

 Open access • Journal Article • DOI:10.1109/TBCAS.2019.2953001

Real-time ultra-low power ECG anomaly detection using an event-driven neuromorphic processor — [Source link](#)

Felix Christian Bauer, Dylan R. Muir, Giacomo Indiveri

Institutions: University of Zurich

Published on: 13 Nov 2019 - arXiv: Signal Processing

Topics: Neuromorphic engineering and Reservoir computing

Related papers:

- [Real-time wearable telecardiology from representative signals](#)
- [A Cascaded Convolutional Neural Network for Assessing Signal Quality of Dynamic ECG.](#)
- [Contactless Real-Time Heartbeat Detection via 24 GHz Continuous-Wave Doppler Radar Using Artificial Neural Networks.](#)
- [Techniques for Event and Feature Detection](#)
- [Multistage Adaptive filter for ECG signal processing](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/real-time-ultra-low-power-ecg-anomaly-detection-using-an-2pjdvgvw8>



Real-Time Ultra-Low Power ECG Anomaly Detection Using an Event-Driven Neuromorphic Processor

Bauer, Felix Christian ; Muir, Dylan Richard ; Indiveri, Giacomo

Abstract: Accurate detection of pathological conditions in human subjects can be achieved through off-line analysis of recorded biological signals such as electrocardiograms (ECGs). However, human diagnosis is time-consuming and expensive, as it requires the time of medical professionals. This is especially inefficient when indicative patterns in the biological signals are infrequent. Moreover, patients with suspected pathologies are often monitored for extended periods, requiring the storage and examination of large amounts of non-pathological data, and entailing a difficult visual search task for diagnosing professionals. In this work we propose a compact and sub-mW low power neural processing system that can be used to perform on-line and real-time preliminary diagnosis of pathological conditions, to raise warnings for the existence of possible pathological conditions, or to trigger an off-line data recording system for further analysis by a medical professional. We apply the system to real-time classification of ECG data for distinguishing between healthy heartbeats and pathological rhythms. Multi-channel analog ECG traces are encoded as asynchronous streams of binary events and processed using a spiking recurrent neural network operated in a reservoir computing paradigm. An event-driven neuron output layer is then trained to recognize one of several pathologies. Finally, the filtered activity of this output layer is used to generate a binary trigger signal indicating the presence or absence of a pathological pattern. We validate the approach proposed using a Dynamic Neuromorphic Asynchronous Processor (DYNAP) chip, implemented using a standard 180 nm CMOS VLSI process, and present experimental results measured from the chip

DOI: <https://doi.org/10.1109/tbcas.2019.2953001>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-184174>

Journal Article

Accepted Version

Originally published at:

Bauer, Felix Christian; Muir, Dylan Richard; Indiveri, Giacomo (2019). Real-Time Ultra-Low Power ECG Anomaly Detection Using an Event-Driven Neuromorphic Processor. *IEEE Transactions on Biomedical Circuits and Systems*, 13(6):1575-1582.

DOI: <https://doi.org/10.1109/tbcas.2019.2953001>

Real-time ultra-low power ECG anomaly detection using an event-driven neuromorphic processor

Felix Christian Bauer

aiCTX AG

Zurich, Switzerland

felix.bauer@aictx.ai

Dylan Richard Muir

aiCTX AG

Zurich, Switzerland

dylan.muir@aictx.ai

Giacomo Indiveri

University of Zurich & ETH Zurich

Zurich, Switzerland

giacomo@ini.uzh.ch

November 14, 2019

Abstract

Accurate detection of pathological conditions in human subjects can be achieved through off-line analysis of recorded biological signals such as electrocardiograms (ECGs). However, human diagnosis is time-consuming and expensive, as it requires the time of medical professionals. This is especially inefficient when indicative patterns in the biological signals are infrequent. Moreover, patients with suspected pathologies are often monitored for extended periods, requiring the storage and examination of large amounts of non-pathological data, and entailing a difficult visual search task for diagnosing professionals.

In this work we propose a compact and sub-mW low power neural processing system that can be used to perform on-line and real-time preliminary diagnosis of pathological conditions, to raise warnings for the existence of possible pathological conditions, or to trigger an off-line data recording system for further analysis by a medical professional. We apply the system to real-time classification of ECG data for distinguishing between healthy heartbeats and pathological rhythms.

Multi-channel analog ECG traces are encoded as asynchronous streams of binary events and processed using a spiking recurrent neural network operated in a reservoir computing paradigm. An event-driven neuron output layer is then trained to recognize one of several pathologies. Finally, the filtered activity of this output layer is used to generate a binary trigger signal indicating the presence or absence of a pathological pattern. We validate the approach proposed using a Dynamic Neuromorphic Asynchronous Processor (DYNAP) chip, implemented using a standard 180 nm CMOS VLSI process, and present experimental results measured from the chip.

Keywords— Neuromorphic, spiking neural network, recurrent network, reservoir computing, ultra-low power, electrocardiogram

1 Introduction

Cardiovascular diseases are among the main causes of premature deaths worldwide [1]. Some pathological conditions can be diagnosed at an early stage by electrocardiography (ECG), which measures the electrical excitation of the heart. In such cases quick interventions may lead to better outcomes and even save lives [2–4].

However, to detect such early indicators in undiagnosed patients it would be necessary to carry out continuous health monitoring over extended periods of time. Until now this has been impractical for several reasons. First of all, long-term monitoring generates large amounts of data, much of which is likely to be non-pathological. Searching such data for indicative segments is a tedious task, which must be performed by trained medical professionals. Furthermore, to reduce inconvenience for patients, recording devices should be portable — ideally wearable — and have a long battery life to avoid frequent recharging or replacement. This is in conflict with the requirement of the device to perform high accuracy recordings and store or transmit the large amounts of recorded data for off-line processing, as these operations typically consume significant energy.

This conflict can be resolved by devices that locally run algorithms to distinguish and filter out irrelevant data continuously and in real-time, without resorting to external- or cloud-computing resources. In this way the amount of data to be stored or transmitted can be drastically reduced and subsequent in-depth analysis by a diagnosing professional is facilitated.

Such a system could, for instance, be part of a wearable ECG monitoring device that is provided to patients with a suspected cardiovascular condition (see Fig. 1). The patient’s electrocardiogram is continuously measured through sensor electrodes and analyzed by our proposed computing system. In this system, the majority of the system components can be in hibernation for low-power operation (shaded components). If a suspicious ECG rhythm is detected, an interrupt is sent to a microcontroller or CPU, which will cause the ECG signal to be recorded onto non-volatile memory or to be transmitted to a server via a wireless connection. In addition, an alarm may be raised if indicative patterns are detected repeatedly, making timely medical treatment more likely. Such a system is designed to perform only preliminary analysis, but can drastically reduce the diagnostic load on the medical professional.

Automated ECG arrhythmia detection has received much interest in the literature. Typical approaches involve neural networks [5–7], support vector machines [8], fuzzy cognitive maps [9] or the extraction and analysis of *a priori* defined features [10]. While many of these algorithms achieve detection accuracies of beyond 90 per cent, they are generally not applicable for ultra-low power hardware implementation. Algorithms based on spiking neural networks (SNNs) are suggested to be well-suited for sub milliwatt ECG heart rate [11] and arrhythmia [12] detection; however these networks are only implemented in simulation on conventional computing hardware.

An alternative cloud computing approach involves sending data to an external, more powerful processor [13, 14]. In spite of advances in data compression [15], continuous

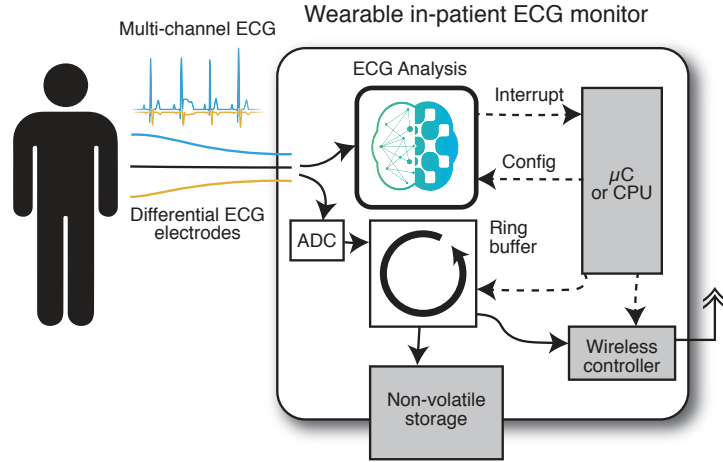


Figure 1: Possible integration of the proposed system into a wearable ECG monitor. A patient’s ECG is measured with differential ECG electrodes and continuously analysed. If an indicative pattern is detected, the system wakes up a microcontroller or CPU to record the relevant data onto a non-volatile storage or transmit it to a server via a wireless connection. Solid lines: ECG signal pathways; dashed lines: command signals; shaded blocks: components in hibernation mode until a pathological signal is detected.

data transmission itself adds considerable power demand to the system. Besides, users may be limited in their mobility if they need to remain within range of a receiving device or of a wireless network.

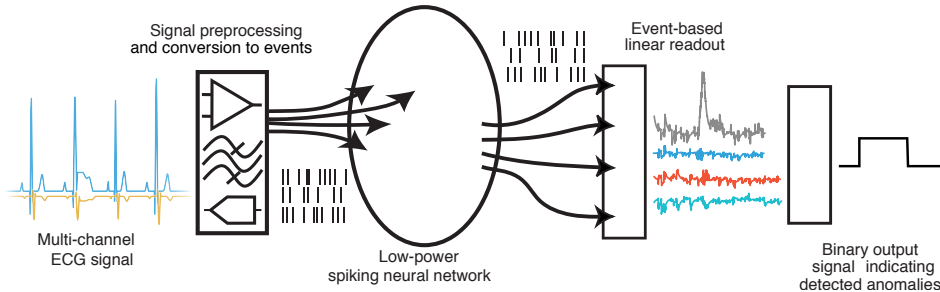


Figure 2: System for real-time ECG anomaly detection. Analog sensor inputs are amplified and filtered in a preprocessing stage and then converted to events using Lebesgue sampling. Resulting event trains are processed by a spiking neural network implemented on a neuromorphic processor. For each type of anomalous pattern a readout unit low-pass filters a weighted sum of neuron firing activities to estimate the likelihood of an anomaly being present in the ECG input. A final binary output signal indicates whether any readout unit is above its detection threshold.

Here we propose a scalable, always-on system architecture for continuous signal processing at microwatt power levels (see Fig. 2), which exploits the ultra-low power characteristics of Very Large Scale Integration (VLSI) neuromorphic processors [16]. Input is evaluated continuously and in real-time. No segmentation of the signal or synchronization to heartbeats is required.

In the approach proposed, analog signals are sensed, amplified, filtered, and then con-

verted to trains of events using Lebesgue sampling via a sigma-delta encoder [17]. The resulting event trains are expanded by a spiking recurrent neural network (sRNN) of leaky integrate-and-fire (LIF) neurons, implemented on the neuromorphic hardware [18]. Neural firing activities are combined and filtered by an event-based linear readout layer that has been trained to detect the presence of specific anomalous patterns in the ECG input. Finally a binary output is generated that indicates the presence or absence of a pathological event.

We validate the architecture proposed by implementing the sRNN on the fully asynchronous, mixed-signal DYNAP-SE device [19], and show experimental results performing real-time ultra-low power analysis of a two-channel ECG signal. In particular, we show that anomalies are detected with high reliability and demonstrate the feasibility of the real-world application of a system as described above. Similar approaches have recently been applied to classify EMG signals with a feed-forward SNN [20] and for classification of ECG anomalies with a sRNN [21], using the same type of hardware.

For other processing stages of the system we refer to existing implementations. Their specifications, together with experimental findings from the DYNAP-SE implementation, allow us to estimate power consumption of a complete system.

To our best knowledge this work presents the first full-system description of an ECG anomaly detector based on a neuromorphic processor and one of the first implementations of a real-world application on the DYNAP-SE system.

2 Methods

2.1 The DYNAP-SE system

The sRNN was implemented on a fully asynchronous, mixed-signal DYNAP-SE chip, first presented and fully characterized in [19] (see Fig. 3a). The DYNAP-SE chip comprises four cores of 256 Adaptive Exponential Integrate-and-Fire (AdExp) neurons, excitatory and inhibitory dynamic synapses implemented using a “Differential-Pair Integrator” (DPI) log-domain circuit [22], and a hierarchical routing architecture for transmitting neural spiking events within cores, among multiple cores and across multiple chips.

Every DYNAP-SE AdExp silicon neuron can subscribe to events from up to 64 presynaptic neurons. For each inter-neuron connection one of four DPI synapse types can be selected, two of which are excitatory and two inhibitory; each with individually tunable characteristics. Neural and synaptic dynamics can be adjusted by setting 25 different circuit parameters, programmed via a temperature-compensated on-chip bias generator [23], shared by neurons on the same core. This also holds for presynaptic weights for same synapse types. To achieve stronger synaptic weights, multiple connections can be implemented between a pair of neurons, thereby effectively multiplying the weight by the number of connections.

Due to the analog nature of the silicon neurons and their device mismatch, temporal and functional characteristics, such as firing thresholds, synaptic efficacy, and time constants, vary between individual neurons and synapses. This intrinsic in-homogeneity is exploited

to introduce variability in neuron parameters and enhance signal dimensionality, as explained in Section 2.3.

For this work, a DYNAP-SE development kit was used, which comprises four DYNAP-SE chips and a Field Programmable Gate Array (FPGA), providing a USB interface to a standard PC for configuring circuit parameters, setting up synaptic connections, sending input events and reading out neural firing activity (see in Fig. 3a).




Figure 3: a) Die photo of the DYNAP-SE multi-core neuromorphic processor [19]. The chip is fabricated in standard 180 nm CMOS technology and comprises four cores with 256 adaptive exponential integrate-and-fire neurons each. b) DYNAP-SE development kit with USB interface to a standard PC.

2.2 ECG data set and conversion to spikes

The set of ECG signals used to evaluate performance in this work is taken from the MIT-BIH Arrhythmia Database [24], provided by Massachusetts Institute of Technology and Beth Israel Hospital through PhysioNet [25]. The data set consists of 48 half-hour excerpts from a set of 4000 ambulatory two-channel ECG recordings from 47 subjects. Every ECG rhythm is labeled either as normal or as exhibiting one of 18 anomalous conditions. 23 of the recordings were picked at random while the remaining 25 were selected to include less common anomaly types and parts of low signal quality to challenge arrhythmia detectors. The signal is band-pass filtered between 0.1 Hz and 100 Hz and digitized with a sampling rate of 360 Hz.

Some anomalies occur very rarely in the data set. In this work only normal rhythms and the five most prevalent anomaly types were considered, which accounts for 95.6 % of total recording time. Examples for normal heartbeats and the anomalies used in this work are shown in Fig. 4.

Three disjoint subsets of the input data were used for training, validation and testing of the system. For the training set, recordings were segmented into individual heartbeats. From these beats 30000 were then drawn at random and concatenated to form the training input. For validation and test sets, recordings were split into longer segments, each comprising five to ten contiguous heartbeats of the same label. In a similar fashion as with the training set, segments were then selected at random and concatenated to form the input signals. For all three sets, samples were drawn such that normal rhythms made up about 75 % of each set and any of the included anomaly types 5 %. Table 1 holds the number of beats in the three subsets for each label, Table 2 the number of continuous anomalous segments in the validation and test sets.

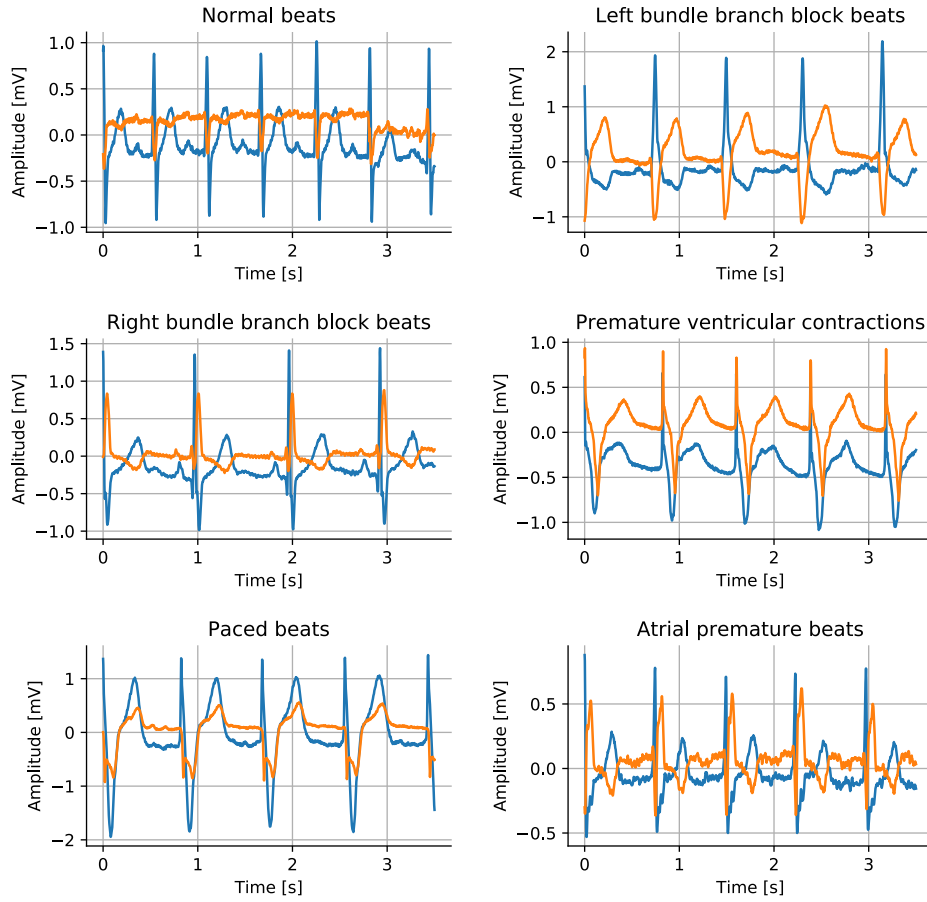


Figure 4: Examples of the ECG signal for normal heartbeats (upper left plot) and each of the five anomaly types used in this work. Blue curves correspond to the first, orange curves to the second channel of the respective recordings.

Beat type (expert label)	Subset		
	Train.	Valid.	Test.
<i>Normal rhythm</i>	22,500	2,315	1,569
<i>Left bundle branch block beat</i>	1,500	175	104
<i>Right bundle branch block beat</i>	1,500	176	103
<i>Premature ventricular contraction</i>	1,500	154	100
<i>Paced beat</i>	1,500	153	98
<i>Atrial premature beat</i>	1,500	168	104
Total	30,00	3,116	2,078

Table 1: Number of heartbeats per label

Anomaly type (expert label)	Subset	
	Validation	Testing
<i>Left bundle branch block beat</i>	24	15
<i>Right bundle branch block beat</i>	22	17
<i>Premature ventricular contraction</i>	22	14
<i>Paced beat</i>	21	16
<i>Atrial premature beat</i>	22	14
Total	111	76

Table 2: *Number of segments per anomaly type*

ECG signals are converted to trains of events through a sigma-delta encoding scheme [17]. For every ECG channel there are two pulse outputs, emitting events either when the input signal increases by a specified amount (up-events) or decreases by a given amount (down-events).

2.3 sRNN architecture

The architecture of the sRNN is illustrated in Fig. 5 and has been inspired by the paradigm of reservoir computing [26, 27]. In these architectures the weights of the recurrent network are initialized randomly and learning takes place only in the output layer, which reads out the state of a hidden layer sampled from the recurrently connected neurons. The network consists of three neuron populations that are implemented on the DYNAP processor: a feed-forward input expansion layer as well as a recurrent excitatory and a non-recurrent inhibitory group that together form the reservoir layer. Random connections and hardware mismatch (see Section 2.1) serve to project the original signal to a high-dimensional state space. Recurrent connections allow for the state to also encode information about the recent past of the input.

The input expansion layer consists of 128 neurons. Each has a fixed number of 1 to 64 presynaptic excitatory connections, drawn uniformly at random, to one of the input channels. Together with the mismatch between the hardware neurons, this connection scheme ensures that each neuron responds differently to the input spike trains, therefore increasing the dimensionality of the signal.

The reservoir layer is partitioned into 512 excitatory and 128 inhibitory neurons. Each excitatory neuron receives presynaptic input through 16 excitatory connections from the input expansion layer, 32 recurrent connections from other excitatory neurons of the same layer and 16 connections from the inhibitory population. Each inhibitory neuron holds presynaptic connections to 64 of the excitatory neurons.

This results in the following connectivities between the three different neuron populations:

<i>Input expansion to excitatory :</i>	12.5%
<i>Excitatory to excitatory :</i>	6.25%
<i>Excitatory to inhibitory :</i>	12.5%
<i>Inhibitory to excitatory :</i>	3.1%




Figure 5: *Recurrent sRNN architecture. Dimensionality of the spiking input is expanded by a feed-forward input expansion layer with post-synaptic connections to the excitatory population of a partitioned reservoir layer. The latter has recurrent connections to itself and additionally stimulates a feed-forward inhibitory population that in turn controls activation of the excitatory population.*

All connections are drawn at random. Multiple connections between the same pair of neurons may exist.

The three above-mentioned neuron populations are implemented on different cores of the neuromorphic processor, which allows setting hardware parameters individually for each and facilitates tuning full-network dynamics. Ideally the network is close to the edge of chaos [28], where dynamics are rich in response to an input signal but return to a stable state after external stimulation ends.

2.4 Learning and readout

Spiking activities of all hardware neurons are constantly monitored with a PC. Spike trains are low-pass filtered by convolution with an exponential kernel $\kappa(t) = \exp -\frac{t}{\tau_{out}}$ with time constant $\tau_{out} = 0.175$ s:

$$\vec{x}(t) = (\vec{s} * \kappa)(t)$$

Here, $\vec{s}(t)$ is a vector containing the spike trains of all hardware neurons and $\vec{x}(t)$ a vector of the filtered spike trains.

For each anomaly type in the input signal one readout unit is trained to detect it by taking a weighted sum of the filtered spiking activities. If this sum exceeds a fixed threshold, the anomaly counts as detected. Training of readout weights and thresholds is done independently for each anomaly type.

Readout weights are trained using linear least squares approximation, such that for all readout units i

$$\|X \cdot \vec{w}^{(i)} - \vec{y}^{(i)}\|$$

is minimized. Here, $\vec{w}^{(i)}$ is a vector holding input weights for unit i . The rows of matrix X are the filtered spike trains \vec{x} , sampled over time. The vector $\vec{y}^{(i)}$ is the target at corresponding time points. It is 1 whenever an anomaly of type i is present in the input signal and 0 otherwise.

During a validation run, a subset of the ECG data is used to find optimal detection thresholds for the readout units. Thresholds are chosen such that the number of false negatives for the respective anomaly and the number of false positives are minimised over the validation set.

3 Results

3.1 Network performance

As described in Section 2.2 the input signal to the network is assembled from segments of contiguous heartbeats. Within a segment, beats share the same label, which is either *normal* or one of the five considered anomalous patterns. The task for the sRNN is to indicate the presence of anomalous conditions in the ECG. Classification of the specific anomaly type is not required. Network output can be interpreted as binary, being positive whenever any of the readout units indicates a detected anomaly and negative otherwise.

Network performance is quantified by means of *sensitivity*, *specificity*, *PPV* (positive predictive value), and *NPV* (negative predictive value), which are defined as follows:

$$\begin{aligned} \textit{sensitivity} &= \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false negatives}} \\ \textit{specificity} &= \frac{\# \text{ true negatives}}{\# \text{ true negatives} + \# \text{ false positives}} \\ \textit{PPV} &= \frac{\# \text{ true positives}}{\# \text{ true positives} + \# \text{ false positives}} \\ \textit{NPV} &= \frac{\# \text{ true negatives}}{\# \text{ true negatives} + \# \text{ false negatives}} \end{aligned}$$

A true positive (false negative) is any segment of anomalous heartbeats during which the presence of an anomaly is (is not) indicated by any of the readout units. A true negative (false positive) is any normal heartbeat during which the network does not (does) indicate the presence of an anomaly. Because only the combined output of all readout units is considered, a unit indicating the presence of an anomaly which does not correspond to its own target label still counts as correct detection.

Overall, 91.3 % of anomalous segments in the test set were detected. Of the normal heartbeats, 2.4 % were falsely indicated as being anomalous. This means that for a normal ECG input, on average there would be a false positive every 31.5 seconds.

As shown in Table 3, for three of the five presented anomaly types the network detected all of the anomalous segments reliably, for another type sensitivity was 88 %. Only for

Anomaly type (expert label)	Sensitivity (in %)	Specificity (in %)	PPV (in %)	NPV (in %)	MTBFP ^a (in seconds)
<i>Left bundle branch block beat</i>	100.0	99.87	88.24	100.0	582.8
<i>Right bundle branch block beat</i>	88.24	99.36	60.0	99.87	116.56
<i>Premature ventricular contraction</i>	71.43	99.24	45.45	99.74	97.1
<i>Paced beat</i>	100.0	99.24	57.14	100.0	97.1
<i>Atrial premature beat</i>	100.0	99.94	93.33	100.0	1165.6
Overall	92.11	97.6	65.42	99.61	31.5

^aMean time between false positives during normal ECG input

Table 3: Anomaly detection performance on test set

Component	Power
<i>DYNAP-SE total</i>	516.1 μW
<i>Amplifiers</i>	$< 2 \times 50 \mu\text{W}^{\text{a}}$
<i>Filters</i>	$< 4 \times 4 \mu\text{W}^{\text{b}}$
<i>Readout</i>	$< 90 \mu\text{W}$
Total	$< 722.1 \mu\text{W}$

^aOne amplifier per ECG channel

^bOne low-pass and one notch filter per ECG channel

Table 4: Power consumption

one anomaly, *premature ventricular contraction*, detection was significantly less reliable with 71.43 %. There is no apparent correlation between individual subjects or ECG recordings and sensitivity for this anomaly.

Specificity is above 99 % for all anomaly types but particularly high for *left bundle branch block beat* and *atrial premature beat*, which resulted in only two and one false positives over the whole test set, respectively.

Values for PPV are lower than the corresponding sensitivities, the lowest being 45.5 % for *premature ventricular contraction*, whereas NPV is high for all conditions.

3.2 Power consumption

Based on the figures provided in [19], the firing activity and architecture of the sRNN on the DYNAP-SE chip translate into a dynamic power consumption of 286.1 μW . Assuming a static power draw of around 230 μW , total power consumption amounts to 516.1 μW . Together with the power estimates for those parts of the system that were simulated on a desktop computer (see Section 4.3), the total power consumption of the proposed system is less than 722.1 μW . An overview of the power figures for different components can be found in Table 4.

4 Discussion

4.1 Performance metric

In this work network performance is evaluated by means of two key values: sensitivity and specificity, which quantify the ability to detect anomalous patterns in the ECG input on the one hand, and on the other hand proneness to produce erroneous detection events.

In the data set, anomalous patterns usually do not occur in an isolated fashion but appear multiple times in close succession. Therefore missing one pathological heartbeat is not problematic as long as others are detected correctly. Hence, sensitivity analysis is done over segments of rhythms with the same label and any detection counts for the whole segment. However, even for the same anomaly type there may be variations between individual beats, in particular for recordings from different subjects. The network might be able to detect the anomaly with high reliability for some subjects and for others not. To make sure that such a scenario does not go undetected, segments always consist of contiguous heartbeats from single recordings.

The ratio of erroneous detection events and number of normal heartbeats provides an easy way to quantify specificity. Together with the combined duration of all normal beats it can be easily translated into the expected mean time between two false positives for an input that is free of anomalies, a number that gives an intuitive idea of specificity.

4.2 Network performance

The network reliably detects most of the presented anomalies, except for *premature ventricular contraction*. There is no apparent correlation between individual subjects and sensitivity for this label and it seems that this pattern is particularly difficult for the network to detect, indicating that performance depends to some extent on the type of anomaly that is to be found.

Which level of detection accuracy is adequate strongly depends on the individual application. For expert systems that autonomously classify heartbeats, high precision is crucial. On the other hand, in an assisted diagnosis scenario where the system only serves as a filter prior to diagnosis by a human expert, precision requirements may be reduced. For instance, long-term monitoring bio-signals from out-patients traditionally generate large amounts of data which raises challenges that can be alleviated by the suggested system: Diagnosing professionals can be directed to relevant sections without having to sift through the complete recording. The amount of data that needs to be stored and further processed can be reduced by filtering out irrelevant non-pathological sequences. This also allows for more compact devices with long battery life.

Regarding mean time between false positives (MTBFP) and specificity, overall values are lower than for individual anomalies. The reason is that taking into account detection events for all classes simultaneously implies that false positives will accumulate over classes. PPV is slightly low for some anomalies, suggesting that positive results are not to be interpreted as definite diagnosis but rather serve to initiate and support further analysis. It should be noted however, that because true positives and false negatives are

defined with respect to segments of multiple heartbeats, their numbers are lower than those for true negatives and false positives, which refer to individual heartbeats. This results in a disproportionately small PPV.

The chosen trade off between accuracy and MTBFP rate can depend on the precise use case. For example, a diagnosing professional can analyze an ECG much more efficiently if only two instead of 70 to 80 rhythms per recorded minute need to be considered. In other scenarios, sensitivity may be weighed up against the number of false alarms by changing detection thresholds. Furthermore, specificity can be improved by only searching for specific anomaly types, e.g. if a specific pathological state is suspected in advance. For example, if only *atrial premature beats* are considered, MTBFP is above 19 minutes with the current setup.

Nevertheless, promising techniques have been proposed that may improve performance significantly, but whose implementation would have exceeded the scope of this work: Instead of using random recurrent connections, network dynamics can be tailored to a given task by spectral analysis of the network [29] or techniques for balancing excitation and inhibition [30, 31], in-reservoir learning algorithms such as [32], or by using more traditional approaches like Backpropagation Through Time (BPTT) [33] to train a non-spiking recurrent neural network and using a transfer algorithm as in [34].

Finally, while the linear regression algorithm used in this work allows for efficient training of the readout weights, other learning models such as support-vector machines (SVMs) [35] may yield more suitable weights for this task. It may even be considered replacing the layer of readout units by a multilayer perceptron (MLP) trained with backpropagation [36].

4.3 Feasibility of preprocessing and readout stages

While the spiking neural network in this work was implemented on neuromorphic hardware, other parts of the system, namely the conversion from analog input to events, the readout stage and the generation of the binary output signal were simulated on a desktop computer. Furthermore, the ECG signal that was used had already been amplified and filtered, so that no preprocessing was necessary. In the following we will argue that an implementation of these stages and therefore of the full proposed system is feasible.

4.3.1 Preprocessing and conversion to events

Typical raw ECG sensor signals have an amplitude of a few millivolts and are generally amplified and noise-filtered. Suitable low-power, low-noise amplifiers have been proposed, for instance, in [37] and [38], each with a power consumption of less than $50 \mu\text{W}$. Similarly, lowpass and notch filters are proposed in [39] and [40], requiring less than $4 \mu\text{W}$ power. Implementations of asynchronous sigma-delta encoders that efficiently convert the analog signal to events are described in [17] and [41].

4.3.2 Readout

The readout units have not been set up as neurons on the neuromorphic processor for two reasons: the limited fan-in of the hardware neurons and the fact that all presynaptic weights of a neuron are the same for a given synapse type.

The fan-in is critical insofar as a readout unit ideally has access to the firing activity of all neurons in the hidden layers of the network. However, it is certainly possible to design neuromorphic processors with readout units that can subscribe to significantly larger numbers of presynaptic neurons. One example is the Reconfigurable On-line Learning Spiking Neuromorphic Processor (ROLLS) [42], which is similar to the DYNAP-SE used in this work and can be configured such that a single neuron has up to 130k synapses. Furthermore, on an algorithmical level the required fan-in can be reduced by variable selection methods and regularization rules that encourage a large number of zero-weights, such as LASSO [43].

Regarding weight quantization, synapses with individually tunable strengths and a precision of a few bits have been implemented in neuromorphic devices like the Dynap-SEL [44]. Although it may not suffice to quantize previously trained high-precision weights, there are learning algorithms such as [45], [46] and [47] that are designed to find suitable low-precision weights. Another potential solution is the application of transfer methods as suggested in [34].

Assuming similar power figures as for DYNAP-SE, we estimate that an implementation of a readout as described above, that produces a binary output indicating whether one of the detection thresholds is exceeded, would consume less than 90 μ W.

5 Conclusion

In this work we propose an always-on ultra-low power system that uses asynchronous neuromorphic hardware to perform real-time anomaly detection on a multi-channel ECG signal with a mean power consumption below one milliwatt. We implement a spiking recurrent neural network on a DYNAP-SE chip and demonstrate that it is able to reliably detect anomalous patterns in the ECG. Remaining parts of the proposed system can be realized with existing circuitry as part of a wearable health monitoring device.

Acknowledgment

This work is supported in part by EU-H2020 grant NeuRAM3 Cube (NEUral computing aRchitectures in Advanced Monolithic 3D-VLSI nano-technologies); by H2020 FET Proactive grant SYNCH (824162); by H2020 ECSEL grant TEMPO (826655); by the European Research Council under the Grant Agreement No. 724295 (NeuroAgents); and by aiCTX AG. Mr Bauer and Dr Muir performed this work as part of their duties at aiCTX AG.

References

- [1] M. Naghavi, A. A. Abajobir, C. Abbafati, K. M. Abbas, F. Abd-Allah, S. F. Abera, V. Aboyans, O. Adetokunboh, A. Afshin, A. Agrawal *et al.*, “Global, regional, and national age-sex specific mortality for 264 causes of death, 1980–2016: a systematic analysis for the global burden of disease study 2016,” *The Lancet*, vol. 390, no. 10100, pp. 1151–1210, 2017.
- [2] A. F. Members, J. J. McMurray, S. Adamopoulos, S. D. Anker, A. Auricchio, M. Böhm, K. Dickstein, V. Falk, G. Filippatos, C. Fonseca *et al.*, “ESC guidelines for the diagnosis and treatment of acute and chronic heart failure 2012: The task force for the diagnosis and treatment of acute and chronic heart failure 2012 of the european society of cardiology. developed in collaboration with the heart failure association (HFA) of the ESC,” *European journal of heart failure*, vol. 14, no. 8, pp. 803–869, 2012.
- [3] S. O. Gottlieb, S. H. Gottlieb, S. C. Achuff, R. Baumgardner, E. D. Mellits, M. L. Weisfeldt, and G. Gerstenblith, “Silent Ischemia on Holter Monitoring Predicts Mortality in High-Risk Postinfarction Patients,” *JAMA*, vol. 259, no. 7, pp. 1030–1035, 02 1988. [Online]. Available: <https://doi.org/10.1001/jama.1988.03720070030029>
- [4] S. O. Gottlieb, M. L. Weisfeldt, P. Ouyang, E. D. Mellits, and G. Gerstenblith, “Silent ischemia as a marker for early unfavorable outcomes in patients with unstable angina,” *New England Journal of Medicine*, vol. 314, no. 19, pp. 1214–1219, 1986, pMID: 2871485. [Online]. Available: <https://doi.org/10.1056/NEJM198605083141903>
- [5] S.-H. Liu, D.-C. Cheng, and C.-M. Lin, “Arrhythmia identification with two-lead electrocardiograms using artificial neural networks and support vector machines for a portable ECG monitor system,” *Sensors*, vol. 13, no. 1, pp. 813–828, 2013. [Online]. Available: <https://www.mdpi.com/1424-8220/13/1/813>
- [6] S. Kiranyaz, T. Ince, and M. Gabbouj, “Real-time patient-specific ECG classification by 1-D convolutional neural networks,” *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 3, pp. 664–675, March 2016.
- [7] N. Wang, J. Zhou, G. Dai, J. Huang, and Y. Xie, “Energy-efficient intelligent ECG monitoring for wearable devices,” *IEEE Transactions on Biomedical Circuits and Systems*, pp. 1–1, 2019.
- [8] E. D. beyli, “ECG beats classification using multiclass support vector machines with error correcting output codes,” *Digital Signal Processing*, vol. 17, no. 3, pp. 675 – 684, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1051200406001941>
- [9] O. A. Crdenas, L. M. Flores Nava, F. G. Castaeda, and J. A. Moreno Cadenas, “ECG arrhythmia classification based on fuzzy cognitive maps,” in *2019 16th International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE)*, Sep. 2019, pp. 1–4.
- [10] S. Gradl, P. Kugler, C. Lohmller, and B. Eskofier, “Real-time ECG monitoring and arrhythmia detection using android-based mobile devices,” in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Aug 2012, pp. 2452–2455.

- [11] A. Das, P. Pradhapan, W. Groenendaal, P. Adiraju, R. T. Rajan, F. Catthoor, S. Schaafsma, J. L. Krichmar, N. Dutt, and C. V. Hoof, “Unsupervised heart-rate estimation in wearables with liquid states and a probabilistic readout,” *Neural Networks*, vol. 99, pp. 134 – 147, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0893608017303003>
- [12] A. Amirshahi and M. Hashemi, “ECG classification algorithm based on STDP and R-STDP neural networks for real-time monitoring on ultra low-power personal wearable devices,” *IEEE Transactions on Biomedical Circuits and Systems*, pp. 1–1, 2019.
- [13] M. Marzencki, K. Tavakolian, Y. Chuo, B. Hung, P. Lin, and B. Kaminska, “Miniature wearable wireless real-time health and activity monitoring system with optimized power consumption,” *Journal of Medical and Biological Engineering*, vol. 30, pp. 227–235, 01 2010.
- [14] I. T. Iliev, I. I. Jekova, S. D. Tabakov, K. G. Koshtikova, and S. T. Iovev, “Telemetry of hospitalized high-risk patients with cardiovascular diseases,” in *2019 IEEE XXVIII International Scientific Conference Electronics (ET)*, Sep. 2019, pp. 1–4.
- [15] H. Mamaghanian, N. Khaled, D. Atienza, and P. Vandergheynst, “Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes,” *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 9, pp. 2456–2466, Sep. 2011.
- [16] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, “Neuromorphic electronic circuits for building autonomous cognitive systems,” *Proceedings of the IEEE*, vol. 102, no. 9, pp. 1367–1388, 9 2014.
- [17] F. Corradi and G. Indiveri, “A neuromorphic event-based neural recording system for smart brain-machine-interfaces,” *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 9, no. 5, pp. 699–709, 2015.
- [18] G. Indiveri, B. Linares-Barranco, T. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Häfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen, “Neuromorphic silicon neuron circuits,” *Frontiers in Neuroscience*, vol. 5, pp. 1–23, 2011. [Online]. Available: <http://www.frontiersin.org/Neuromorphic-Engineering/10.3389/fnins.2011.00073/abstract>
- [19] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, “A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs),” *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 12, no. 1, pp. 106–122, Feb. 2018.
- [20] E. Donati, M. Payvand, N. Risi, R. Krause, and G. Indiveri, “Discrimination of EMG signals using a neuromorphic implementation of a spiking neural network,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 13, no. 5, pp. 795–803, Oct 2019.
- [21] F. Corradi, S. Pande, J. Stuijt, N. Qiao, S. Schaafsma, G. Indiveri, and F. Catthoor, “ECG-based heartbeat classification in neuromorphic hardware,” in *2019 International Joint Conference on Neural Networks (IJCNN)*, July 2019, pp. 1–8.
- [22] C. Bartolozzi and G. Indiveri, “Synaptic dynamics in analog VLSI,” *Neural Computation*, vol. 19, no. 10, pp. 2581–2603, Oct 2007.

- [23] T. Delbruck, R. Berner, P. Lichtsteiner, and C. Dualibe, “32-bit configurable bias current generator with sub-off-current capability,” in *International Symposium on Circuits and Systems, (ISCAS), 2010*, IEEE. Paris, France: IEEE, 2010, pp. 1647–1650.
- [24] G. B. Moody and R. G. Mark, “MIT-BIH arrhythmia database,” 1992. [Online]. Available: <https://physionet.org/physiobank/database/mitdb/>
- [25] A. L. Goldberger, L. A. N. 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 (June 13), circulation Electronic Pages: <http://circ.ahajournals.org/content/101/23/e215.full> PMID:1085218; doi: 10.1161/01.CIR.101.23.e215.
- [26] W. Maass, T. Natschläger, and H. Markram, “Real-time computing without stable states: A new framework for neural computation based on perturbations,” *Neural Computation*, vol. 14, no. 11, pp. 2531–2560, 2002.
- [27] H. Jaeger, *Tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the “echo state network” approach*. GMD-Forschungszentrum Informationstechnik, 2002.
- [28] C. Langton, “Computation at the edge of chaos: Phase transitions and emergent computation,” *Physica D: Nonlinear Phenomena*, vol. 42, no. 1, pp. 12–37, 1990.
- [29] P. V. Aceituno, Y. Gang, and Y. Liu, “Tailoring artificial neural networks for optimal learning,” *CoRR*, vol. abs/1707.02469, 2017. [Online]. Available: <http://arxiv.org/abs/1707.02469>
- [30] M. Boerlin, C. K. Machens, and S. Denève, “Predictive coding of dynamical variables in balanced spiking networks,” *PLoS computational biology*, vol. 9, no. 11, p. e1003258, 2013.
- [31] A. Alemi, C. K. Machens, S. Deneve, and J.-J. Slotine, “Learning nonlinear dynamics in efficient, balanced spiking networks using local plasticity rules,” in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [32] W. Nicola and C. Clopath, “Supervised learning in spiking neural networks with force training,” *Nature communications*, vol. 8, no. 1, p. 2208, 2017.
- [33] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [34] X. He, T. Liu, F. Hadaeghi, and H. Jaeger, “Reservoir transfer on analog neuro-morphic hardware,” in *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 2019, pp. 1234–1238.
- [35] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [36] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning internal representations by error propagation,” California Univ San Diego La Jolla Inst for Cognitive Science, Tech. Rep., 1985.
- [37] Xiao Yang, Qi Cheng, Li-fei Lin, Wei-wei Huang, and Chao-dong Ling, “Design of low power low noise amplifier for portable electrocardiogram recording system applications,” in *2011 IEEE International Conference on Anti-Counterfeiting, Security and Identification*, June 2011, pp. 89–92.

- [38] M. Ghamati and M. Maymandi-Nejad, “A low-noise low-power MOSFET only electrocardiogram amplifier,” in *2013 21st Iranian Conference on Electrical Engineering (ICEE)*, May 2013, pp. 1–5.
- [39] R. Kumar, S. Sharma, and R. Goyal, “A low power low-noise low-pass filter for portable ECG detection system,” in *IJCTA*. International Science Press, 2016, vol. 9, no. 41, pp. 95–103.
- [40] C. Yehoshuva, R. Rakhi, D. Anto, and S. Kaurati, “0.5 V, ultra low power multi standard Gm-C filter for biomedical applications,” in *2016 IEEE International Conference on Recent Trends in Electronics, Information Communication Technology (RTEICT)*, May 2016, pp. 165–169.
- [41] P. Lichtsteiner, C. Posch, and T. Delbruck, “A 128x128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor,” *IEEE Journal of Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, Feb 2008.
- [42] N. Qiao, H. Mostafa, F. Corradi, M. Osswald, F. Stefanini, D. Sumislawska, and G. Indiveri, “A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses,” *Frontiers in Neuroscience*, vol. 9, no. 141, pp. 1–17, 2015.
- [43] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996. [Online]. Available: <http://www.jstor.org/stable/2346178>
- [44] C. S. Thakur, J. L. Molin, G. Cauwenberghs, G. Indiveri, K. Kumar, N. Qiao, J. Schemmel, R. Wang, E. Chicca, J. Olson Hasler, J.-s. Seo, S. Yu, Y. Cao, A. van Schaik, and R. Etienne-Cummings, “Large-scale neuromorphic spiking array processors: A quest to mimic the brain,” *Frontiers in Neuroscience*, vol. 12, p. 891, 2018. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fnins.2018.00891>
- [45] C. Zhu, S. Han, H. Mao, and W. J. Dally, “Trained ternary quantization,” *arXiv preprint arXiv:1612.01064*, 2016.
- [46] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, “Quantized neural networks: Training neural networks with low precision weights and activations,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6869–6898, 2017.
- [47] K. Helwegen, J. Widdicombe, L. Geiger, Z. Liu, K. Cheng, and R. Nusselder, “Latent weights do not exist: Rethinking binarized neural network optimization,” *CoRR*, vol. abs/1906.02107, 2019. [Online]. Available: <http://arxiv.org/abs/1906.02107>