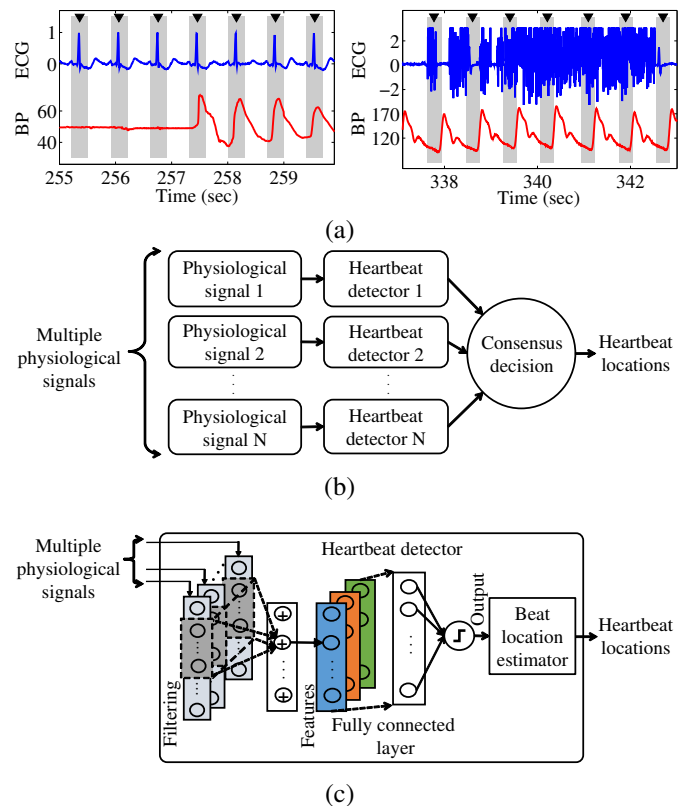


Fig. 1. (a) ECG and BP signals from PhysioNet 2014 Challenge database with annotations at the top [6]: left—segment from record 1071 with clean ECG and missing BP signals, right—segment from record 2800 with corrupted ECG and clean BP signals; (b) Traditional multimodal heartbeat detector [3]–[5]; (c) Proposed multimodal detector: CNN-based information fusion (CIF).

could be poor, when other signals maintain high quality. However, a typical multisignal detector [2]–[5], simply combining multiple single-signal detections, does not systematically exploit inter-signal correlation. To fill this gap, we propose as a heartbeat detector a convolutional neural network (CNN) that makes a final decision by directly fusing information from multiple signals, without needing intermediate detections.

Automated estimation of heartbeat location from ECG signals has received considerable attention. Various techniques including signal differencing, filterbanks, wavelet transform and Hilbert transform followed by threshold comparison have been reported [7]–[13]. To fulfill real-time requirements, adaptive thresholding has been employed [14]–[16]. Not surprisingly, detection performance is generally improved when multi-channel ECG data are used instead of single channel [9], [17], [18]. Simultaneously, attempts have also been made to

estimate heartbeat locations from BP signal [19]. In general, the aforementioned techniques exploit temporal dependency in each signal. Recently, detection algorithms combining information from ECG and BP signals have been reported [2], [3]. Multimodal detectors generally function as shown in Figure 1b [3]–[5]. An intermediate estimate of heartbeat location is obtained from each signal using hand-picked signal-specific features. Based on such intermediate estimates, a consensus decision is heuristically made, albeit without directly considering inter-signal dependency. An advanced method following this paradigm considers upto three channels each of ECG and blood pressure signals [18]. However, exploitation of inter-signal dependency assumes significance because various clinical conditions (e.g., arrhythmias and morphological anomalies) and non-clinical conditions (e.g., noise and pacemakers) affect different signal modes differently. Noting this, a recent work employs a hidden semi-Markov model (HSMM) to statistically relate ECG and BP features [20]. While this method desirably accounts for temporal dependency, HSMM could be restrictive as a model, and the said features are still hand picked for specific signals and not generalizable. In contrast, we propose as a heartbeat detector a CNN, shown in Figure 1c, which learns suitable filters to extract features, temporally fusing information from multiple signals. Finally, a fully connected network maps these features to possible heartbeat locations. As the filter coefficients and the network weights are learned from training data, the proposed CNN-based information fusion (CIF) algorithm generalizes to arbitrary set of signals.

Generally, a detector performance is evaluated on a public database, which is divided into two sets, used for training and testing. Notionally, the algorithm should be tuned based on the training set, and the performance be recorded on the test set. However, this framework would mimic the real life, only when the detector operates on unforeseen but representative data. In this vein, the “PhysioNet 2014 Challenge” (hereafter “Challenge”) made an extensive training dataset of multiple signals (including ECG and BP) publicly available, while keeping the representative test set hidden. The Challenge held a time-bound competition among detectors, and identified the leaders reporting scores up to 93.64% [3], [18]. Fortunately, the portal continues to score submitted algorithms [4], [5], and the proposed CIF algorithm achieved a score of 94%, which compares favorably with hitherto reported scores, indicating its practical significance. Further, as alluded earlier, our CIF algorithm generalizes to an arbitrary set of physiological signals. To illustrate this, we moved to two channels of ECG. However, in absence of a scoring service that uses undisclosed test data, we turned to MIT-BIH arrhythmia database, based on which performance of a number of existing algorithms has already been reported [7]–[18]. As the test data are not hidden, a performance index here may not directly represent the anticipated real-life accuracy, although performance comparison can still be illuminating. Adopting the Challenge scoring rule, and respectively using only one ECG channel and both the channels, our CIF algorithm achieved scores of 99.89% and 99.92%, which ranked among the highest reported so far. Here, we hasten to add that multi-database evaluation has been attempted before. Specifically, a previously mentioned

algorithm operating on upto three channels each of ECG and blood pressure signals has been evaluated on five databases, including the Challenge and MIT/BIH databases [18]. However, as that algorithm depends on features specific to ECG and blood pressure signals, its applicability is accordingly restricted to records consisting of the associated signals only. In contrast, the proposed algorithm generalizes to arbitrary set of signals and records.

In summary, our CIF algorithm for heartbeat detection

- 1) fuses information from multiple physiological signals;
- 2) is data-driven and generalizable;
- 3) is efficient and robust for different sets of signals;
- 4) performs favorably versus signal-specific methods.

The rest of the paper is organized as follows. In Section II, we describe the databases used and the proposed heartbeat detector. Experimental results are provided in Section III. Finally, Section IV concludes the paper.

II. MATERIALS AND METHODS

We now formalize the problem of heartbeat detection from multiple signals. A number K (≥ 2) of physiological signals $x_{t_1:t_2}^{(1)}, x_{t_1:t_2}^{(2)}, \dots, x_{t_1:t_2}^{(K)}$ from a subject are simultaneously recorded from time t_1 to time t_2 . Our task is to estimate heartbeat locations $\tau_1, \tau_2, \dots, \tau_m$ (m denoting the number of heartbeats) from those signals. Next we describe the databases, outline our CIF heartbeat detector, and present the performance evaluation strategy.

A. Databases

We make use of PhysioNet 2014 challenge and MIT-BIH arrhythmia databases for performance evaluation [3], [6]. In each database, a patient record is annotated with consensus heartbeat locations marked by two or more cardiologists. In particular, location of the R-peak of an ECG is taken as the heartbeat location. Specific details of the databases are described below.

1) *Challenge database*: The Challenge aims at benchmarking heartbeat detectors operating on multimodal data including ECG and BP [3]. The publicly available training set consists of 200 patient records, while the hidden test set has 210 records. Signals from those records have varying duration and sampling frequency varying between 250 and 360 Hz. Also available is a public portal that evaluates and benchmarks the performance of submitted algorithms. Specifically, a score is assigned based on the classification accuracy on the entire database consisting of 410 records including training and hidden test data.

2) *MIT-BIH arrhythmia database*: The MIT-BIH arrhythmia database, often used for evaluating heartbeat detection [7]–[18] and related algorithms [21], consists of 48 public records, each comprising half-hour two-channel ECG signals sampled at 360Hz. Of those, 25 records, numbered 200–234, contain less common but clinically significant cardiac abnormalities, which we used for training. As customary [7]–[18], the algorithmic performance was evaluated on the entire database consisting of an additional 23 records, numbered 100–124, containing normal beats and a representative set of routine arrhythmias. To maintain consistency, we followed the Challenge rules for final scoring.

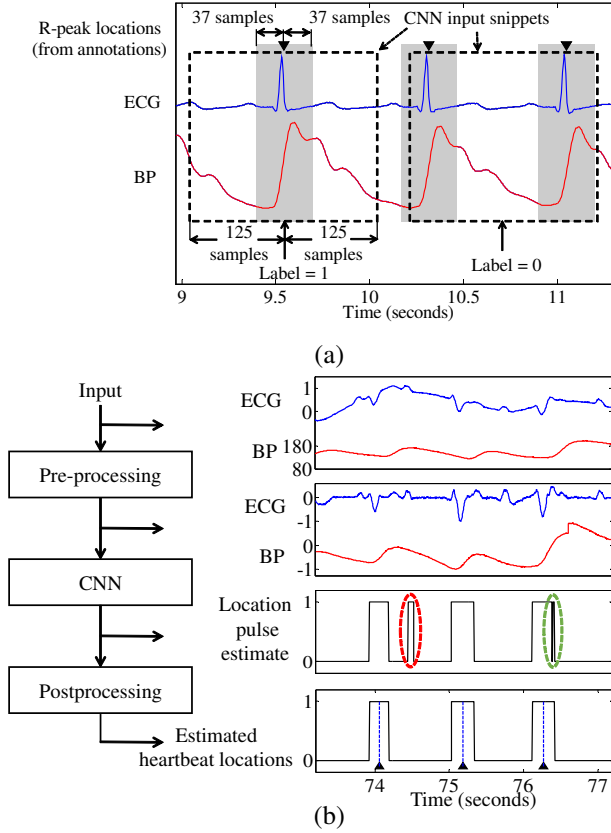


Fig. 2. (a) CNN input snippets, labeled based on location pulses (shaded gray); (b) Block diagram of the proposed algorithm with intermediate outputs.

B. Detection based on CNN-based Information Fusion (CIF)

While localizing heartbeats, an estimated location is deemed accurate if it is no further than 150ms from the corresponding annotated location, in accordance with ANSI/AAMI EC38 and EC57 standards [22]. Thus, a 300ms location pulse can be formed around the annotated location, and an accurate location estimate falls within such a pulse. As sampling rate varies within as well as across databases, for our analysis, all signals were up or downsampled to a standard rate of 250Hz. At this rate, the location pulse spanned 75 samples, depicted by gray shade in Figures 1a, 2a, 4a-d and 5a-f.

To systematically fuse information from multiple signals, we made use of CNNs. Recall that a CNN adds convolution layers between the input layer and (the hidden as well as) the output layer of a fully-connected neural network [23]. We restricted to one convolution layer consisting of linear filters, whose outputs were passed through a sigmoid nonlinearity onto the fully connected network (see Figure 1c). Assuming K signals, each aforementioned filter consists of K component 1D filters, each with a common length L operating on a specific signal, such that the overall filter output (which is nothing but the desired feature set) is obtained by summing component filter outputs. Formally, assume p filters altogether, and denote by $\{h_l^{(k,j)}\}_{l=0}^{L-1}$ the coefficients of the k -th ($1 \leq k \leq K$) 1D channel of the j -th ($1 \leq j \leq p$) filter. The corresponding 1D filter output is given by $y_n^{(j)} = \sum_{l=0}^{L-1} x_{n-l}^{(k)} h_l^{(k,j)}$. Hence, the

overall output of the j -th filter is given by $y_n^{(j)} = \sum_{k=1}^K y_n^{(k,j)}$. The feature set consists of $y_n^{(j)}$, where $j = 1, 2, \dots, p$ and n varies as follows. At a time, we consider a window of length M of signal vector $x_{t:t+M-1}^{(k)}$ for the k -th channel with some start time t . Further, we avoid zero-padding, so that the j -th filter output $y_n^{(j)}$ is obtained only for $n = t+L-1, \dots, t+M-1$ amounting to $M - L + 1$ samples. Thus the length of the feature vector produced by all p filters equals $p(M - L + 1)$. Interestingly absence of zero-padding allows a short filter to produce a larger number of outputs than a long filter. The aforementioned filter coefficients as well as network weights were optimized using multimodal training data such that the optimized values essentially capture inter-signal as well as temporal dependence. The said optimization was performed using back propagation algorithm under the cross entropy cost function [23].

Our CNN input comprises of snippets of $M = 251$ samples (duration 1.004s) each of constituent signals. Each snippet is assigned a label as shown in Figure 2a. Specifically, if the middle sample of any snippet falls within any location pulse, the said snippet is labeled 1. It is labeled 0, otherwise. For training, we considered successive snippets with an overlap of 150 samples (duration of about 0.6s). As expected, we obtained more training snippets with label 0 than those with label 1 in the process. Subsequently, we maintained class balance by removing randomly chosen snippets with label 0, as required. Alongside original snippets, we also populate the training dataset with modified snippets with only one component signal intact and the rest zeroed so as to enable heartbeat detection even when only one signal is available. The overall training set consists of 551154 snippets for the Challenge database, and 451050 snippets for the MIT-BIH arrhythmia database.

To use CIF efficiently, we preprocessed the data as follows. Baseline wander, common in ECG signals, was removed by passing the signal through a cascade of median filters of length 50 and 150 (respective duration 200ms and 600ms) [24]. Next, range of each signal was normalized to $[-1, 1]$. Specifically, each non-overlapping window of 500 samples (duration 2s) was scaled by the maximum magnitude. Finally, time synchronization was ensured. For example, in the Challenge database, we compensated for the lag between the ECG and BP signals, which others had found to be around 200ms [25]. We optimized CNN architecture and weights at the training and validation phase. To this end, we used k (integer)-fold randomized cross validation on the training set [26].

Test snippets were generated from a test record by moving a window of length 251 one sample at a time. Such snippets are presented sequentially to the trained and optimized CNN, and the output sequence was viewed as location pulse estimate. We propose a two-step postprocessing for such initial estimate to improve accuracy. Firstly, we considered inverted pulses, i.e., intervals with label 0 preceded and succeeded by regions with label 1. Very short inverted pulses lasting less than 3 samples (duration 10ms) arise likely due to noise. Such an inverted pulse can be located within the region circled in green in Figure 2b. We propose to remove each such spurious inverted

TABLE I
PERFORMANCE OF VARIOUS CNN ARCHITECTURES (BOLDFACE INDICATES OPTIMIZED VALUES)

Architecture	No. of convolution layers = 1, No. of hidden layers = 0												No. of convolution layer = 1 No. of filters = 2 Filter length (L) = 20 No. of hidden layers = 1		
	No. of filters = 1				No. of filters = 2				No. of filters = 3				No. of hidden nodes		
	Filter length (L)				Filter length (L)				Filter length (L)				No. of hidden nodes		
	1	20	150	250	1	20	150	250	1	20	150	250	100	200	500
Se_{avg} (%)	86.2	96.6	95.7	86.3	90.8	96.9	95.9	88.2	92.5	95.9	95.8	88.1	92.7	93.2	92.1
PPV_{avg} (%)	92.0	97.7	97.5	91.0	93.7	97.5	97.2	92.2	93.7	97.5	97.2	92.2	95.5	95.6	95.1
Se_{gross} (%)	85.6	96.3	95.6	86.1	90.5	96.7	95.3	88.0	91.9	95.3	95.3	87.6	92.0	92.6	91.1
PPV_{gross} (%)	93.3	97.8	97.8	93.6	95.0	97.7	97.4	94.4	96.2	98.1	97.4	93.2	96.2	96.3	95.8
Score (%)	89.3	97.1	96.7	89.2	92.5	97.2	96.4	90.7	93.6	96.7	96.4	90.3	94.1	94.4	93.5
Exec. time (ms/beat)	3.1	3.0	3.1	3.2	3.6	3.6	3.6	3.6	3.8	3.7	3.8	3.9	51.5	51.5	52.4

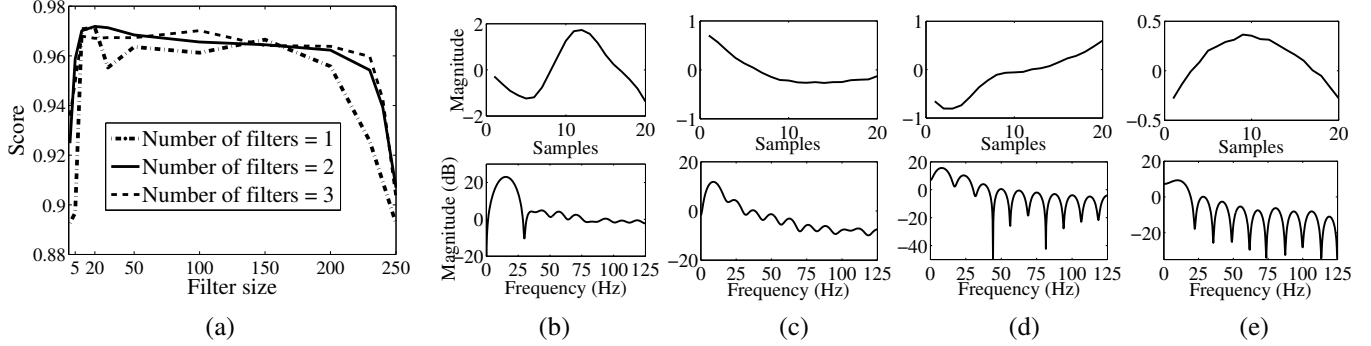


Fig. 3. (a) Performance variation with filter size and number of filters (peak score with two filters of length 20); Learned filters (top) and corresponding frequency response (bottom) for (b) ECG channel in filter-1, (c) BP channel in filter-1, (d) ECG channel in filter-2, (e) BP channel in filter-2.

pulse by switching the label from 0 to 1. Secondly, recall that a bona fide location pulse is 75 samples long and separation between two heartbeats is generally greater than 100 samples (corresponding to 150 beats/min). Therefore, an estimated pulse lasting less than 50 samples (duration 200ms) is likely spurious. For an example, refer to the region circled in red in Figure 2b. We propose to remove such a pulse by switching the label from 1 to 0. The resulting pulses were taken as the final estimates of location pulses, and the midpoints of such pulse estimates were taken as estimated heartbeat locations.

C. Performance evaluation

For performance evaluation using either database, we adopted the Challenge evaluation scheme, which is described below. Recall that a location estimate was deemed correct if it fell within 150ms of the annotated location [22]. For the i -th record, we counted the numbers of correctly detected beats (true positives, TP_i), missed beats (false negatives, FN_i) and spurious detections (false positives, FP_i), and hence calculated sensitivity $Se_i = TP_i / (TP_i + FN_i)$ and positive predictive value (PPV) $PPV_i = TP_i / (TP_i + FP_i)$. Subsequently, recordwise average sensitivity $Se_{avg} = \sum_{i=1}^N Se_i / N$ and average PPV $PPV_{avg} = \sum_{i=1}^N PPV_i / N$ were computed, where N denotes the total number of records. We also calculated gross sensitivity $Se_{gross} = TP / (TP + FN)$ and gross PPV $PPV_{gross} = TP / (TP + FP)$, where $TP = \sum_{i=1}^N TP_i$, $FN = \sum_{i=1}^N FN_i$ and $FP = \sum_{i=1}^N FP_i$. Finally, the algorithm was assigned an overall performance score of $(Se_{avg} + PPV_{avg} + Se_{gross} + PPV_{gross}) / 4$.

TABLE II
PERFORMANCE COMPARISON ON PHYSIONET 2014 CHALLENGE HIDDEN DATASET (HIGHEST VALUES ARE BOLDFACED)

Algorithm	Se_{avg}	PPV_{avg}	Se_{gross}	PPV_{gross}	Score
Galeotti <i>et al</i> 2015 [27]	91.08	86.96	92.74	87.37	89.53
DeCooman <i>et al</i> 2015 [28]	89.59	89.62	90.74	90.15	90.02
Antink <i>et al</i> 2015 [29]	89.40	90.70	91.02	91.87	90.70
Johnson <i>et al</i> 2015 [25]	92.61	89.03	95.07	89.30	91.50
Pangerc <i>et al</i> 2015 [18]	93.86	91.57	95.65	93.48	93.64
Ding <i>et al</i> 2016 [4]	89.36	87.15	91.06	87.09	88.66
Rankawat <i>et al</i> 2016 [5]	91.60	88.85	92.74	90.39	90.89
Proposed CIF algorithm	92.85	94.29	93.41	95.47	94.00
CIF algorithm (only ECG)	85.72	89.70	87.07	91.20	88.42
CIF algorithm (only BP)	67.00	88.66	70.11	94.26	80.00

III. EXPERIMENTAL RESULTS

Now we present the experimental results and compare the performance of the proposed CIF algorithm with existing heartbeat detectors. To this end, we made use of the ECG and BP signals of PhysioNet 2014 Challenge database and later two ECG channels of MIT-BIH arrhythmia database.

A. Results for Challenge Database

We first optimized the CNN for the Challenge database and evaluated its performance.

CNN Optimization: We optimized the CNN architecture as well as weights. For the sake of simplicity, we considered shallow architectures with single convolution and at most one hidden layers. Using two channels, ECG and BP, of the Challenge training set, any CNN architecture with optimized weights was given a 10-fold randomized cross validation score




Fig. 4. R-peak detection with (a) missing data in BP channel; (b) missing data in ECG channel; (c) noisy ECG channel; (d) paced beats.

TABLE III
PERFORMANCE COMPARISON WITH OTHER METHODS ON MIT-BIH ARRHYTHMIA DATASET (SUPERIOR VALUE IN BOLDFACE)

Algorithm	No. of channels	Total beats	TP	FP	FN	Se_{gross}	PPV_{gross}	Se_{avg}	PPV_{avg}	Score
Hamilton <i>et al.</i> (1986) [7]	1	109267	108927	248	340	99.69	99.77	99.70	99.76	99.73
Afonso <i>et al.</i> (1999) [17]	2	91283	90909	406	374	99.59	99.56	99.61	99.53	99.57
Christov: Algorithm-1 (2004) [9]	2	110050	109756	215	294	99.73	99.80	99.75	99.80	99.77
Christov: Algorithm-2 (2004) [9]	2	110050	109810	239	240	99.78	99.78	99.80	99.78	99.79
Arzeno <i>et al.</i> (2008) [14]	1	109453	109099	405	354	99.68	99.63	–	–	–
Zhang <i>et al.</i> (2009) [10]	1	109510	109297	204	213	99.80	99.81	99.80	99.81	99.81
Ghaffari <i>et al.</i> (2009) [11]	1	109428	109327	129	101	99.91	99.88	–	–	–
Ghaffari <i>et al.</i> (2010) [12]	1	109428	109367	89	61	99.94	99.91	–	–	–
Manikandan <i>et al.</i> (2012) [13]	1	109496	109417	140	79	99.93	99.87	99.93	99.86	99.90
Ding <i>et al.</i> (2014) [16]	1	109494	109360	134	73	99.93	99.88	99.93	99.88	99.91
Gutiérrez-Rivas <i>et al.</i> (2015) [15]	1	109949	109447	289	502	99.54	99.74	–	–	–
Pangerc <i>et al.</i> (2015) (channel-1) [18]	1	109494	109369	125	125	99.89	99.89	–	–	–
Pangerc <i>et al.</i> (2015) (both channels) [18]	2	109494	109380	92	114	99.90	99.92	–	–	–
Elgendi <i>et al.</i> (2017) [30]	1	109985	109775	82	247	99.78	99.93	99.78	99.92	99.85
CIF algorithm (channel-1)	1	109494	109322	64	172	99.84	99.94	99.84	99.95	99.89
Proposed CIF algorithm	2	109494	109422	103	72	99.93	99.91	99.94	99.91	99.92

[26], based on which various configurations were compared. First, considering no hidden layers, the number of filters and the common filter length were optimized. To simplify matters, only one-, two- and three-filter configurations were explored. Also, the filter length was varied from 1 to 250 (see Table I and Figure 3a). A short filter scores low because, while outputting many features, it considers only a short temporal span, and does not exploit temporal dependency efficiently. A long filter also scores low because, while observing for a significant duration, has only a few outputs to encode the temporal dependency into. In general, as the filter length increases, scores rise sharply, remain somewhat steady, and then fall sharply for each filter. Here, the plateau region indicates that a filter length taken from a wide interval performs satisfactorily. The peak scores for one-filter and two-filter cases are both achieved at a filter length of 20 and are nearly equal, while the score for three filters is somewhat less. Of the first two, we picked the latter because the plateau around the peak is relatively flatter, facilitating robust operation. Next, using the above two-filter configuration with filter length 20 as reference, a hidden layer was introduced, and the aforementioned score was recorded vis-à-vis varying number of hidden nodes in anticipation of possible improvement, which did not materialize (Table I). Also, with the introduction of the hidden layer, the execution time per beat, using MATLAB v.2014b run on a desktop computer with an Intel core i7 3.4 GHz 64-bit processor and 16 GB memory, increased by more than an order of magnitude. So, we persisted with the aforementioned two-filter

configuration without hidden layers.

Recall that each filter operates on the two-channel data (ECG and BP) and generates a feature vector that captures both the temporal and inter-signal dependencies. Let us now inspect the optimized filter characteristics in time as well as frequency domain as presented in Figures 3b-e. Here the frequency response corresponding to one (ECG or BP) channel of the filter is obtained by zeroing the other (BP or ECG) channel. Observe that the ECG channel of filter-1 learned to pass a 3db frequency band of approximately 9.25–21.25Hz. This passband generally matches the recommended frequency range for detecting QRS complexes [31], and approximates the passband of 10–24Hz of ECG-specific slope- and peak-sensitive filters [18]. Turning to time domain, the present impulse response bears similarity with Haar-like matched filters and mother wavelets that have been used in representing the shape of QRS complexes [12], [16]. In short, the ECG channel of our learned filter-1 appears to capture much of human intuition. On the other hand, the BP channel of learned filter-1 passes a narrower and lower frequency band of 4.15–14.16Hz with more pronounced sidelobes. Interestingly, this passband does not overlap with the passband of 1.2–3.5Hz of a filter used to detect the steep-slope segment of blood pressure signals [18]. Thus, while the learned filter-1 appears to corroborate the effectiveness of a slope- and peak-sensitive filter for the ECG channel, it does not for the BP channel. In general, the CNN presents an alternative paradigm, where two filters seem to play complementary roles. Specifically,

TABLE IV
RECORDWISE PERFORMANCE ON MIT-BIH ARRHYTHMIA DATABASE

Rec	Peaks	Using only channel-1			Using both channels			FP FN
		FP FN	Se	PPV	FP FN	Se	PPV	
100	2273	0 1	99.96	100	0 1	99.96	100	— —
101	1865	0 1	99.95	100	1 1	99.95	99.95	↑ —
102	2187	0 0	100	100	1 1	99.95	99.95	↑ ↑
103	2084	0 0	100	100	0 0	100	100	— —
104	2229	0 1	99.96	100	0 0	100	100	— ↓
105	2572	16 22	99.14	99.38	11 12	99.53	99.57	↓ ↓
106	2027	2 2	99.90	99.90	2 2	99.90	99.90	— —
107	2137	1 1	99.95	99.95	1 1	99.95	99.95	— —
108	1763	2 6	99.66	99.89	2 5	99.72	99.89	— ↓
109	2532	0 0	100	100	1 0	100	99.96	↑ —
111	2124	0 2	99.91	100	0 0	100	100	— ↓
112	2539	0 0	100	100	0 0	100	100	— —
113	1795	0 1	99.94	100	0 1	99.94	100	— —
114	1879	2 0	100	99.89	0 0	100	100	↓ —
115	1953	0 0	100	100	6 1	99.95	99.69	↑ ↑
116	2412	0 20	99.17	100	0 0	100	100	— ↓
117	1535	0 0	100	100	0 0	100	100	— —
118	2278	1 0	100	99.96	1 0	100	99.96	— —
119	1987	0 0	100	100	0 0	100	100	— —
121	1863	0 2	99.89	100	2 0	100	99.89	↑ ↓
122	2476	0 0	100	100	0 0	100	100	— —
123	1518	1 0	100	99.93	1 0	100	99.93	— —
124	1619	0 0	100	100	0 0	100	100	— —
200	2601	2 2	99.92	99.92	2 2	99.92	99.92	— —
201	1963	1 22	98.88	99.95	1 1	99.95	99.95	— ↓
202	2136	0 5	99.77	100	1 0	100	99.95	↑ ↓
203	2958	22 18	99.40	99.26	36 22	99.26	98.8	↑ ↑
205	2656	0 2	99.92	100	0 0	100	100	— ↓
207	1860	5 11	99.41	99.73	11 6	99.68	99.41	↑ ↓
208	2955	1 12	99.59	99.97	6 9	99.70	99.80	↑ ↓
209	3005	0 0	100	100	1 1	99.97	99.97	↑ ↑
210	2650	3 4	99.85	99.89	4 1	99.96	99.85	↑ ↓
212	2748	1 0	100	99.96	2 0	100	99.93	↑ —
213	3251	0 1	99.97	100	0 1	99.97	100	— —
214	2262	0 4	99.82	100	0 3	99.87	100	— ↓
215	3363	0 0	100	100	0 0	100	100	— —
217	2208	0 3	99.86	100	0 0	100	100	— ↓
219	2154	0 0	100	100	0 0	100	100	— —
220	2048	0 0	100	100	0 0	100	100	— —
221	2427	0 0	100	100	2 0	100	99.92	↑ —
222	2483	1 5	99.80	99.96	2 0	100	99.92	↑ ↓
223	2605	0 0	100	100	1 0	100	99.96	↑ —
228	2053	3 18	99.12	99.85	0 1	99.95	100	— ↓
230	2256	0 0	100	100	0 0	100	100	— —
231	1571	0 0	100	100	0 0	100	100	— —
232	1780	0 5	99.72	100	5 0	100	99.72	↑ ↓
233	3079	0 1	99.97	100	0 0	100	100	— ↓
234	2753	0 0	100	100	0 0	100	100	— —
Gross	109494	64 172	99.84	99.94	103 72	99.93	99.91	↑ ↓
Avg			99.84	99.95		99.94	99.91	

compared to the ECG channel of filter-1, filter-2 has a passband 3.66–12.38Hz of significantly lower frequencies and more pronounced sidelobes. Turning to the BP channel, filter-2 has lowpass response with a cutoff frequency of 15.88Hz, complementing the bandpass response of filter-1.

Heartbeat detection performance: As mentioned earlier, Challenge score reports are generated based on a public training set and a hidden test set, designed to reflect real-life complexities. In Table II, we report sensitivity and PPV values, both gross and average, as well as the overall score obtained by the proposed CIF as well as existing algorithms. Our overall score of 94.00% is significantly higher than 88.42% and 80.00%, respectively obtained using only ECG and only BP,

illustrates the benefit of fusion, and compares favorably with hitherto reported scores. Importantly, our heartbeat detector was found to be robust to various non-clinical conditions shown in Figures 4a-c, where one of the BP and the ECG signals is either missing or heavily corrupted. There, filter-1 and filter-2 appeared to prominently encode ECG and BP features, respectively, whose simultaneous use ensured robust detection. Further, consider paced beats, induced by pace maker implants, which are not genuine heartbeats, but conventional detectors often confuse those with normal beats [25]. Desirably, our detector correctly ignored paced beats, as shown in Figure 4d, as a pace maker leaves no corresponding BP signature.

B. Results for MIT-BIH Arrhythmia Database

Next we evaluate the proposed CIF algorithm on MIT-BIH arrhythmia database consisting of two-channel ECG records, and hence demonstrate its generalizability. We first optimized the CNN architecture and weights using 5-fold randomized cross validation on training records 200-234 (refer Section II-A2) following the steps outlined in Section III-A. The optimized CNN consisted of one filter of length 30 and had no hidden layers, and was used in our CIF heartbeat detector.

Heartbeat detection performance: In Table III, we provide comparative performance analysis of the proposed CIF algorithm vis-à-vis published results. Specifically, we make comparisons in terms of the total numbers of beats, true positives (TP), false positives (FP) and false negatives (FN), as well as sensitivity and PPV values, both gross and average (if available), and the overall score (evaluated using the Challenge formula). We also indicate how many channels each competing algorithm made use of. Notice that the scores, obtained by even the earliest methods, are high. Also, the performance superiority of 2-channel methods over 1-channel methods remains unclear. Those facts possibly indicate that most recorded signals are well behaved. Unsurprisingly, the overall advantage enjoyed by the proposed 2-channel CIF over 1-channel CIF and other competing algorithms was slender, as well. At the same time, compared to 1-channel CIF, the 2-channel CIF reduced FNs from 172 to 72, and increased FPs from 64 to 103, indicating that a recordwise performance comparison could be illuminating.

Such comparison is presented in Table IV, where increase (↑), decrease (↓) or no change (—) in FPs and FNs are indicated. Of the 48 records, 20 see no change at all, and a significant number of those witness minor changes. Let us now inspect the records producing major changes. First, we show in Figures 5a-c illustrative segments from respective records 116, 105 and 203, where 2-channel CIF corrected errors committed by 1-channel CIF with the help of channel-2. In record 116, FNs incurred using only channel-1 got corrected when channel-2 was also used, because those correspond to episodes of missing signal in channel-1 only. Similar phenomenon also occurs in records 201 and 228 (not shown). In record 105, FPs and FNs incurred by 1-channel CIF got rectified by 2-channel CIF, despite both channels being noisy. In record 203, tagged as ‘very difficult record, even for humans’ [32], correction

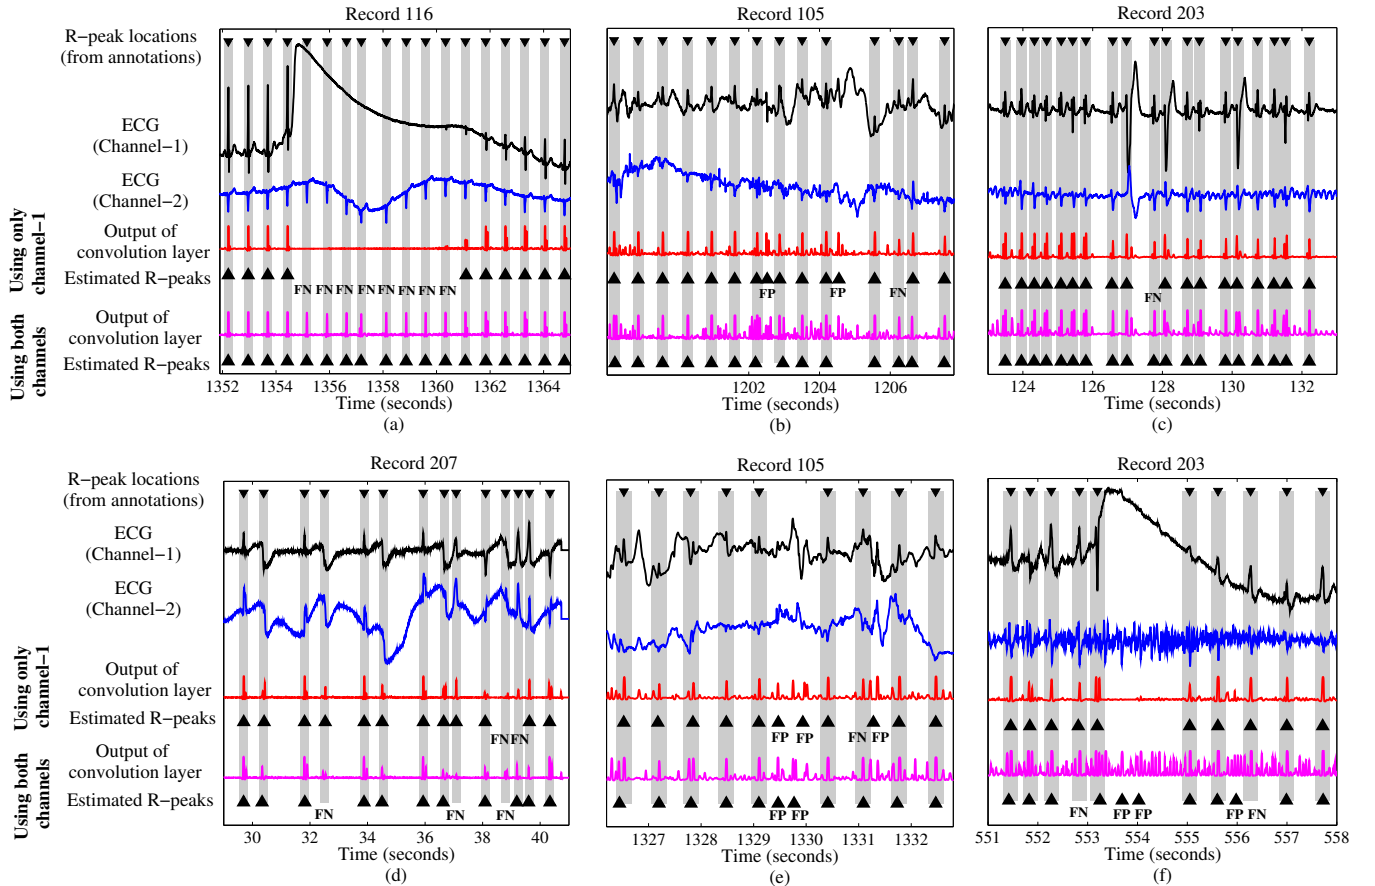


Fig. 5. Robust heartbeat detection using 2-channel CIF in various cases: (a) missing signal in channel-1, (b) noisy signal, (c) arrhythmic condition; Erroneous detection: (d)-(f) illustrative cases.

of a certain FN is shown. The aforementioned examples demonstrate the robustness of 2-channel CIF in scenarios of varying difficulty levels.

Despite its general success, our algorithm still fails in certain difficult circumstances. As illustration, we depict in Figures 5d-f segments from records 207, 105 and 203, where the intended results were not obtained. In record 207, tagged ‘extremely difficult’ [32], while one FN was corrected, one remained unrectified, and two new ones were introduced, because channel-2 appears to convey confusing information. In record 205, a heartbeat location estimate shifted slightly to correct both an FN and an FP. At the same time, two FPs remained uncorrected due to ‘high grade noise and artifact’ [32]. In the ‘very difficult’ record 203, a shift in location estimate introduced an FP and an FN. Those along side another FN and two FPs resulted because the channel-2 signal, buried in noise, distracted the 2-channel CIF algorithm. In summary, in a variety of circumstances, the proposed 2-channel CIF method tends to correct estimation errors made based on channel-1 using additional information from channel-2, except in certain difficult cases. In the process, while achieving a slightly higher overall score, the 2-channel CIF tended to also bring the four performance indicators, namely, Se_{avg} , PPV_{avg} , Se_{gross} and PPV_{gross} , closer in value to one another, compared to 1-channel CIF.

IV. CONCLUSION

In this paper, we proposed a CNN-based information fusion (CIF) algorithm for heartbeat detection from multiple physiological signals that systematically exploits both temporal and inter-signal dependencies. Specifically, our detector learns a set of linear filters to extract fused features, which are then mapped to estimated heartbeat locations. On the ECG and BP channels of the PhysioNet 2014 Challenge database, our CIF technique achieved a score of 94%, which is superior to hitherto reported scores. Further, on the MIT-BIH arrhythmia database of 2-channel ECG records, we achieved a score of 99.92%, which also compares favorably with competing scores. At this point, note that unlike the proposed technique, the existing methods do not generalize to arbitrary databases. For instance, algorithms applied to the Challenge database generally use hand-picked features specific to ECG and BP and do not generalize to 2-channel ECG of the arrhythmia database, and *vice versa*. One notable exception to the above is found in an aforementioned algorithm [18], which uses hand-picked features but allows upto three channels each of ECG and blood pressure signals, and hence can be applied to the Challenge, MIT/BIH and other databases consisting of ECG and BP records. In contrast, our algorithm is designed to generalize to records consisting of arbitrary physiological signals. The superior performance of our generalizable CIF

technique is especially remarkable, because it did not use human knowledge of any signal set, and is based only on features directly learnt from the database at hand. In addition to generalizability and improved accuracy, our method also demonstrated desirable robustness to various clinical anomalies and non-clinical distortions. Accordingly, we envisage CIF-based heartbeat detection being incorporated in medical devices monitoring multiple physiological signals. As a technical improvement, we plan to study the sensitivity-versus-PPV tradeoff by varying the CNN decision threshold (which has been kept constant in this paper), and optimize such threshold subject to operating constraints.

REFERENCES

- [1] World Health Organization, Fact sheet on CVDs, March 2017.
- [2] Q. Li, R. G. Mark, and G. D. Clifford, "Robust heart rate estimation from multiple asynchronous noisy sources using signal quality indices and a kalman filter," *Physiol. Meas.*, vol. 29, no. 1, p. 15, 2007.
- [3] I. Silva, B. Moody, J. Behar, A. Johnson, J. Oster, G. D. Clifford, and G. B. Moody, "Robust detection of heart beats in multimodal data," *Physiol. Meas.*, vol. 36, no. 8, p. 1629, 2015.
- [4] Q. Ding, Y. Bai, Y. B. Erol, R. Salas-Boni, X. Zhang, and X. Hu, "Robust qrs peak detection by multimodal information fusion of ecg and blood pressure signals," *Physiol. Meas.*, vol. 37, no. 11, p. N84, 2016.
- [5] S. A. Rankawat and R. Dubey, "Robust heart rate estimation from multimodal physiological signals using beat signal quality index based majority voting fusion method," *Biomedical Signal Processing and Control*, vol. 33, pp. 201–212, 2017.
- [6] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [7] P. S. Hamilton and W. J. Tompkins, "Quantitative investigation of qrs detection rules using the mit/bih arrhythmia database," *IEEE Trans. Biomed. Eng.*, no. 12, pp. 1157–1165, 1986.
- [8] B. U. Kohler, C. Hennig, and R. Orglmeister, "The principles of software qrs detection," *IEEE Eng. Med. Biol. Mag.*, vol. 21, no. 1, pp. 42–57, 2002.
- [9] I. I. Christov, "Real time electrocardiogram qrs detection using combined adaptive threshold," *Biomed. Eng. Online*, vol. 3, no. 1, p. 28, 2004.
- [10] F. Zhang and Y. Lian, "Qrs detection based on multiscale mathematical morphology for wearable ecg devices in body area networks," *IEEE Trans. Biomed. Eng.*, vol. 3, no. 4, pp. 220–228, 2009.
- [11] A. Ghaffari, M. Homaeinezhad, M. Akraminia, M. Atarod, and M. Daevaeiha, "A robust wavelet-based multi-lead electrocardiogram delineation algorithm," *Med. Eng. Phys.*, vol. 31, no. 10, pp. 1219–1227, 2009.
- [12] A. Ghaffari, M. Homaeinezhad, M. Khazraee, and M. Daevaeiha, "Segmentation of holter ecg waves via analysis of a discrete wavelet-derived multiple skewness–kurtosis based metric," *Annals of biomedical engineering*, vol. 38, no. 4, pp. 1497–1510, 2010.
- [13] M. S. Manikandan and K. Soman, "A novel method for detecting r-peaks in electrocardiogram (ecg) signal," *Biomedical Signal Processing and Control*, vol. 7, no. 2, pp. 118–128, 2012.
- [14] N. M. Arzeno, Z. D. Deng, and C. S. Poon, "Analysis of first-derivative based qrs detection algorithms," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 2, pp. 478–484, 2008.
- [15] R. Gutiérrez-Rivas, J. J. García, W. P. Marnane, and Á. Hernández, "Novel real-time low-complexity qrs complex detector based on adaptive thresholding," *IEEE Sensors J.*, vol. 15, no. 10, pp. 6036–6043, 2015.
- [16] J. Ding, C. Huang, Y. Ho, C. Hung, Y. Lin, and Y. Chen, "An efficient selection, scoring, and variability ratio test algorithm for ecg r-wave peak detection," *Exp Clin Cardiol*, vol. 20, no. 8, pp. 4256–4263, 2014.
- [17] V. X. Afonso, W. J. Tompkins, T. Q. Nguyen, and S. Luo, "Ecg beat detection using filter banks," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 2, pp. 192–202, 1999.
- [18] U. Pangerc and F. Jager, "Robust detection of heart beats in multimodal records using slope-and peak-sensitive band-pass filters," *Physiol. Meas.*, vol. 36, no. 8, p. 1645, 2015.
- [19] W. Zong, T. Heldt, G. Moody, and R. Mark, "An open-source algorithm to detect onset of arterial blood pressure pulses," in *IEEE Computers in Cardiology*, 2003, pp. 259–262.
- [20] M. A. Pimentel, M. D. Santos, D. B. Springer, and G. D. Clifford, "Heart beat detection in multimodal physiological data using a hidden semi-markov model and signal quality indices," *Physiol. Meas.*, vol. 36, no. 8, p. 1717, 2015.
- [21] S. Kiranyaz, T. Ince, and M. Gabbouj, "Real-time patient-specific ecg classification by 1-d convolutional neural networks," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 3, pp. 664–675, 2016.
- [22] ECAR, AAMI, "Recommended practice for testing and reporting performance results of ventricular arrhythmia detection algorithms," *Association for the Advancement of Medical Instrumentation*, 1987.
- [23] Y. LeCun and Y. Bengio, "Convolutional networks for images, speech, and time series," *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [24] P. De Chazal, M. O'Dwyer, and R. B. Reilly, "Automatic classification of heartbeats using ecg morphology and heartbeat interval features," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 7, pp. 1196–1206, 2004.
- [25] A. E. Johnson, J. Behar, F. Andreotti, G. D. Clifford, and J. Oster, "Multimodal heart beat detection using signal quality indices," *Physiol. Meas.*, vol. 36, no. 8, p. 1665, 2015.
- [26] S. Arlot, A. Celisse *et al.*, "A survey of cross-validation procedures for model selection," *Statistics surveys*, vol. 4, pp. 40–79, 2010.
- [27] L. Galeotti, C. G. Scully, J. Vicente, L. Johannesen, and D. G. Strauss, "Robust algorithm to locate heart beats from multiple physiological waveforms by individual signal detector voting," *Physiol. Meas.*, vol. 36, no. 8, p. 1705, 2015.
- [28] T. De Cooman, G. Goovaerts, C. Varon, D. Widjaja, T. Willemen, and S. Van Huffel, "Heart beat detection in multimodal data using automatic relevant signal detection," *Physiol. Meas.*, vol. 36, no. 8, p. 1691, 2015.
- [29] C. H. Antink, C. Brüser, and S. Leonhardt, "Detection of heart beats in multimodal data: a robust beat-to-beat interval estimation approach," *Physiol. Meas.*, vol. 36, no. 8, p. 1679, 2015.
- [30] M. Elgendi, A. Mohamed, and R. Ward, "Efficient ecg compression and qrs detection for e-health applications," *Sci. Rep.*, vol. 7, no. 459, 2017.
- [31] D. S. Benitez, P. Gaydecki, A. Zaidi, and A. Fitzpatrick, "A new qrs detection algorithm based on the hilbert transform," in *IEEE Computers in Cardiology*, 2000, pp. 379–382.
- [32] www.physionet.org/physiobank/database/html/mitdbdir/records.htm.