Extraction of ECG features with spiking neurons for decreased power consumption in embedded devices

Zonglong Li, Laurie E Calvet

Extraction of ECG features with spiking neurons for decreased power consumption in embedded devices

Zonglong LI
C2N (UMR 9001) CNRS / Université Paris-Saclay
Palaiseau, France
zonglong.li@universite-paris-saclay.fr

Laurie E. Calvet
C2N (UMR 9001) CNRS / Université Paris-Saclay
Palaiseau, France
laurie.calvet@c2n.upsaclay.fr

***Abstract*—In recent years, the computational efficiency of biomimetic spiking neural networks have received increasing interest. Here we show how a biomimetic spiking neuron circuit can be used to reduce the data transfer from the embedded device to a nearby computer for an electrocardiogram classification.**

Keywords— spiking neuron circuit, embedded systems, electrocardiogram, supervised learning

I. INTRODUCTION

Cardiovascular disease is the leading cause of death globally and early detection of heart dysfunction can save lives. Developing wearable devices that allow continuous monitoring of heart health would prevent catastrophic events and is therefore an important goal. A major challenge for all wearable health monitoring devices is to achieve low power consumption while allowing continuous operation.

One possible solution is to employ neuromorphic circuits, which exhibit high energy efficiency and are therefore well-suited for embedded electronics applications, including wearable devices. Recently, researchers have designed a spiking recurrent neural network (sRNN) consisting of leaky integrate-and-fire (LIF) neurons to detect anomalous patterns in electrocardiograms (ECGs) [1]. Energy consumption in such continuously monitoring devices, however, is dominated by the wireless transmission of the raw ECG data. Therefore, to extend the use time, their system only transmits ECG data for in-depth analysis when an unspecified anomaly is detected. We previously showed how transmission of EEG data could be significantly reduced by finding the most salient features with a spiking neuron circuit and then transmitting just these features [2]. Here we propose via simulations that ultra-low power spiking neurons can also be used to detect ECG features and explore whether this information is sufficient to perform a classification in a nearby device.

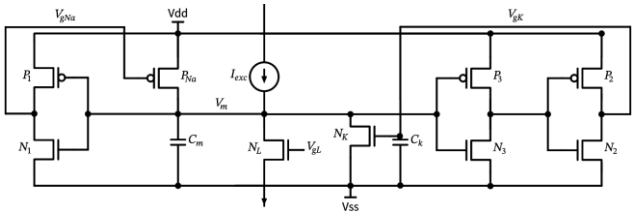

Fig. 1. Morris-Lecar inspired neuronal model of CMOS.

This work was supported by ANR contract ANR-21-FAI1-0006-01

II. METHODS

A. Spiking neuron circuit

We consider a Morris-Lecar inspired neuron circuit. The Morris-Lecar model is one of the simplest models of biological neurons. It consists of two ion channels Ca^{++} or Na^{+} and K^{+}, which allow the flow of ions into and out of the cell. The resulting change in the electrical potential of the cell membrane, V_m, leads to a sequence of action potentials, also known as a spike train. It results in two coupled non-linear differential equations that describe a dynamic system and variations of its parameters lead to different types of bifurcation. It allows, for instance, the observation of type I (where the frequency is a continuous of the applied current) and type II neurons (where the frequency versus applied current can be discontinuous). *Sourikopoulos et al* [3] showed how this model can be schematized into a very low energy consuming CMOS circuit. It is depicted *Fig. 1,* with the addition of the leak transistor N_L, and is the basis for the simulations considered here. Note that our simulations were benchmarked against the devices reported in [2, 3]. The neurons used are of type I.

We describe the basic operation of the circuit in *Fig. 1*. If V_m is at resting potential, the transistors P_{na} and N_k are off. As the input current I_{exc} charges the capacitor C_m, V_m increases gradually. Once V_m reaches a certain value, P_{na} turns on and allows a larger current to charge C_m. This causes V_m to increase even faster until it is large enough to turn on N_k. N_k then discharges C_m, causing V_m to drop quickly until it reverts to resting potential. If there is still a non-zero input current, the cycle will restart, creating an oscillating voltage or spike train. In a specific range of I_{exc}, the firing rate of the circuit increases. This property allows the firing rate to encode the features of the input data.

To evaluate the circuit's response to ECG signals, we simulate its differential equations using Python. We set the supply voltage V_{dd} to *100 mV* to ensure the transistors operate in the deep sub-threshold regime, which optimizes energy efficiency. It can be described by an exponential dependence:

$$I_{ds} = G_0 * \exp\left(\frac{V_{gs}}{\eta\, Vt}\right) * V_{ds} \qquad (1)$$

G_0 is the device conductance which is proportional to the ratio of the transistor length and width, η is the subthreshold slope

factor, V_t is the thermal voltage, V_{gs} and V_{ds} are the gate-to-source and drain-to-source voltages. Note that the threshold voltage of the transistors is taken to be *0V*. We define the different transistors by changing G_0 proportionately to the transistor size. Note that V_{gL} controls the leak transistor and is used to adjust the initial current at which the neuron begins to generate pulses, henceforth termed the neuron threshold. This adjustment allows optimizing the circuit so that it focuses on different parts of the input signal. However, for the sake of convenience in the simulation, we directly define a value for the leakage current instead of using a leakage transistor.

We choose to use python code rather than SPICE because simulating such a dynamic system is computationally intensive. In order to gain in computation time, we do not need the more complex models found in SPICE simulators but can use the relatively simple description of the sub-threshold behavior.

Table Ⅰ. Parameters for the neuron

Parameters	Values
C_m	0.05 (pF)
C_k	0.1 (pF)
G_{na}	60 (pS)
G_k	200 (pS)
G_{p1}	16 (pS)
G_{n1}	15 (pS)
G_{p2}	10 (pS)
G_{n2}	15 (pS)
G_{p3}	26 (pS)
G_{n3}	15 (pS)
V_{dd}	100 (mV)
V_{ss}	-100 (mV)
V_{et}	40 (mV)
IL	low threshold neuron: 0.1 (pA) high threshold neurons: 35-130(pA)

Fig. 2 shows simulations of the neuron firing rate as it varies with the input current. The neuron circuit generates spikes within the current input range of ~10 - 600 pA. The curve shows how beyond a certain threshold the neuron is no longer bistable. Note that the relationship between firing rate and input current is not linear. The firing rate changes faster when the input current is small, and slower when the input current is large. These are crucial properties that allow us to identify electrocardiogram (ECG) features.

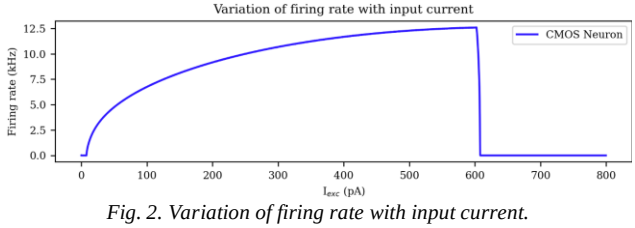

Fig. 2. Variation of firing rate with input current.

B. ECG classification task

The ECG allows a specialist to understand the functioning of the heart by depicting the propagation of the cardiac dipole resulting from the activity of polarization-depolarization of cardiac cells. In *Fig. 3* [4] shows how a typical ECG signal

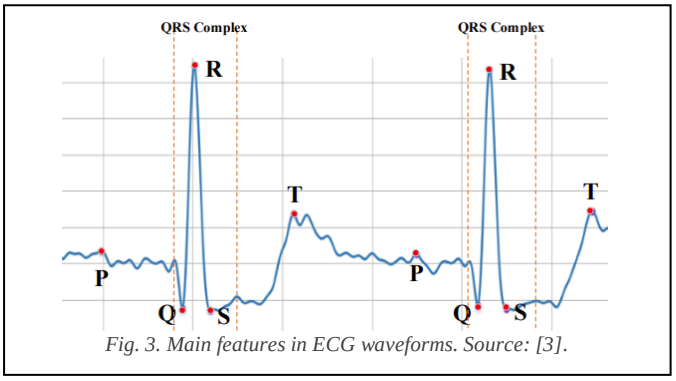

Fig. 3. Main features in ECG waveforms. Source: [3].

consists of *5* 'waves': P, Q, R, S and T. In this polarization of the electrodes, the R wave is clearly the highest amplitude and is the most straightforward and crucial part of detection. The distance between R waves gives the heart rate.

Building on previous work that explored how EEG data could be classified through the utilization of firing rates produced by spiking neurons [2], here we explore how to obtain ECG features from the spikes generated by spiking neurons and classify heartbeats based on the obtained features.

We use a standard ECG dataset, MIT-BIH Arrhythmia Database [5]. This data set has 48 half-hour excerpts of two-channel, 24-hour, ECG recordings obtained from 47 subjects. It includes normal heartbeats and a large variety of different anomalies, some of which are quite rare. The sampling rate of the signal is 360 Hz, and each heartbeat is labeled by cardiologists. We explore how to classify normal heartbeats and the four most common types of anomalous heartbeats, as depicted in *Table II*. From the database, we chose the patients that exhibit many of these anomalies, as listed in *Table III*.

Table Ⅱ. Our databases with 5 types of heartbeats

Heartbeat type	Training set (Beats)		Test set (Beats)	
	A	B	A	B
1. Normal rhythm (N)	10969 / 9626		5484 / 4812	
2. Atrial premature beat (A)	10969 / 9626		652 / 567	
3. Premature ventricular contraction (P)	10969 / 9626		568 / 397	
4. Left bundle branch block beat (L)	10969 / 9626		2184 / 2128	
5. Right bundle branch block beat (R)	10969 / 9626		1875 / 1806	
Total	54845 / 9626		10763 / 9710	

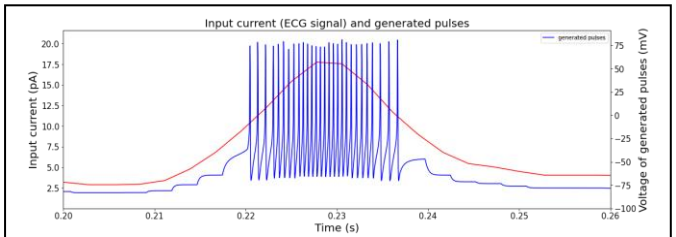

Fig. 4. Neuron with high threshold for detecting R-waves. The red curve (left axis) is the processed ECG data that is input into the circuit. The blue is the circuit output (right axis).

Before feeding the ECG signal to the neuron, we pre-process the signal. Our process of obtaining ECG features relies on the amplitude of ECG signal, making it important for us to correct the baseline. In our simulation, we used an algorithm [6] to correct the baseline. Meanwhile, using a suitable bandpass filter like low cut-off frequency 5Hz and high cut-off frequency 15Hz, in embedded devices can also reduce the influence of noise and baseline wander [7]. We then adjust the ECG signal to a range of input currents that can be processed by the neuron, in a real circuit this would be done simply by an additional transistor.

After this data processing, we input the ECG into two neuronal circuits with different neuron thresholds, adjusted using the leakage current *IL*, from *35pA* to *130pA* for different patients. Due to individual differences, there are differences in the amplitude of the R-wave, and in a practical device, we need an adaptive system to eliminate the effects of this variability. Homeostasis circuit is a good direction to solve the problem [8].

We use a high threshold neuron so that pulses are produced only in the region of each R-wave, as shown in *Fig. 4.* As the amplitude of the red ECG signal increases, it becomes apparent that the spacing between the blue pulses decreases gradually. For each R-wave generated pulse group, we find the position with the closest spacing between spikes, which represents the highest signal amplitude of the R-wave.

We lower the threshold by decreasing the value of leakage current *IL* to *0.1pA* to detect other small-amplitude features, The low threshold neuron produces pulses in all parts of the ECG and can detect the small amplitude waves. To determine the Q-wave, for example, from our observations of the database, for a normal ECG the distance between the Q- and R-waves of the same heartbeat is typically within *0.01*s ~ *0.03*s. We extend this range to *0.01*s ~ *0.07*s to detect the possible locations of the Q-wave when the heartbeat is abnormal like PVCs, as shown by the orange lines in *Fig. 5*. We then find the change in the distance between spikes, as described for the R-wave, and identify the location with the lowest amplitude as that with the largest inter-pulse interval.

In our initial study, we obtained 15 features for each heartbeat: the time of P, Q, S and T waves relative to the R-wave of the same heartbeat, the amplitude of the P, Q, R, S and T wave, expressed here as the frequency of the spikes, the duration of R-wave, Pre-RR and Post-RR: the time difference between each R-wave and the previous and next R-wave, Variable-RR: average of current Pre-RR and Post-RR, Local-RR: average of the previous 10 Pre-RRs, Average RR: average of the previous 60 Pre-RRs. *Fig. 6* shows an example of the positions of the P, Q, R, S and T waves obtained by this method.

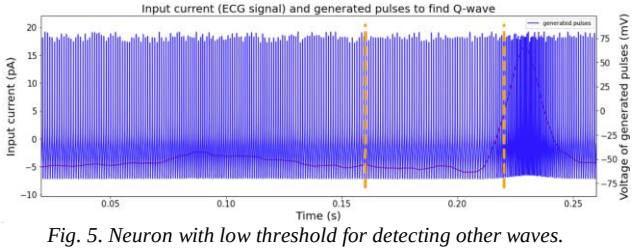

Fig. 5. Neuron with low threshold for detecting other waves.

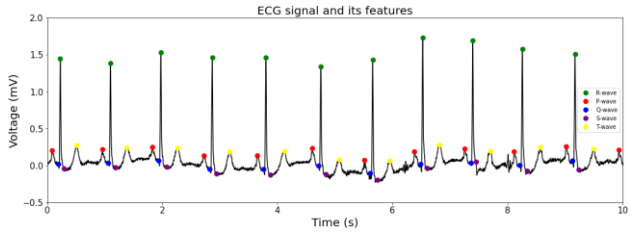

Fig. 6. Example of the positions of P, Q, R, S, T waves obtained using the method of the spiking neurons.

However, to compare the traditional method of extracting ECG features offline in the next step, we do not use information about P-wave, T-wave, and duration of R-wave.

To determine the detection accuracy, we set an error tolerance for determining whether the found R-wave is that of an R-wave. If it occurs within 0,05s of the time of the real R-wave, then it is considered to be correctly detected. The results for the R-waves (present in the original database classification) are summarized in *Table III*, under A in the table).

To perform the machine learning classification, we prepared a training set by padding the observed anomalies so that they all had an equal number of occurrences and left the initial ratio in the test set as shown in *Table II* (label A). To ensure standardization, only the first 1800 seconds (30 minutes) of each file were extracted, leaving approximately 5 seconds unused.

Using the features obtained from our method, we fit the data using the sci-kit learn package [9]. In order to use same data to compare other people's technologies, we chose an open-source Python implementation of QRS complex detector for ECG signals [10], based on the Pan-Tomkins algorithm [7]. Importantly, the programs provided online did not implement all the features in the article, like fiducial mark on filtered data, use of another set of thresholds based on the filtered ECG, irregular heart rate detection, and missing QRS complex detection search-back mechanism. In addition, we added two functions to detect Q-wave and S-wave, the principle is to find the minimum value within a reasonable range for each R-wave.

III. RESULTS

Table III displays the number of unrecognizable heartbeats for each file: our approach(A) / Pan-Tomkins's algorithm(B). The results show that our method has a very high accuracy in detecting the QRS complex with > 99% overall accuracy, confirming the interest for this method. While here we observed our accuracy to be greater than the Pan-Tomkins algorithm, note that, as discussed in the methods section, the QRS detector was missing some of the features in [10]. Most notably, the lack the dual-threshold technique (to find missed beats, and thereby reduce false negatives) and irregular heart rate detection mechanism in the original paper, there may be a considerable number of undetectable heartbeats when premature atrial contractions (PACs) and premature ventricular contractions (PVCs) are present in the ECG like file 209, 214, 221 and 228, which will impact the comparison of the classifications.

Table III. Our method to detect QRS complex (A) compared with [8] (B) FP means false positive: a heardbeat was detected but was not identified by the cardiologist and FN means false negative means a heartbeat was not detected but was identified by a cardiologist.

File	Total (1800s) (Beats)	FP (Beats)		FN (Beats)		Failed Detection (Beats)		Failed Detection (%)	
		A	B	A	B	A	B	A	B
109	2522	16 / 6		19 / 21		35 / 27		1.4 / 1.1	
111	2118	4 / 7		36 / 7		40 / 14		1.9 / 0.7	
116	2404	7 / 3		24 / 30		31 / 33		1.3 / 1.4	
118	2271	7 / 2		16 / 8		23 / 10		1.0 / 0.4	
119	1981	2 / 1		4 / 0		6 / 1		0.3 / 0.1	
209	2997	3 / 1		3 / 181		6 / 182		0.2 / 6.1	
212	2740	0 / 1		0 / 0		0 / 1		0 / 0	
214	2251	24 / 10		26 / 135		50 / 145		2.2 / 6.4	
215	3352	2 / 1		6 / 133		8 / 134		0.2 / 4.0	
220	2041	0 / 0		0 / 21		0 / 21		0 / 1.0	
221	2420	10 / 1		18 / 185		28 / 186		1.2 / 7.7	
228	2047	57 / 1		48 / 1737		105 / 1738		5.1 / 84.9	
231	1565	0 / 2		0 / 2		0 / 4		0 / 0.3	
232	1775	2 / 4		2 / 4		4 / 8		0.2 / 0.5	
Total	32484	134 / 40		202 / 2464		336 / 2504		1.0 / 7.7	
Total excluding 228	30437	77 / 39		154 / 727		231 / 766		0.7 / 2.5	

We report the classified the heartbeat with ten normalized features based on the QRS detection: Q-time, S-time, Pre-RR, Post-RR, Q-amplitude, R-amplitude, S-amplitude, Variable-RR, Local-RR, Average-RR. The confusion matrix is shown in *Fig. 7*. We find that the k-nearest neighbors (KNN) algorithm has a 96% overall accuracy using the ECG features obtained through the neurons and 95% overall accuracy using the ECG features obtained through [8].

IV. DISCUSSION AND CONCLUSION

In this first result, our neuron parameters are not very energy efficient, as the large number of spikes generated by the low threshold neuron compared to the high threshold neuron would result in it consuming a larger energy. This could easily be optimized, however, by changing the neuron parameters in *Table I* so that fewer spikes are emitted. Although reducing the firing rate results in fewer spikes, it should be noted that excessively low firing rate may results in larger errors. A compromise must be found between the firing rate and the desired level of accuracy.

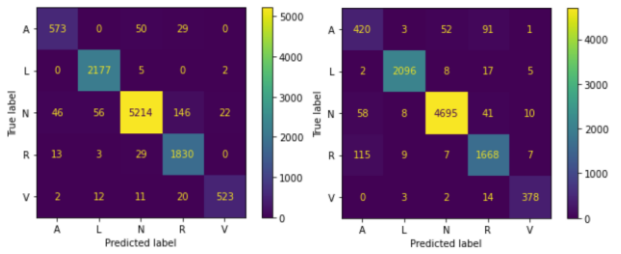

Fig. 7. Left: confusion matrix using features generated from the spiking neurons. Right: confusion matrix based on features generated using the QRS detection based on the method in [8] (For more information about the categories A,L,N,R,V see Table II)

Despite the relatively high number of spikes generated by the low threshold neuron, the energy consumption per spike will be very low (~ 4-fJ/Spike in an ideal case) [3] and would be on the order of 10s of pJ resulting in ~10 nW per heartbeat cycle. While the circuits that would enable the feature detection to be obtained from the spiking neurons would need to be detailed to obtain the total energetic cost of this implementation, our purpose here, however, is to demonstrate the validity of using rate encoding to extract ECG features and show that they can provide an accuracy comparable to an automated detection of features.

We note that the primary source of energy consumption is still likely to originate from the transfer of data to a nearby computer or portable device and is ~ 10 μW for the best detector. With just 6 features transmitted (Time and amplitude of Q, R and S waves) for each heartbeat instead of ~300, the total power consumption will be significantly reduced. In the future tuning of the neuron, different neuron parameters would enable different firing rates so that the firing rate of the low threshold neuron is reduced.

We have explored the use of spiking neurons to probe features of ECG data. We find that we can realize and excellent classification of the different heart anomalies using the features from the spiking neurons and performing a nearest neighbor classification. This suggests that spiking neurons provide a low energy route towards realizing ECG detection for continuous heart monitoring.

REFERENCES

[1] Bauer, Felix Christian et al. “Real-time ultra-low power ECG anomaly detection using an event-driven neuromorphic processor.” IEEE transactions on biomedical circuits and systems 13.6 (2019): 1575-1582.

[2] Calvet, Laurie E, et al. “Spiking sensory neurons for analyzing electrophysiological data.” *ECS Journal of Solid State Science and Technology* 9.11 (2020): 115004.

[3] Sourikopoulos, Ilias, et al. “A 4-fJ/spike artificial neuron in 65 nm CMOS technology.” Frontiers in neuroscience 11 (2017): 123.

[4] Wijaya, Chandra, et al. “Abnormalities state detection from P-wave, QRS complex, and T-wave in noisy ECG.” Journal of Physics: Conference Series. Vol. 1230. No. 1. IOP Publishing, 2019.

[5] Moody, George B., and Roger G. Mark. “The impact of the MIT-BIH arrhythmia database.” IEEE engineering in medicine and biology magazine 20.3 (2001): 45-50.

[6] Md Azimul Haque (2022). Feature Engineering & Selection for Explainable Models A Second Course for Data Scientists

[7] Pan, Jiapu, et al. "A real-time QRS detection algorithm." IEEE transactions on biomedical engineering 3 (1985): 230-236.

[8] Bartolozzi, Chiara, et al. "Implementing homeostatic plasticity in VLSI networks of spiking neurons." 2008 15th IEEE International Conference on Electronics, Circuits and Systems. IEEE, 2008.

[9] Pedregosa, Fabian, et al. "Scikit-learn: Machine learning in Python." the Journal of machine Learning research 12 (2011): 2825-2830.

[10] Michał Sznajder, & Marta Łukowska. (2017). Python Online and Offline ECG QRS Detector based on the Pan-Tomkins algorithm (v1.1.0).