Analysis of the Impact of Hardware Front-End Limitations on Vagus Nerve Electroneurogram-Based Seizure Detection

Sylvain Favresse, *Graduate Student Member, IEEE*, Antoine Nonclercq, *Senior Member, IEEE*, Riëm El Tahry, David Bol, *Senior Member, IEEE*, and Denis Flandre, *Senior Member, IEEE*

***Abstract*— *Objective:* The detection of seizures from the vagus nerve electroneurogram (VENG) is an emerging application allowing minimally invasive closed-loop vagus nerve stimulation for the treatment of epilepsy. Several VENG-specific front-end designs were presented earlier, but their specifications were approximate or arbitrary, leading to inefficient or sub-optimal designs in an applicative context. In this work, we analyze how front-end design choices and non-idealities impact seizure detection performance. The sensitivity analysis leads to the presentation of potentially optimal front-end designs with low power consumption. *Methods:* This work uses experimental VENG data from 8 rats, a behavioral model of the front-end circuits, and a seizure detection algorithm based on template matching. The studied front-end limitations are intrinsic noise, common-mode rejection, dynamic range, quantization resolution, and oversampling rate, studied individually in a sensitivity analysis. Optimal designs are then proposed based on sensitivity results. The performance of the seizure detection algorithm is characterized by metrics that are least affected by the small dataset size. *Results:* Noise and dynamic range are the main factors impacting algorithm performance. An optimal front-end design with 3-μV_{RMS} noise, 6-bit quantization, and 2.1-µW power consumption is presented and achieves perfect seizure classification on the dataset. *Conclusion:* The presented analysis enables the optimization of front-end specifications and a significant reduction in power consumption. *Significance:* This work bridges the gap between the performance of a seizure detection algorithm and the circuit-level specifications for ultra-low-power integrated bio-interface design.**

***Index Terms*— action potential (AP), analog circuit, digital circuit, power consumption, seizure detection algorithm, signal acquisition, vagus nerve electroneurogram (VENG).**

Received xxxxxxxxxxxx; revised xxxxxxxxxxxx and xxxxxxxxxxxx; accepted xxxxxxxxxxxx. Date of publication xxxxxxxxxxxx; date of current version xxxxxxxxxxxx.

Sylvain Favresse is a FRIA grantee of the Fonds de la Recherche Scientifique - FNRS.

S. Favresse, D. Bol and D. Flandre are with the ICTEAM institute, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium (emails: {sylvain.favresse, david.bol, denis.flandre}@uclouvain.be). A. Nonclercq is with the BEAMS department, Université Libre de Bruxelles, 1050 Brussels, Belgium (email: antoine.nonclercq@ulb.be). R. El Tahry is with the Department of Neurology, Center for Refractory Epilepsy, Cliniques Universitaires Saint-Luc, 1200 Brussels, Belgium (email: riem.eltahry@saintluc.uclouvain.be).

Digital Object Identifier [missing DOI]

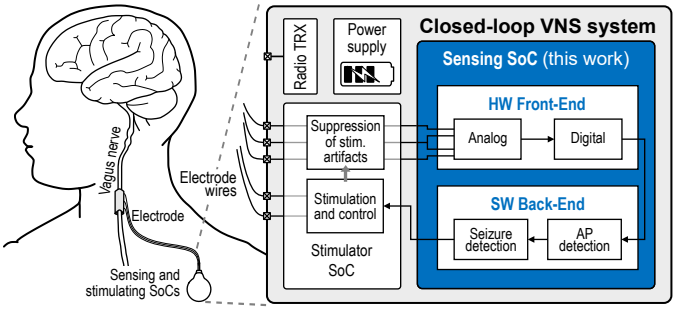

Fig. 1. Overview of an implanted closed-loop VNS system with details on the architecture of the sensing system-on-chip.

I. Introduction

CLOSED-loop vagus nerve stimulation (VNS) is a promising treatment for patients who have drug-resistant epilepsy [1]–[5]. Whereas open-loop VNS uses an ON-OFF stimulation pattern with a fixed duty cycle, the closed-loop counterpart adapts the stimulation parameters to the patient's instantaneous state. In this configuration, a stimulation is delivered when the device detects the onset of a seizure. A closed-loop VNS system thus requires accurate seizure detection, which can be achieved using several biomarkers. The AspireSR device delivers VNS in open loop or based on the detection of ictal tachycardia. The latter modality achieves better efficacy [6] but also suffers from a high rate of false alarms [7]. Electroencephalogram (EEG)-based seizure detection is the gold standard and provides good detection results [8]–[12], but scalp EEG monitoring devices are impractical to wear daily [13]. Intracranial EEG recordings can be highly accurate but require electrodes that are invasive and have a short lifespan [2], [14]. Several studies have shown that seizures could also be detected from the vagus nerve electroneurogram (VENG) [15]–[17], paving the way for a minimally invasive closed-loop VNS modality using the same electrode for both sensing and stimulation.

A potential system allowing such therapy is shown in Fig. 1. As illustrated, the stimulator is implanted in the chest and connected to the left vagus nerve via a cuff electrode [18]. The system includes a power supply, radio modules, and

two systems-on-chip (SoCs) for stimulation and sensing. The sensing SoC, the focus of this work, includes a hardware (HW) front-end that combines analog and digital blocks for signal acquisition. The software (SW) back-end processes the acquired signal to identify seizures and sends a trigger signal to the stimulator upon detection. Assuming the use of a 1,700-mAh 3.3-V battery as in commercial VNS devices [19], [20], the entire implanted system must operate on a power budget of around 50 µW for a 10-year lifetime, making low power consumption a primary design specification.

Our previous works focused on the design of the integrated front-end and showed that its power consumption is linked to non-idealities such as intrinsic noise [21]–[23]. Those trade-offs highlight the need to define accurate specifications before the design so as to achieve optimal and efficient operation.

Defining performance targets for biomedical front-ends is often done in a bottom-up approach, using approximate or arbitrary specifications such as a target signal-to-noise ratio (SNR). Other metrics (e.g., dynamic range and common-mode rejection ratio (CMRR)) are often maximized as a precaution. However, these approaches may result in non-functional or non-optimal designs at the system level. Accurately defining specifications based on knowledge of the application allows for minimizing the power consumption for a target functionality. Previous works proposed different methodologies to define accurate specifications. In [24], the impact of digital pre-processing parameters on EEG-based seizure detection is evaluated, but the effects of analog non-idealities are not addressed. In [25], neural recording front-ends are optimized for maximum spike sorting accuracy, but noise is not taken into account, and no explicit link is made with front-end power consumption. A cross-domain optimization is presented in [26] for a mixed-signal SoC performing electrocardiogram signal acquisition and machine learning (ML)-based arrhythmia classification. The analog front-end and the digital back-end are combined into a single optimization process to maximize ML algorithm accuracy while minimizing system energy consumption. This high-level HW/SW co-optimization allows for reaching a global optimum but requires a large amount of data for training and testing the circuits and algorithms. In the case of VENG-based seizure detection, no large dataset exists, preventing us from developing large ML algorithms. This work thus focuses on optimizing the front-end design for minimum power consumption while ensuring accurate seizure detection results using a limited dataset and a pre-existing algorithm.

A wide variety of algorithms exist to detect seizures from EEG signals, using namely template matching, wavelet transforms, and ML techniques [27]–[30]. Algorithms for VENG-based detection are scarcer. Although wavelets and ML could also be used in this context, only algorithms based on template matching have been proposed for seizure detection from the VENG [17], [31], [32]. In this work, we propose to use a previously validated template-matching algorithm and to test its sensitivity to non-idealities in signal acquisition HW. The algorithm used in this work first detects every action potential (AP) in the signal with a patient-specific approach [31], then uses the amplitude and frequency of APs to identify seizures [17]. The final classification parameters are trained

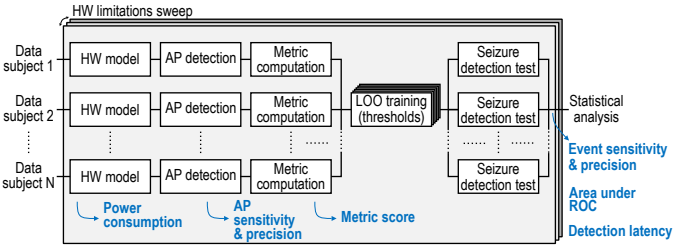

Fig. 2. Diagram of the different steps in the proposed study. Extracted performance metrics highlighted in blue.

based on the available dataset. Given that APs sensed on the vagus nerve with a cuff electrode typically have an RMS (root-mean-square) amplitude of around 7 µV [33], a first-order approach would consist in specifying a maximum input-referred noise (IRN) level below 0.7 µV to achieve an arbitrary SNR of 20 dB. Reaching such a low noise level would require a high current consumption in the instrumentation amplifier (IA), possibly wasting power [22].

This paper presents the evaluation of the impact of limitations, design choices, and non-idealities in the HW front-end on the performance of the algorithm that performs VENG-based seizure detection. Pre-recorded VENG data is used as input to a configurable model of the HW front-end. The output signal, altered by HW limitations, is fed to a seizure detection algorithm. The algorithm's accuracy is finally evaluated to assess the impact of the limitations inherent in the HW front-end. A sensitivity analysis is first performed on each limitation, one by one, by changing the HW model configuration. Then, new designs are proposed to define optimal front-end specifications that combine minimum front-end power consumption with maximum processing accuracy. The analysis and optimization proposed in this work are specific to the selected front-end architecture, algorithm, and application. The workflow and methodology can however be applied to any system that combines front-end circuits and signal-processing algorithms. For easy generalization, the implementations of the HW front-end model, the seizure detection algorithm, the classifier, and the raw performance results are available on GitHub [34].

In this paper, Section II presents the methodology employed in this work. Section III describes the front-end model and its validation. Section IV discusses the results of the sensitivity analysis and proposes and verifies optimal front-end designs. Section V discusses the transposition of this work's conclusions to human use, and Section VI concludes this work.

II. Analysis Methodology

A. Workflow

Fig. 2 presents the workflow used in this work to characterize the impact of HW front-end limitations. A set of experimental VENG data on 8 rats, including 2 controls, is used as input. A configurable HW model is built to characterize the impact of the HW front-end on the VENG signals. The output of the HW model, that represents the data stored in the sensing SoC's memory, is fed to the AP detection algorithm, which identifies APs in the signal. Then follows the seizure detection algorithm, which computes a

time-varying metric M_{det} that detects rapid changes in firing rate and AP amplitude. When M_{det} exceeds a predefined threshold, a seizure alarm is triggered. Leave-one-out (LOO) training is used to optimize the threshold parameters, and the algorithm performance is evaluated on the test set. This process is repeated to evaluate seizure detection performance across different HW limitations. The sensitivity of seizure detection performance to HW limitations is then evaluated. To do so, multiple performance indicators are extracted at different stages in the workflow, as illustrated in Fig. 2. Power consumption estimations are included in the HW front-end model and are based on its configuration. For example, the model includes a reduction in power consumption if the tolerated noise level increases, as discussed in Section III-B. AP detection accuracy is computed at the first algorithm stage, and a score is attributed to M_{det} after computing that metric. Next, the accuracy of seizure classification is also calculated. All performance metrics are detailed in Section II-E. The sensitivity of algorithm accuracy to HW configuration is finally characterized statistically to extract significant differences and trends.

B. Experimental data

The experimental VENG recordings used in this work were obtained from 8 anaesthetized male Wistar rats under the experimental conditions described in [17], [32], and approved by the University of Health Sciences Sector Laboratory Animal Protection Committee (2018/UCL/MD/001). The EEG was recorded through two differential channels using custom-made epidural electrodes, and the VENG was obtained with a micro cuff tripolar electrode with 2-mm contact spacing. The animals were placed in a Faraday cage, stabilized for 15 min after the surgery, and injected with a ketamine booster. Afterwards, a saline solution was injected into each animal. Acute seizures were induced in 6 animals (P1-6) with Pentylenetetrazol (PTZ), and the two remaining animals (S1-2) were used as controls. For rats P1-6, PTZ was injected until the tonic-clonic stage of the seizure, while the control rats received an equal volume of saline for 10 min. All recordings were kept to a duration of 31 min. Seizure events were annotated by an epileptologist using EEG and video recordings. Fig. 3(a) presents the timing of the recording experiments and the annotated seizure stages. Focusing only on generalized tonic-clonic seizures is a limitation of this work, as other types of seizures might be more difficult to detect [13].

A custom signal acquisition setup was made for this acute experiment, using off-the-shelf components. Power consumption was not a concern in this case, as the main objective was to acquire VENG signals as accurately as possible. The estimated performance of the recording setup for this experiment is a front-end IRN of 0.5 μV_{RMS}, an input impedance of 10 MΩ, a CMRR of 125 dB, a bandwidth (BW) of [0.012-12] kHz, a resolution of 16 bits, and a sampling rate of 80 kS/s [33]. The cuff electrode interface also generates noise, which is around 0.5 μV_{RMS} [35] and included in the recorded signal. It is assumed in this work that the experimental setup is ideal and that the recorded data, divided in amplitude by the acquisition gain, is identical to the input VENG signal.

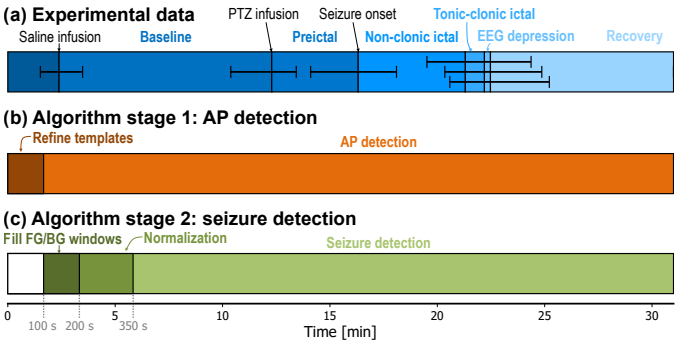

Fig. 3. Timing overview of (a) the experimental VENG data, (b) the AP detection stage of the algorithm and (c) the seizure detection stage of the algorithm. The VENG data intervals are reported as mean values, with error bars showing the minimum and maximum values across subjects. A detection is considered valid if it occurs after the PTZ infusion and up to the end of the EEG depression plus the duration of the BG window.

C. Seizure detection algorithm

The algorithm for embedded seizure detection implemented in this work is divided into two pipelined stages. The first stage, inspired by [31], detects spikes in the input signals using template matching. The second stage, inspired by [17], uses the frequency and amplitude of the detected APs to create a time-dependent metric and identify seizures. Those two stages of the algorithm execution are detailed next.

For the first stage, template matching works by computing the normalized cross-correlation (NCC) between an AP template and the signal in a sliding window. In case of multiple templates, the maximum NCC across all templates is selected at each time step. The signal is considered an AP when the NCC is above a predefined threshold of 75%. The NCC is computed on a 2-ms window moving with 0.05-ms steps. The AP detection algorithm from [31] adapts its AP templates to the patient by performing the detection on the entire signal once with generic templates, then clustering the detected APs. The cluster centroids are used as templates to run a second round of AP detection on the same signal. Using multiple templates allows to adapt to AP shapes from various nerve fibers, and doing a cluster-based template generation allows for a patient-specific AP detection [36]. However, running the algorithm twice on the same recording is not possible in an implant working in real time. In this work, another real-time-compatible method is used to make the templates patient-specific. As shown in Fig. 3(b), the first 100 s of each recording are used to refine the AP templates. That is done by starting the execution with a large number of templates (12 in this implementation), generated with the method from [31] by varying the AP width parameter. For each template, the number of detected APs is counted. At 100 s through the recording, the templates that detected less than 5% of the total number of detected APs are discarded for the rest of the execution. For long-term chronic recording, template adaptation could be performed regularly. The AP detection algorithm analyzes the input signal in 0.5-s non-overlapping time frames and outputs both the number of APs per frame and the mean AP amplitude. APs that feature an amplitude below

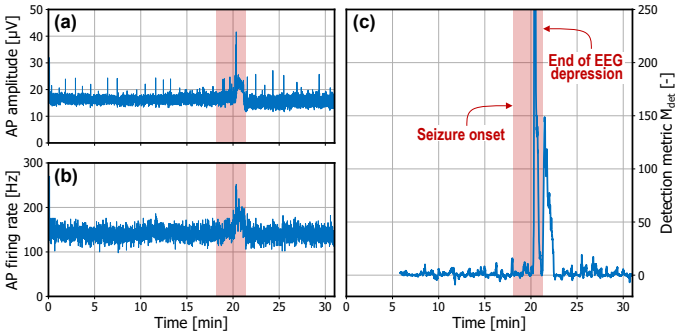

Fig. 4. Example of algorithm outputs on subject P4 with (a) the AP amplitude, (b) the AP firing rate and (c) the resulting detection metric that detects rapid changes in amplitude and firing rate. The annotated reference seizure event is highlighted in light red from onset to the end of EEG depression.

the RMS value of the signal in the time frame are considered to be noise and are discarded. APs with amplitudes above 5 times the RMS value are also discarded as they are associated with artifacts.

For the second algorithm stage, [17] tracks the firing rate and the AP amplitude using two abutting sliding windows: 10-s foreground (FG) and 90-s background (BG). Different trials in this work showed that these window lengths yield the best results in this implementation and correspond to the optimum reported in [17]. After that, the algorithm combines information on firing rate and AP amplitude. It first computes the relative slope of each quantity:

$$\Delta f_r = \frac{f_{\mathrm{FG}}}{f_{\mathrm{BG}}}, \quad \Delta U_r = \frac{U_{\mathrm{FG}}}{U_{\mathrm{BG}}}, \tag{1}$$

where $f_{\mathrm{FG/BG}}$ (resp. $U_{\mathrm{FG/BG}}$) is the mean firing rate (resp. AP amplitude) in each window, and Δf_r (resp. ΔU_r) is the relative slope in firing rate (resp. amplitude). The firing rate and amplitude slopes are normalized with

$$\Delta f_N = \frac{\Delta f_r - 1}{\mathrm{STD}[\Delta f_{r,\mathrm{norm}}]}, \quad \Delta U_N = \frac{\Delta U_r - 1}{\mathrm{STD}[\Delta U_{r,\mathrm{norm}}]}, \tag{2}$$

where $\mathrm{STD}[\Delta f_{r,\mathrm{norm}}]$ (resp. $\mathrm{STD}[\Delta U_{r,\mathrm{norm}}]$) is the standard deviation of the relative slope in firing rate (resp. amplitude) evaluated during the normalization time period illustrated in Fig. 3(c) and lasting 150 s. Since the actual seizure detection performance can only be evaluated after normalization, the final detection assessment starts at 350 s through the recording. For that purpose, the seizure detection metric M_{det} is computed from the normalized slopes:

$$M_{\mathrm{det}} = (\Delta f_N + 1)(\Delta U_N + 1). \tag{3}$$

Algorithm output examples are shown in Fig. 4, featuring the raw firing rate and mean amplitude in each time frame, and the seizure detection metric M_{det}. A thresholding process is applied to that metric for the binary classification. If M_{det} remains above the threshold T_M for a duration equal to or higher than a parameter D_M, a seizure alarm is triggered.

D. Leave-one-out training

The threshold parameters T_M and D_M are trained on available data to obtain a binary classification of seizure events. For the small dataset used in this work, LOO training is ideal to avoid training and testing on the same subject [37]. T_M and D_M are swept from 1 to 100 and 0.5 s to 10 s, respectively. D_M is kept at low values to favor short detection latencies. For each pair of threshold parameters, the performance on each subject is computed following the SzCORE framework for EEG-based seizure detection presented in [37]. The reference seizure annotations used in this work are the ones defined by an epileptologist, as described in Section II-B. Any overlap between reference and hypothesis is considered a valid detection. Pre-ictal tolerance is defined as the time difference between the start of the pre-ictal interval and the seizure onset, so that a detection during the pre-ictal interval is considered valid and results in a negative latency. The post-ictal tolerance is fixed to the duration of the BG window. Events separated by less than 90 s are merged, and the maximum hypothesis event duration is set to 5 min. With these parameters, each pair of thresholds gives a sensitivity and a precision result for seizure event detection. Those metrics are defined as

$$\text{Sensitivity} = \mathrm{TP}/(\mathrm{TP} + \mathrm{FN}), \tag{4}$$

$$\text{Precision} = \mathrm{TP}/(\mathrm{TP} + \mathrm{FP}), \tag{5}$$

where TP is the number of correct detections, FN the number of missed detections, and FP the number of false detections.

The sweep of threshold parameters yields a list of performance results, which are plotted on a sensitivity-precision curve. Only the nondominated results are kept. For evaluation on the test set, one single pair of threshold parameters is selected. If multiple results achieve the same optimal sensitivity and precision, the pair of threshold parameters with the lowest detection latency is selected. The remaining pairs of nondominated threshold parameters are also saved to calculate a receiver operating characteristic (ROC) curve on the test set. The leave-one-out training is repeated for each tested front-end configuration.

E. Performance indicators

Many performance indicators can be defined from the detection of seizure events, as detailed in [37]. For the objective of closed-loop neurostimulation, event-based metrics such as sensitivity and precision are most relevant. However, due to the limited size of the dataset used in this work, a single false positive or negative detection has a large impact on the mean performance, resulting in low statistical confidence in the event-based metrics. Other metrics are thus extracted earlier in the workflow, as illustrated in Fig. 2.

APs are frequent events, from 50 to 200 APs per second identified in the recordings. Thanks to those large numbers, statistically-significant performance indicators can be extracted from the AP detection step. AP detection accuracy is defined relatively to an ideal reference case in which the HW model is bypassed and the experimental VENG data is directly fed to the algorithm. APs are considered to be correctly detected if the time difference between the reference and the hypothesis AP is less than 1 ms, which is half the duration of an AP. Sensitivity and precision of AP detection are computed according to (4), (5).

After computing the seizure metric M_{det}, a metric score is evaluated. It is defined as the ratio of the RMS value of M_{det} during seizure to the RMS value of the same seizure metric outside of the valid detection range:

$$\text{Metric score} = \frac{\text{RMS}[M_{\text{det}}(\text{seizure})]}{\text{RMS}[M_{\text{det}}(\text{non-seizure})]}. \quad (6)$$

Expressed in dB, it describes changes in the detection metric during seizures, independently of the selected threshold parameters.

On top of that, the procedure for the selection of one pair of threshold parameters from many with the same performance also has a large impact on the final metrics. In this work, selection is performed to minimize detection latency, but other selection criteria could be used. Applying the selected pair of thresholds to the test set yields hypothesis seizure events that we compare with the reference events. Sensitivity, precision and detection latency are computed.

As explained in Section II-D, the entire set of nondominated threshold parameters is saved and applied to the test set to obtain a ROC curve. The area under the ROC curve (AUC) on the test set is a reported performance indicator that results from the LOO training but is independent of the selection of a single pair (T_M; D_M). The test set AUC is considered in this work to be a more reliable performance indicator than seizure sensitivity and precision.

F. Statistical analysis

The statistical framework in our study is designed to reach good statistical significance despite the limited size of the experimental dataset. The Spearman's rank correlation coefficient (r_s) is computed to assess whether a tested front-end limitation systematically degrades or improves algorithm performance. A non-parametric paired permutation test is then performed on the correlation coefficient to evaluate its statistical significance. For those tests, a p-value (p) is reported. To assess the magnitude of the impact of hardware limitations, the difference (Δ) in algorithm performance between the two extreme test conditions is computed. Its statistical significance is evaluated using a non-parametric paired permutation test on the difference, yielding a reported p-value. A significance level of 0.05 is used in this work. Due to the small size of the dataset, the performance indicators based on seizure events (metric score, seizure sensitivity and precision, detection latency, and test AUC) cannot be evaluated with statistical significance in most cases. The impact of each hardware limitation on seizure detection performance is thus studied qualitatively. It is concluded that a hardware limitation impacts seizure detection performance if correlated impacts on AP detection, metric score, and AUC are observed.

III. Hardware Front-End

A. Architecture

The integrated front-end considered in this work is the one presented in [23], which was simulated and experimentally characterized on a prototyped 65-nm CMOS chip called BREWER. The block diagram of the mixed-signal front-end is shown in Fig. 5. The analog blocks follow a conventional biomedical front-end structure with a low-noise band-pass IA, a 2nd-order band-pass analog filter (A-BPF), and an oversampling low-power analog-to-digital converter (ADC). The digital circuits include a 2nd-order infinite-impulse-response band-pass digital filter (D-BPF) and a decimation operation. The BIAS block generates analog reference currents for the IA and A-BPF modules. To reject electromyogram (EMG) interference in VENG recordings, two channels are necessary at the input and are summed by the IA. As highlighted in Fig. 5, the design choices and non-idealities of the hardware front-end whose impacts are studied in this work are (i) intrinsic noise of the IA, (ii) CMRR of that same amplifier, (iii) saturation in the analog filters, (iv) resolution of the ADC, and (v) oversampling rate (OSR). The BW required for AP acquisition is from 0.3 to 3.0 kHz [31], [38]–[40]. The cut-off frequencies of the A-BPF and D-BPF modules are thus fixed at those BW limits. For a correct definition of the frequency response of the filters, the sample rate at the output of the digital filters is fixed at 20 kS/s. The amplification gain is also fixed for the sensitivity analysis since the study of its impact would be redundant with the sensitivities to saturation voltage and ADC resolution.

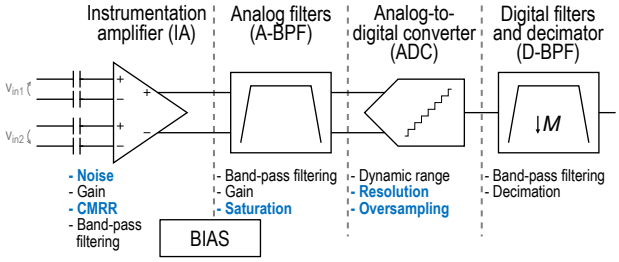

Fig. 5. Block diagram of the considered hardware front-end with annotations of the functions and non-idealities included in the behavioral model for each block. The limitations whose impact is studied in this work are highlighted in bold blue.

The proposed workflow involves sweeping the parameters and performance of the integrated circuits with experimental data as inputs. Such analysis cannot be done in measurements due to the limited reconfigurability of the prototyped front-end. The performance could be modified in transistor-level simulations, but the runtime would be far too high. A behavioral model of the front-end was thus constructed, including the modules and functions shown in Fig. 5. In the following sections, the behavioral model for each circuit block is detailed, along with the underlying assumptions. Each circuit block also has an associated power consumption model that is calibrated to match the simulated power breakdown reported in [23].

B. Instrumentation amplifier

In the IA model, noise is added to the input signal. The IA noise is modeled as a combination of white and pink random noise sources, generated with the Box-Muller transformation and the Voss-McCartney algorithm, respectively. The noise corner frequency f_c is fixed at 1 kHz, and it is assumed that f_c remains constant over the entire range of IRN values

considered in this work. A parasitic common-mode input signal is also included in the model, and attenuated by a finite CMRR. In biopotential recording applications, the common-mode input is typically composed of 50/60 Hz power line interference, EMG artifacts, and other unwanted signals. Those are located at low frequency and can take a wide range of amplitudes [41], [42]. The power line interference results in voltages with an amplitude in the 100 μV_{PP} range [43]. Considering recordings with a cuff electrode, the electromyogram signal has a power spectral density that decreases with the frequency [44], and an integrated amplitude around 1 mV [35]. In this work, a common-mode input is applied, representing all those parasitic signals. It is generated as a random signal with a $1/f$ power spectral density (PSD) and an amplitude of 1 mV over the [1-100k]-Hz BW.

The IA's power consumption is primarily determined by its noise requirements. As discussed in [22], [45], current consumption is inversely proportional to the square of the noise level assuming a constant noise efficiency factor (NEF). Therefore, the power consumption of the IA is modeled as

$$P_{\text{IA}} = (1.16\mu\text{V}/\text{IRN})^2 \times 8.06\mu\text{W}, \tag{7}$$

where IRN is the IA input-referred noise in the $\pi/2\times$[0.3-3.0]-kHz BW. The 1.16 µV and 8.06 µW values are selected to match the IRN and power consumption of the IA measured on the prototype chip, with an NEF of 2.2. The constant-NEF model relies on the assumption that flicker noise decreases with the same proportions as thermal noise, such that f_c remains constant over a wide range of IRN values. As discussed in [22], thermal noise decreases with an increase in power consumption, whereas flicker noise is mainly dependent on transistor gate area. In the lowest-noise conditions tested in this work, f_c remains constant provided that a sufficient area is allocated to the IA. In the highest-noise conditions, thermal noise is so high that f_c tends to be lower than 1 kHz. However, that effect is neglected in this first-order model.

C. Analog filters

The A-BPF module includes gain and filtering operations. Additionally, saturation of the output signal is modeled by a clipped sine function defined as

$$v_2 = \begin{cases} -v_{\text{sat}} & \text{if } v_1 < -v_{\text{sat}}\pi/2, \\ v_{\text{sat}} & \text{if } v_1 > v_{\text{sat}}\pi/2, \\ v_{\text{sat}} \sin\left(\frac{v_1}{v_{\text{sat}}}\right) & \text{otherwise}, \end{cases} \tag{8}$$

where v_1 is the unsaturated output, v_2 is the output after saturation and v_{sat} is the saturation voltage. This saturation limits the dynamic range and results in distorted output signals if the input amplitude is close to or above v_{sat}. To avoid this saturation effect, two options are possible: (i) increase v_{sat}, or (ii) decrease the gain of the analog front-end. Decreasing the gain also impacts the quality of the digital representation of the signal if the ADC range and resolution are not adapted.

The power consumption of the analog filters depends on the amplifier's gain-bandwidth product (GBW). It is modeled as

$$P_{\text{A-BPF}} = \text{GBW}/130\text{kHz} \times 2^{N_b}/2^{12} \times 2.25\mu\text{W}, \tag{9}$$

where N_b is the number of bits in the ADC. The impact of the ADC resolution is due to variations in the input capacitance of the ADC with N_b, impacting the load of the analog filters. The considered capacitive load is limited to values above 0.5 pF to account for parasitic capacitances. At a 12-bit resolution, the ADC input capacitance is 1.2 pF.

D. Analog-to-digital converter

The ADC is modeled as a digitization operation (i.e., sampling and quantization) on the signal. No degradation of the effective number of bits is modeled compared to the target resolution. The OSR is defined as the ratio between the ADC sampling rate and the fixed 20-kS/s output rate. The default OSR of 8 thus results in an ADC sampling rate of 160 kS/s. For the quantization part, a fixed differential input range of $\pm$0.91 V is considered. The default resolution is 12 bits.

The ADC power consumption model is given by

$$P_{\text{ADC}} = f_s/20\text{kSps} \times \text{OSR}/8 \times 2^{N_b}/2^{12} \times 0.83\mu\text{W}, \tag{10}$$

where f_s is the output sample rate. It is considered that the ADC has a power consumption proportional to the sampling rate and to the number of quantization levels.

E. Digital filters

The D-BPF module performs band-pass filtering in the digital domain and applies decimation when the ADC operates at a sampling rate above 20 kS/s. No non-idealities are added to this module. The power consumption is modeled as

$$P_{\text{D-BPF}} = f_s/20\text{kSps} \times \text{OSR}/8 \times N_b/12 \times 1.50\mu\text{W}, \tag{11}$$

where the 1.50-µW power consumption with default parameters is obtained from post-synthesis simulations of the prototype chip.

F. Biasing

The BIAS module generates reference currents for the IA and the analog filters. In the behavioral model, it only adds power consumption to the front-end, modeled as

$$P_{\text{BIAS}} = (P_{\text{IA}}+P_{\text{A-BPF}})/(8.06\mu\text{W}+2.25\mu\text{W})\times 1.76\mu\text{W}. \tag{12}$$

In multi-channel front-ends, the bias module is shared across channels and its power consumption is only accounted once.

G. Behavioral model summary and validation

The behavioral model combines the modules listed previously. It works with a time step of 1.56 µs. The total power consumption of the front-end is the sum of the power consumption of each module.

The model is validated by comparing it to measurements and transistor-level SPICE simulations of the prototyped front-end using Siemens Eldo. Fig. 6 shows the validation results for the noise, frequency response, and output waveforms. In Fig. 6(a)-(b), it can be seen that the generated noise is Gaussian and closely follows the target PSD. The frequency response of the behavioral model also faithfully follows the

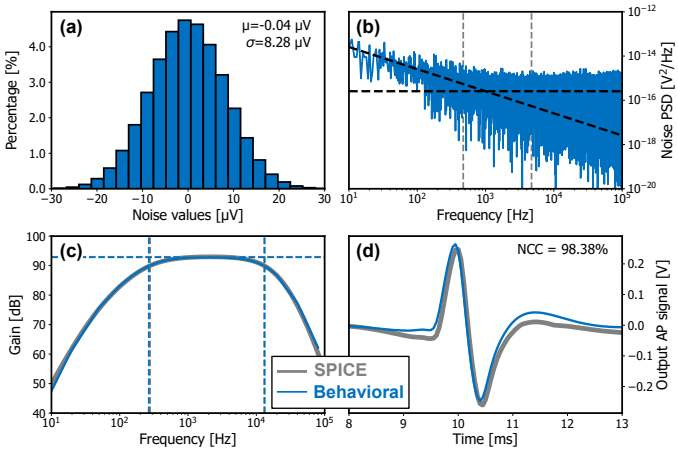

Fig. 6. Validation of the behavioral model for the HW front-end. (a) Histogram of the generated noise. (b) Power spectral density of the same noise, with black dashed lines indicating the target white and pink PSDs and the grey dashed lines indicating the equivalent noise bandwidth. (c) Comparison of the Bode diagram from SPICE and behavioral simulations. (d) Front-end output, for noiseless SPICE and behavioral simulations with 20-μV_{PP} AP as input, reaching an NCC of 98.38%.

TABLE I
PERFORMANCE AND PARAMETERS COMPARISON BETWEEN THE PROTOTYPE CHIP AND THE DEFAULT BEHAVIORAL MODEL

	Prototype chip	**Behavioral model**
IA BW [kHz]	[0.27-43]†	[0.27-43.5]
IA IRN* [μV_{RMS}]	1.16	**1.16**
IA CMRR [dB]	74.8	**75**
A-BPF BW [kHz]	[0.05-13.9]†	[0.05-3.0]
Saturation voltage [V]	0.76†	**0.75**
Total gain [dB]	93.0	92.8
ADC f_s [kS/s]	160▷	**160**▷
ADC resolution [#bits]	12	**12**
D-BPF BW [kHz]	[0.18-4.0]†	[0.2-3.0]
Output f_s [kS/s]	20	20

*Integrated on $\pi/2\times$[0.3-3.0] kHz BW. †Simulated. ▷OSR=8

one obtained from SPICE simulations in Fig. 6(c), matching the total gain and cut-off frequencies. The noiseless transistor-level and behavioral simulations are also compared with an AP waveform at the input, as shown in Fig. 6(d). For AP amplitudes above 5 μV_{PP}, the NCC between the two models' outputs exceeds 95%, confirming that the effects on APs are well represented in the behavioral model.

The model's default parameters used in this work correspond to the prototype measurements and post-layout simulations with reduced bandwidth. Table I details the simulated performance and the corresponding default model parameters. The values highlighted in the behavioral model are those that vary in the sensitivity analysis in this work. The behavioral model takes about 45 min of CPU time to simulate 1h of operation of the front-end, compared to about 23,000 years of CPU time for the transistor-level post-layout simulations. Relying on behavioral models enables the capture of effects on different time scales in a single simulation.

IV. IMPACT OF HW LIMITATIONS

This section studies the impact of HW front-end limitations on the performance of the seizure detection algorithm following the workflow presented in Fig. 2. First, a sensitivity analysis of each HW limitation is conducted to analyze the impact on power consumption and seizure detection. As illustrated in Fig. 2, the whole sequence is executed multiple times to sweep over the considered HW limitations that are highlighted in Fig. 5. Each HW parameter is studied independently: for each sweep, all parameters are fixed to the default value given in Table I except for the one limitation that is studied. Second, the results of the sensitivity analysis are used to propose new designs that achieve good seizure detection performance while minimizing power consumption.

A. Sensitivity analysis

The sensitivity analysis studies the performance of the seizure detection algorithm and the power consumption of the front-end when sweeping different front-end design choices and non-idealities. IA noise, IA CMRR, A-BPF saturation, ADC resolution, and ADC OSR are studied. The sweep of the number of bits of the ADC is performed at a fixed full scale, which affects the ADC resolution. Fig. 7 shows the results of the sensitivity analysis including each performance indicator (rows) and each limitation sweep (columns).

1) Noise: IA noise significantly impacts both power consumption and algorithm performance, as shown in Fig. 7(a). At low noise levels, the IA is the bottleneck of the system regarding power consumption. Closely studying the maximum noise level that can be tolerated is thus essential to improve the energy efficiency of the front-end. Increasing IA noise reduces AP sensitivity ($r_s = -1.00$, $p = 0.00$) and precision ($r_s = -1.00$, $p = 0.00$). The metric score and AUC remain approximately stable up to noise levels of around 8 μV. It can be noted that for that IRN value, AP sensitivity and precision are equal to 41.1% and 24.4% respectively. This indicates that poor overall AP detection accuracy can still lead to good seizure detection. Fig. 8(a)-(b) demonstrates that effect, where the mean AP amplitude and firing rate in the foreground window are plotted for different noise levels. For an IRN of 0.5 or 2.0 μV, it can be seen that noise has almost no additional effect: even APs with the lowest amplitude can be detected. For an IRN of 8.0 μV, only APs with the highest amplitude are correctly identified, increasing the mean detected AP amplitude. False detections also increase the firing rate. Seizures are still visible, as the mean AP amplitude is higher during those events, leading to a larger number of detectable APs and a higher observed amplitude. Up to that noise level, the effects of noise on AP amplitude and firing rate are approximately constant in time, leading to almost no impact on the seizure detection metric in Fig. 8(c), thanks to the slope computations in (1) that act similarly to a derivative operator. With a noise level of 32 μV, most APs are invisible, and no significant temporal variations are observed, resulting in M_{det} being close to zero at all times. When comparing the AP detection metrics with the minimum and maximum noise levels tested, a significant difference is observed in sensitivity ($\Delta = -54.5\%$, $p = 0.02$) and precision ($\Delta = -44.5\%$, $p = 0.02$). The decrease in AP detection performance with increasing noise reflects in a decrease in the metric score,

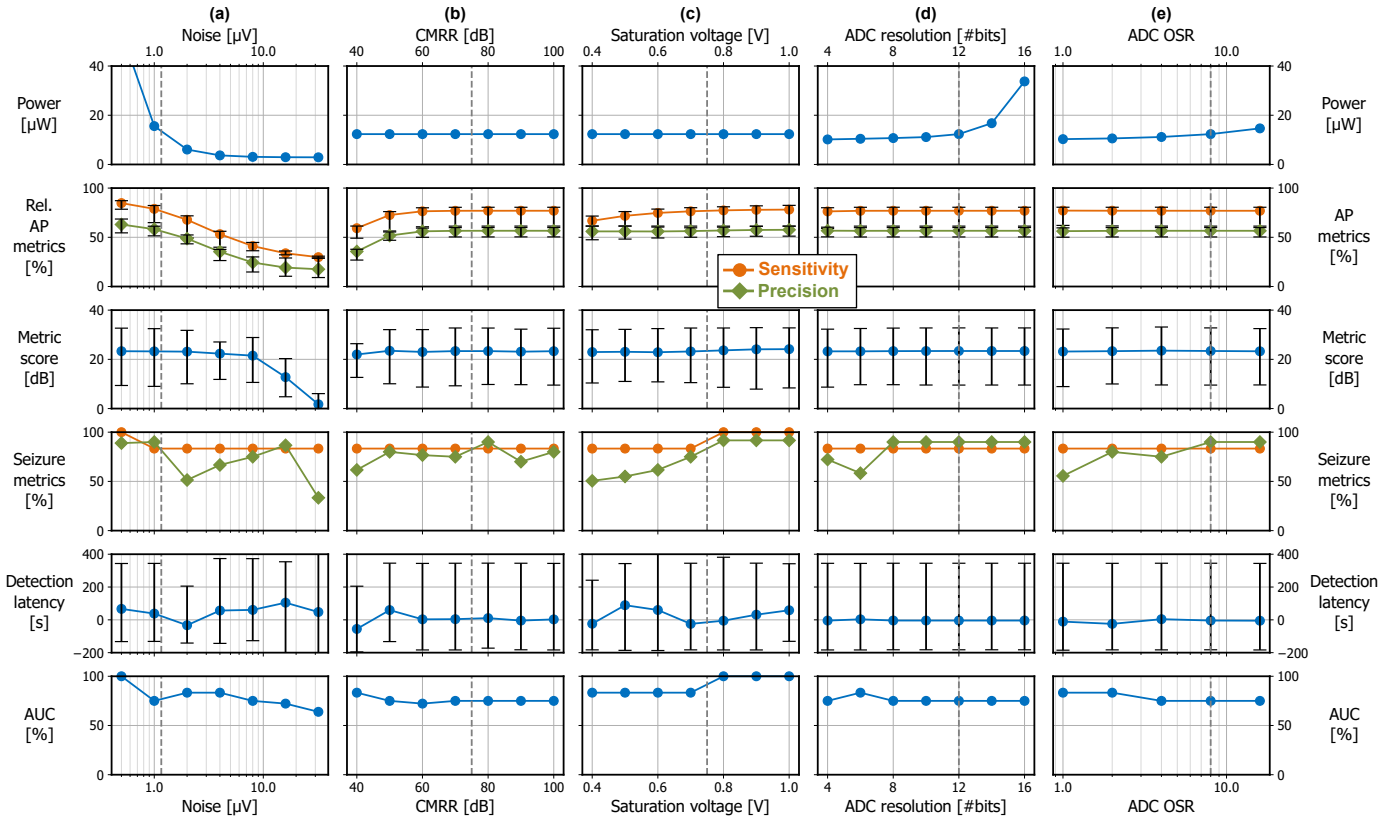

Fig. 7. Sensitivity analysis results for the impact of hardware limitations on front-end power consumption and seizure detection performance. The gray lines indicate the default values for each limitation, used in the other sensitivity analyzes. AP and seizure metrics combine sensitivity and precision. The error bars for AP metrics, metric score and detection latency indicate the min/max values among all tested subjects. The main plot for each performance metric indicates the mean value. Tested hardware limitations: (a) noise, (b) CMRR, (c) saturation voltage, (d) ADC resolution, and (e) ADC OSR.

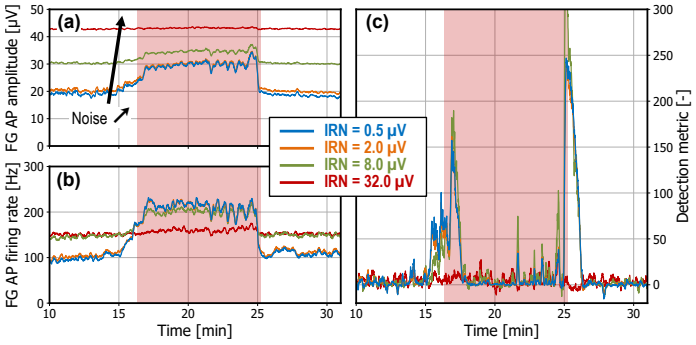

Fig. 8. Seizure detection on subject P1 for different levels of added noise. (a) mean AP amplitude in the foreground window, (b) mean AP firing rate in the foreground window, (c) seizure detection metric.

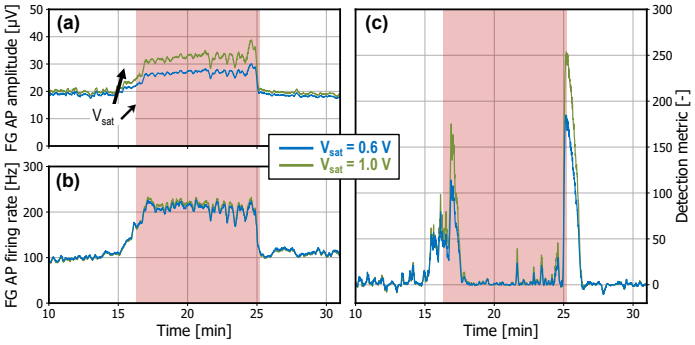

Fig. 9. Seizure detection on subject P1 for different levels of saturation voltage. (a) mean AP amplitude in the foreground window, (b) mean AP firing rate in the foreground window, (c) seizure detection metric.

although not significant ($r_s = -0.48$, $p = 0.07$). Between the lowest and highest tested noise levels, there is a significant difference in metric score ($\Delta = -18.8$ dB, $p = 0.02$). The observed effects on AP detection metrics, metric score, and AUC are correlated, showing the strong impact of front-end noise on seizure detection accuracy.

2) Common-mode rejection: The effect of finite IA CMRR is modeled by adding a random 1-mV signal as a common-mode input. For low CMRR values, Fig. 7(b) shows that the effects on the various performance indicators are similar to those of noise. For CMRR $>$ 60 dB, the common-mode input is negligible compared to the intrinsic noise of the IA. By comparing the noise and common-mode signal PSDs, it can be shown analytically that a CMRR larger than 57 dB is sufficient to keep the common-mode signal below the intrinsic noise at all frequencies, considering the numerical values assumed in this work. Overall, increasing the CMRR improves AP sensitivity ($r_s = 0.95$, $p = 0.00$) and precision ($r_s = 0.95$, $p = 0.00$). Comparing the case of CMRR $= 40$ dB to CMRR $= 100$ dB gives significant differences in AP sensitivity ($\Delta = 17.4\%$, $p = 0.02$) and precision ($\Delta = 20.6\%$, $p = 0.02$). Regarding the metric score, no monotonic impact

($r_s = -0.29$, $p = 0.07$) and no significant difference ($\Delta = -0.39$ dB, $p = 0.39$) can be determined. No impact on the AUC is observed either. It is concluded that the CMRR level should be chosen to keep the common-mode signal impact below the noise level, but targeting a CMRR value in the range of 100 dB is not required. The impact of CMRR on power consumption is not modeled, but it should be noted that higher CMRR implies greater analog design complexity. The quantitative results with each CMRR value are only valid with the common-mode input amplitude assumed in the model. If a higher common-mode input is expected, a higher CMRR should be targeted, i.e., for a 10-mV common-mode input, the minimum CMRR limit increases from 60 dB to 80 dB.

3) Saturation: Limiting the saturation voltage is equivalent to reducing the dynamic range of the analog blocks in the front-end. Due to the high gain, it clips the waveforms of high-amplitude APs. It affects the estimation of AP amplitude and reduces the mean amplitude of detected APs (Fig. 9(a)), but also reduces the correlation between the APs and the templates, slightly decreasing the mean detected AP firing rate (Fig. 9(b)). Increasing the saturation voltage thus increases AP sensitivity ($r_s = 0.67$, $p = 0.00$) and precision ($r_s = 0.74$, $p = 0.00$). Between $v_{\text{sat}} = 0.4$ V and $v_{\text{sat}} = 1.0$ V, there is also a significant difference in AP sensitivity ($\Delta = 10.4\%$, $p = 0.03$) and precision ($\Delta = 1.9\%$, $p = 0.01$). As shown in Fig. 9(c), a reduced dynamic range affects the seizure detection metric, but with a lesser impact than IA noise. In Fig. 7(c), it can be seen that the impact of the saturation voltage on the metric score is systematic ($r_s = 0.36$, $p = 0.01$), but the difference between the extreme values is not statistically significant ($\Delta = 1.13$ dB, $p = 0.17$). Given that the impacts on AP detection, metric score and seizure metrics are correlated, it is qualitatively concluded that increasing the saturation voltage improves seizure detection performance. The impact of the saturation voltage on power consumption is not modeled. However, increasing the saturation voltage increases the design complexity and may have second-order impacts on the power consumption.

4) ADC resolution: Fig. 7(d) indicates that, due to the low SNR, ADC resolution has little impact on AP and seizure detection performance. Increasing the number of bits systematically improves AP sensitivity ($r_s = 0.78$, $p = 0.00$) and results in a significant difference between 4 and 16 bits ($\Delta = 0.7\%$, $p = 0.02$). AP precision, however, does not significantly improve with N_b ($r_s = 0.23$, $p = 0.16$), and the difference between the extreme values cannot be proven ($\Delta = -0.1\%$, $p = 0.28$). The correlation with the metric score is not significant ($r_s = 0.05$, $p = 0.30$), and that metric shows no significant difference between the smallest and largest tested numbers of bits ($\Delta = -0.0$ dB, $p = 0.50$). At high N_b, ADC resolution becomes the bottleneck in power consumption due to its impact on the power of the analog filters, the ADC and the digital filters. Low numbers of bits should thus be preferred.

5) ADC oversampling: The OSR also has little impact on algorithm performance, as shown in Fig. 7(e). Increasing the OSR reduces AP sensitivity ($r_s = -0.63$, $p = 0.02$), but improves AP precision ($r_s = 0.77$, $p = 0.02$). When comparing OSR = 1 to OSR = 16, there is a small difference in AP sensitivity ($\Delta = -0.2\%$, $p = 0.05$) and AP precision ($\Delta = 0.4\%$, $p = 0.02$), but that is not visible on the metric score ($\Delta = 0.1$ dB, $p = 0.32$). The OSR impacts the power consumption and should thus be kept low.

TABLE II
QUALITATIVE EVALUATION OF THE IMPACTS ON PERFORMANCE INDICATORS WHEN HW LIMITATIONS INCREASE

	Noise ↗	CMRR ↗	v_{sat} ↗	ADC res. ↗	ADC OSR ↗
Power efficiency	++	/	/	--	-
AP detection	--	+	+	/	/
Metric score	--	/	+	/	/
Seizure detection	--	/	+	/	/

6) Qualitative assessment: Table II summarizes the qualitative observations that can be made from the sensitivity analysis presented in Fig. 7. It can be deduced that noise is the major limitation impacting performance. The saturation voltage also has a more limited impact. The ADC resolution and OSR do not significantly affect seizure detection performance and should therefore be chosen based on power consumption. The CMRR only degrades the algorithm performance with very low CMRR values.

B. Candidate optimal designs

The quantitative and qualitative conclusions from the sensitivity analysis can be leveraged to propose new designs that achieve good seizure detection performance while minimizing power consumption. The front-end configuration is optimized manually rather than with an automated optimization algorithm to avoid finding local optima that are not based on previous observations and are not statistically representative. The optimal operating points are shown on a power-AUC curve in Fig. 10, with a color bar indicating AP sensitivity. Table III details the front-end configuration for each design presented, with a highlight on each parameter modified compared to the previous design.

Starting from the prototype chip (O), the design choices that do not significantly impact the seizure detection performance are first adapted. The CMRR is changed to 60 dB to simplify the analog design process while keeping the common-mode signal negligible compared to noise. The ADC resolution and OSR are reduced to 4 bits and 1, respectively, to reduce power consumption. The saturation voltage is then increased to 0.9 V at point A to reach an AUC of 100%, as expected from Fig. 7(c). Targeting a saturation voltage higher than 0.9 V would not be useful since the ADC's full-scale range is limited to ± 0.91 V. Design A represents a 20% improvement in power consumption without changing the noise specifications. However, increasing the dynamic range of an analog amplifier can be challenging at low supply voltages. An alternative would be to keep the saturation voltage at 0.75 V while reducing the AFE gain by 3 dB. As a result, the output signal amplitude is reduced and less affected by the limited dynamic range. Nevertheless, reducing the output signal amplitude means that the full ADC range is no longer utilized, thereby

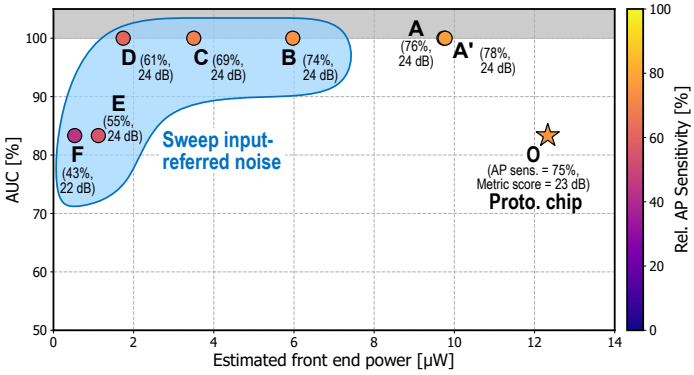

Fig. 10. Power, AUC and AP sensitivity of the proposed optimal front-end specifications, compared to the prototype chip. AP sensitivity and metric score values are indicated for each proposed design.

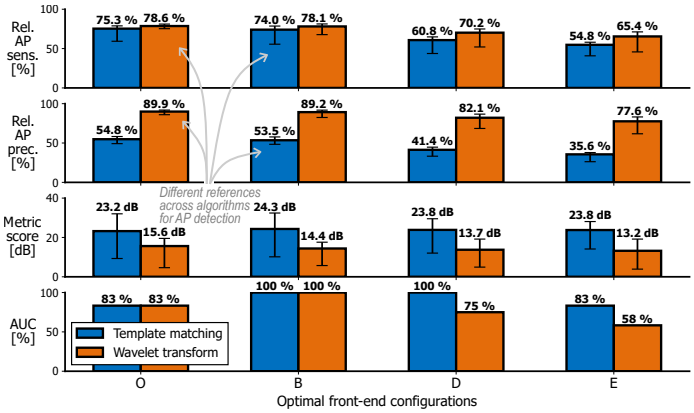

Fig. 11. AP and seizure detection accuracy using template matching and wavelet transform algorithms, in four different optimal front-end configurations. AP detection accuracy is relative to algorithm runs when bypassing the front-end model. Reference AP detection is thus different across algorithms.

TABLE III
FRONT-END PARAMETERS OF THE PROPOSED OPTIMAL DESIGNS

	Noise [µV]	CMRR [dB]	Sat. voltage [V]	ADC res. [# bits]	ADC OSR [-]	Gain [dB]
O	1.16	75	0.75	12	8	92.8
A	1.16	**60**	**0.9**	**4**	**1**	92.8
A'	1.16	**60**	0.75	**6**	**1**	**89.8**
B	**1.5**	60	0.75	6	1	89.8
C	**2.0**	60	0.75	6	1	89.8
D	**3.0**	60	0.75	6	1	89.8
E	**4.0**	60	0.75	6	1	89.8
F	**8.0**	60	0.75	6	1	89.8

reducing the input-referred minimum detectable signal. To account for this effect, the ADC resolution is increased to 6 bits with a negligible power overhead, as can be seen with point A' in Fig. 10. For designs B to F, all front-end parameters are kept identical to the ones in point A' except for noise, which is progressively increased. Due to the strong dependence of power consumption on IA noise, the former is drastically reduced by this operation. Design D offers the best trade-off, with an AUC of 100% on the dataset and a power consumption of 1.7 µW, achieved by tolerating an input-referred noise of up to 3 μV_{RMS}. The quantitative results are approximate, but several conclusions can be made. First, the CMRR and the saturation voltage are kept at low values to ease the analog design, with no modeled impact on power consumption. Second, the ADC resolution and OSR can be significantly reduced compared to those of the prototype chip to reduce front-end power consumption. Finally, noise is the major factor impacting both power consumption and seizure detection performance, and a relatively low SNR can be tolerated to improve power efficiency.

To confirm the trends in seizure detection accuracy across the proposed optimal designs, they were tested with another seizure detection algorithm. Inspired by [46], it uses wavelet transforms instead of template matching to detect APs. The seizure detection and classification principles based on AP features remain the same as those for the template-matching algorithm. Fig. 11 shows the side-by-side results from both algorithms for four of the optimal designs introduced in Fig. 10. It can first be seen that the wavelet-based algorithm maintains good AP detection sensitivity and precision in all configurations. The wavelet-based algorithm performs well for AP detection in the presence of noise. However, no absolute comparison can be made with the template-matching technique, since the AP detection ground truth differs between the two tested algorithms. Second, all accuracy metrics in Fig. 11 indicate that the wavelet-based algorithm's accuracy is degraded when noise increases, although less sharply than the template-matching algorithm's for AP detection. The optimal front-end configuration with the wavelet-based algorithm is close to design B, with an IRN level of 1.5 μV_{RMS} and an estimated front-end power of 6 µW. While the exact optimal results are thus different from one algorithm to the other, the trends in algorithm performance evolution over the tested designs remains similar. The better noise immunity of the template-matching algorithm in terms of metric score and AUC might be linked to its poorer AP detection accuracy by forcing the algorithm to only focus on APs with higher amplitude for seizure classification, which are the most visible seizure markers.

Fig. 12 shows the power breakdown in the configuration of the prototype chip, and for design D. Thanks to the less stringent specifications in design D, the power consumption of all modules is strongly reduced. The major absolute decrease in power consumption is on the IA, mainly through the increase in tolerated IRN, bringing the estimated IA power down to 1.2 µW. To ensure the feasibility of the ultra-low-power front-end, design D is implemented at transistor level. Compared to the prototype chip, the same architecture is followed but various components such as resistors and transistors are resized. As a main difference, the bias current in the IA is decreased, resulting in a higher noise level. SPICE simulations reveal that the proposed front-end meets the specifications with a power budget of 2.1 µW, which is only slightly higher than the expected 1.7 µW. As shown in Fig. 12, the main difference comes from the power consumption of the BIAS module, which cannot be decreased as much as anticipated. To validate the behavioral model in this optimal configuration, its outputs are compared to SPICE simulations when using AP stimuli at the inputs, as done in Fig. 6(d)

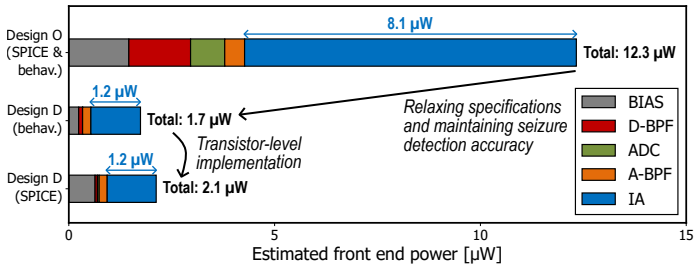

Fig. 12. Evolution of power consumption breakdown from the prototyped front-end to the optimal proposed design.

TABLE IV
BEHAV./SPICE COMPARISON OF OPTIMAL DESIGN D

	Behavioral model	**SPICE**
IA BW [kHz]	[0.27-43]	[0.27-15.5]
IA IRN [μV_{RMS}]	3.0	**2.55***
IA CMRR [dB]	60	60
A-BPF BW [kHz]	[0.05-3.0]	[0.05-3.6]
Saturation voltage [V]	0.75	0.78
Total gain [dB]	89.8	89.8
D-BPF BW [kHz]	[0.20-3.0]	[0.21-3.2]
P_{IA} [µW]	1.20	**1.19**
$P_{A\text{-}BPF}$ [µW]	0.20	**0.20**
P_{ADC} [µW]	0.002	**0.04**
$P_{D\text{-}BPF}$ [µW]	0.09	**0.06**
P_{BIAS} [µW]	0.28	**0.64**
Total power [µW]	1.74	2.12

* Expected to be higher in measurements

for the front-end on the prototype chip. The comparison reveals an NCC higher than 90% for AP amplitudes equal to or above 10 μV_{PP}, confirming that the behavioral model correctly represents analog non-idealities in that ultra-low-power configuration. Table IV compares the behavioral and SPICE simulations of design D. It can be seen that the IA consumes the same power as anticipated, with an IRN lower than 3.0 μV_{RMS}. Maintaining a margin on the noise level is a good practice in schematic simulations to account for additional noise from layout parasitics and other noise sources. The power consumption of the ADC is higher than predicted by the behavioral model, mainly due to components such as the comparator and the digital core whose power consumption does not scale with 2^{N_b}. However, the impact of the ADC on the total power consumption remains negligible.

V. PERSPECTIVES FOR HUMAN VENG RECORDING

The analyses proposed in this paper depend on experimental VENG data recorded on rats. Although the PTZ rat model is widely used to study epileptic seizures [16], [47], [48], it presents some differences with chronic human recordings. While those differences cannot be evaluated quantitatively due to the lack of human data, several qualitative observations can be made. First, PTZ experiments are conducted on anaesthetized rats. Devices implanted for chronic use suffer from significantly more motion artifacts, which can increase common-mode disturbance and degrade signal quality. Second, the vagus nerve is larger in humans than in rats, by a factor up to 10 [49]. Consequently, a larger cuff electrode can be used in humans, which results in a larger signal amplitude [33]. Therefore, a higher noise level could be tolerated in humans. Finally, the larger nerve diameter also results in lower fiber selectivity when using extraneural electrodes due to the larger distance between electrode and nerve fiber [50]. As a result, APs from the innermost fibers are more complicated to detect, impacting the recorded firing rate and AP shape diversity.

VI. CONCLUSION

To the author's best knowledge, this work is the first to study the minimum required performance of the integrated front-end for the detection of seizures from the vagus nerve electroneurogram. It aims to define accurate front-end design specifications to minimize its power consumption. For that purpose, the presented work uses experimental data, an accurate model of the integrated front end, and a custom seizure detection algorithm. The major analog design choices and non-idealities were considered one by one in a sensitivity analysis, showing that noise and dynamic range are the major factors impacting seizure detection performance, and that accurate seizure detection is possible even in case of poor AP identification. Potentially optimal front-end configurations were then proposed. Compared to the prototyped front-end, the ADC resolution and oversampling rate can be relaxed to 6 bits and 1 respectively, without sacrificing detection performance. Similarly, the IA can tolerate a noise up to 3 μV_{RMS} and a CMRR of 60 dB. A transistor-level implementation of an optimal front-end meeting those specifications was proposed, achieving a power consumption of only 2.1 µW. The quantitative results presented are dependent on the assumptions used in this work. For instance, the noise results are valid for the signal amplitude obtained in the experimental recordings, and the CMRR values are associated with an assumed 1-mV common-mode input. The methodology employed in this work is portable to other applications. It can be used to validate the performance of an acquisition circuit for specific signal-processing tasks, as well as to define specifications for acquisition circuits before the transistor-level design. Further works should focus on obtaining chronic recordings in human subjects to address the differences between rodent and human devices, and to better tailor the front-end performance to the final application. Larger volumes of data could also help optimize the algorithm and integrate more advanced machine learning methods for the classification. The proposed optimal design could also be prototyped to validate its performance.

ACKNOWLEDGMENT

The authors would like to thank the members of the ECS group at UCLouvain and of the BEAMS department at ULB for their valuable feedback on this work.

REFERENCES

[1] M. Ravan, S. Sabesan, and O. D'Cruz, "On quantitative biomarkers of VNS therapy using EEG and ECG signals," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 2, pp. 419–428, 2017.
[2] H. Kassiri *et al.*, "Closed-loop implantable neurostimulators for individualized treatment of intractable epilepsy: A review of recent developments, ongoing challenges, and future opportunities," *IEEE Trans. Biomed. Circuits Syst.*, vol. 18, no. 6, pp. 1268–1295, 2024.
[3] P. Ryvlin *et al.*, "Neuromodulation in epilepsy: state-of-the-art approved therapies," *The Lancet Neurology*, vol. 20, no. 12, pp. 1038–1047, 2021.

[4] P. Afra *et al.*, "Evolution of the vagus nerve stimulation (VNS) therapy system technology for drug-resistant epilepsy," *Frontiers Med. Technol.*, vol. 3, 2021.

[5] B. Fisher *et al.*, "Responsive vagus nerve stimulation for drug resistant epilepsy: A review of new features and practical guidance for advanced practice providers," *Frontiers Neurology*, vol. 11, 2021.

[6] P. Hamilton *et al.*, "Clinical outcomes of VNS therapy with AspireSR® (including cardiac-based seizure detection) at a large complex epilepsy and surgery centre," *Seizure*, vol. 58, p. 120–126, May 2018.

[7] P. Boon *et al.*, "A prospective, multicenter study of cardiac-based seizure detection to activate vagus nerve stimulation," *Seizure*, vol. 32, pp. 52–61, 2015.

[8] D. Wang *et al.*, "Epileptic seizure detection in long-term EEG recordings by using wavelet-based directed transfer function," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 11, pp. 2591–2599, 2018.

[9] A. Bhattacharyya and R. B. Pachori, "A multivariate approach for patient-specific EEG seizure detection using empirical wavelet transform," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 9, pp. 2003–2015, 2017.

[10] S. Zhao, J. Yang, and M. Sawan, "Energy-efficient neural network for epileptic seizure prediction," *IEEE Trans. Biomed. Eng.*, vol. 69, no. 1, pp. 401–411, 2022.

[11] Y. S. Park *et al.*, "Early detection of human epileptic seizures based on intracortical microelectrode array signals," *IEEE Trans. Biomed. Eng.*, vol. 67, no. 3, pp. 817–831, 2020.

[12] P. Detti *et al.*, "A patient-specific approach for short-term epileptic seizures prediction through the analysis of EEG synchronization," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 6, pp. 1494–1504, 2019.

[13] A. Ulate-Campos *et al.*, "Automated seizure detection systems and their effectiveness for each type of seizure," *Seizure*, vol. 40, p. 88–101, 2016.

[14] L. Luan *et al.*, "Emerging penetrating neural electrodes: In pursuit of large scale and longevity," *Annu. Rev. Biomed. Eng.*, vol. 25, no. 1, p. 185–205, Jun. 2023.

[15] E. Acedo Reina *et al.*, "Vagus nerve electroneurogram-based detection of acute kainic acid induced seizures," *Frontiers Neuroscience*, vol. 18, Aug. 2024.

[16] K. R. Harreby, C. Sevcencu, and J. J. Struijk, "Early seizure detection in rats based on vagus nerve activity," *Med. Biol. Eng. Comput.*, vol. 49, no. 2, p. 143–151, Oct. 2010.

[17] L. Stumpp *et al.*, "Vagus nerve electroneurogram-based detection of acute pentylenetetrazol induced seizures in rats," *Int. J. Neural Syst.*, vol. 31, no. 07, p. 2150024, May 2021.

[18] P. Doguet *et al.*, "An optoelectronic implantable neurostimulation platform allowing full MRI safety and optical sensing and communication," *Scientific Rep.*, vol. 14, no. 1, May 2024.

[19] S. Martínez *et al.*, "Trends in volumetric-energy efficiency of implantable neurostimulators: A review from a circuits and systems perspective," *IEEE Trans. Biomed. Circuits Syst.*, vol. 17, no. 1, pp. 2–20, 2023.

[20] *AspireSR™ Technical Sheet*, LivaNova USA, Inc, Houston, Texas, 2022.

[21] R. Dekimpe and D. Bol, "Mixed-signal compensation of tripolar cuff electrode imbalance in a low-noise ENG analog front-end," in *Proc. IEEE 48th Eur. Solid-State Circuits Conf. (ESSCIRC)*, 2022, pp. 445–448.

[22] S. Favresse, D. Bol, and D. Flandre, "A combined analytical and simulation-based methodology for quantifying the noise-power-area trade-offs in biomedical amplifiers," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 71, no. 12, pp. 6177–6189, 2024.

[23] S. Favresse *et al.*, "An ultra-low-power dual-channel analog front end for vagus nerve sensing with chopping-minimized gain mismatch," in *IEEE Biomed. Circuits Syst. Conf. (BioCAS)*, 2025, pp. 121–125.

[24] P. Reisinger *et al.*, "The impact of EEG preprocessing parameters on ultra-low-power seizure detection," *Epilepsia*, vol. 66, no. 10, pp. 3895–3906, 2025.

[25] Y. Chen *et al.*, "Optimizing neural recording front-ends toward enhanced spike sorting accuracy in high-channel-count systems," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 33, pp. 2180–2191, 2025.

[26] R. Dekimpe and D. Bol, "Cross-domain optimization of low-power mixed-signal sensor systems under classification accuracy constraints," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 43, no. 8, pp. 2507–2517, 2024.

[27] H. S. Liu, T. Zhang, and F. S. Yang, "A multistage, multimethod approach for automatic detection and classification of epileptiform EEG," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 12, p. 1557–1566, Dec. 2002.

[28] O. Kocadagli and R. Langari, "Classification of EEG signals for epileptic seizures using hybrid artificial neural networks based wavelet transforms and fuzzy relations," *Expert Syst. Appl.*, vol. 88, pp. 419–434, 2017.

[29] A. A. Ein Shoka *et al.*, "EEG seizure detection: concepts, techniques, challenges, and future trends," *Multimedia Tools Appl.*, vol. 82, no. 27, p. 42021–42051, Apr. 2023.

[30] A. Sharmila, "Epilepsy detection from EEG signals: a review," *J. Med. Eng. & Technol.*, vol. 42, no. 5, pp. 368–380, 2018.

[31] R. Raffoul *et al.*, "Action potential detection algorithm adaptable to individual nerve and recording setup," in *Proc. IEEE Biomed. Circuits Syst. Conf. (BioCAS)*, 2022, pp. 655–659.

[32] M.-A. Ngan Yamb *et al.*, "Seizure detection based on vagus nerve activity: Translation to an implantable setting," in *Proc. IEEE Int. Symp. Med. Meas. Appl. (MeMeA)*, 2024, pp. 1–6.

[33] L. Stumpp *et al.*, "Recording of spontaneous vagus nerve activity during pentylenetetrazol-induced seizures in rats," *J. Neuroscience Methods*, vol. 343, p. 108832, 2020.

[34] S. Favresse, "VENG-HW-limitations-seizure-detection," 2026. [Online]. Available: https://github.com/SylvainFav/VENG-HW-limitations-seizure-detection/tree/main

[35] A. Demosthenous *et al.*, "An integrated amplifier with passive neutralization of myoelectric interference from neural recording tripoles," *IEEE Sensors J.*, vol. 13, no. 9, pp. 3236–3248, 2013.

[36] A. Nonclercq *et al.*, "Cluster-based spike detection algorithm adapts to interpatient and intrapatient variation in spike morphology," *J. Neuroscience Methods*, vol. 210, no. 2, p. 259–265, Sep. 2012.

[37] J. Dan *et al.*, "SzCORE: Seizure community open-source research evaluation framework for the validation of electroencephalography-based automated seizure detection algorithms," *Epilepsia*, vol. 66, no. S3, pp. 14–24, 2025.

[38] K. A. Ng *et al.*, "A wireless multi-channel peripheral nerve signal acquisition system-on-chip," *IEEE J. Solid-State Circuits*, vol. 54, no. 8, pp. 2266–2280, 2019.

[39] Y. Liu *et al.*, "A 64-channel versatile neural recording SoC with activity-dependent data throughput," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 6, pp. 1344–1355, 2017.

[40] S. Luan *et al.*, "Compact standalone platform for neural recording with real-time spike sorting and data logging," *J. Neural Eng.*, vol. 15, no. 4, p. 046014, May 2018.

[41] T. J. Levy *et al.*, "An impedance matching algorithm for common-mode interference removal in vagus nerve recordings," *J. Neuroscience Methods*, vol. 330, p. 108467, Jan. 2020.

[42] A. Demosthenous and I. Triantis, "An adaptive ENG amplifier for tripolar cuff electrodes," *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 412–421, 2005.

[43] A. Nonclercq and P. Mathys, "Reduction of power line interference using active electrodes and a driven-right-leg circuit in electroencephalographic recording with a minimum number of electrodes," in *26th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, vol. 1, 2004, pp. 2247–2250.

[44] R. Rieger *et al.*, "Design of a low-noise preamplifier for nerve cuff electrode recording," *IEEE J. Solid-State Circuits*, vol. 38, no. 8, pp. 1373–1379, 2003.

[45] D. A. Hall, K. A. Makinwa, and T. Jang, "Quantifying biomedical amplifier efficiency: The noise efficiency factor," *IEEE Solid-State Circuits Mag.*, vol. 15, no. 2, pp. 28–33, 2023.

[46] S. F. Thrall *et al.*, "A comparison of wavelet-based action potential detection from the NeuroAmp and the Iowa Bioengineering Nerve Traffic Analysis system," *J. Neurophysiol.*, vol. 131, no. 6, pp. 1168–1174, 2024.

[47] A. Lüttjohann, P. F. Fabene, and G. van Luijtelaar, "A revised racine's scale for PTZ-induced seizures in rats," *Physiol. & Behav.*, vol. 98, no. 5, p. 579–586, Dec. 2009.

[48] T. Singh, A. Mishra, and R. K. Goel, "PTZ kindling model for epileptogenesis, refractory epilepsy, and associated comorbidities: relevance and reliability," *Metabolic Brain Disease*, vol. 36, no. 7, p. 1573–1590, Aug. 2021.

[49] N. A. Pelot *et al.*, "Quantified morphology of the cervical and subdiaphragmatic vagus nerves of human, pig, and rat," *Frontiers Neurosci.*, vol. 14, Nov. 2020.

[50] H. Smets *et al.*, "Analysing vagus nerve spontaneous activity using finite element modelling," *J. Neural Eng.*, vol. 18, no. 5, p. 056008, apr 2021.