

Received July 17, 2020, accepted August 11, 2020, date of publication August 17, 2020, date of current version August 28, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3017394

Cyberattacks on Miniature Brain Implants to Disrupt Spontaneous Neural Signaling

SERGIO LÓPEZ BERNAL¹, ALBERTO HUERTAS CELDRÁN²,
LORENZO FERNÁNDEZ MAIMÓ³, MICHAEL TAYNNAN BARROS^{4,5}, (Member, IEEE),
SASITHARAN BALASUBRAMANIAM², (Senior Member, IEEE),
AND GREGORIO MARTÍNEZ PÉREZ¹, (Member, IEEE)

¹Departamento de Ingeniería de la Información y las Comunicaciones, University of Murcia, 30100 Murcia, Spain

²Telecommunication Software and Systems Group, Waterford Institute of Technology, X91 K0EK Waterford, Ireland

³Departamento de Ingeniería y Tecnología de Computadores, University of Murcia, 30100 Murcia, Spain

⁴School of Computer Science and Electronic Engineering, University of Essex, Essex CO4 3SQ, U.K.

⁵CBIG/BioMediTech, Faculty of Medicine and Health Technology, Tampere University, 33014 Tampere, Finland

Corresponding author: Sergio López Bernal (slopez@um.es)

This work was supported by the Irish Research Council, under the Government of Ireland Postdoctoral Fellowship under Grant GOIPD/2018/466.

ABSTRACT Brain-Computer Interfaces (BCI) arose as systems that merge computing systems with the human brain to facilitate recording, stimulation, and inhibition of neural activity. Over the years, the development of BCI technologies has shifted towards miniaturization of devices that can be seamlessly embedded into the brain and can target single neuron or small population sensing and control. We present a motivating example highlighting vulnerabilities of two promising micron-scale BCI technologies, demonstrating the lack of security and privacy principles in existing solutions. This situation opens the door to a novel family of cyberattacks, called neuronal cyberattacks, affecting neuronal signaling. This article defines the first two neural cyberattacks, Neuronal Flooding (FLO) and Neuronal Scanning (SCA), where each threat can affect the natural activity of neurons. This work implements these attacks in a neuronal simulator to determine their impact over the spontaneous neuronal behavior, defining three metrics: number of spikes, percentage of shifts, and dispersion of spikes. Several experiments demonstrate that both cyberattacks produce a reduction of spikes compared to spontaneous behavior, generating a rise in temporal shifts and a dispersion increase. Mainly, SCA presents a higher impact than FLO in the metrics focused on the number of spikes and dispersion, where FLO is slightly more damaging, considering the percentage of shifts. Nevertheless, the intrinsic behavior of each attack generates a differentiation on how they alter neuronal signaling. FLO is adequate to generate an immediate impact on the neuronal activity, whereas SCA presents higher effectiveness for damages to the neural signaling in the long-term.

INDEX TERMS Brain computer interfaces, security, artificial neural networks, biological neural networks.

I. INTRODUCTION

Brain-computer Interfaces (BCIs) are considered as bidirectional communication systems between the brain and external computational devices. Although BCIs arose as systems focused on controlling external devices such as prosthetic limbs [1], they have gone one step further, enabling *artificial* stimulation and inhibition of neuronal activity [2]. In the last years, neuronal stimulation has already been applied in

The associate editor coordinating the review of this manuscript and approving it for publication was Yassine Maleh¹.

different scenarios such as the provision of sensory feedback to prosthetic or robotic limbs [3], treatment of neurodegenerative diseases or disorders like Alzheimer's or depression [4], and even futuristic applications such as interconnected networks of brains [5] or brains connected to the Internet [6].

New BCI technologies are emerging, allowing a precise acquisition, stimulation, and inhibition of neuronal signaling. It reduces the brain damage caused by traditional invasive BCI systems and improves the limitations of non-invasive technologies such as attenuation, resolution, and distortion constraints [7], [8]. One of the most recent and promising

BCI technique focuses on the use of nanodevices allocated across the brain cortex [9]. Specifically, a relevant task of nanodevices equipped with optogenetic technology is the use of light to stimulate or inhibit engineered neurons according to different firing patterns sent by external transceivers [10]. Promising initiatives such as Neuralink aim to accelerate the development of these technologies [11].

The previous BCI technologies hold the promise of changing our society by improving the cognitive, sensory, and communications skills of their users. However, they also open the door to critical cyberattacks affecting the subjects' safety and data security. In this context, essential vulnerabilities of current non-invasive BCI systems have been documented, exploited, and partially solved in the literature [12]. As an example, the authors of [13], [14] demonstrated the feasibility of presenting malicious visual stimuli to extract subjects' sensitive data like thoughts. Besides, Sundararajan [15] conducted a successful jamming attack over the wireless communication used by the BCI, compromising its availability. However, the irruption of invasive and non-invasive stimulation and inhibition techniques, without security nor privacy capabilities, brings to the reality a novel family of cyberattacks affecting the neuronal activity. We call them *Neural cyberattacks*, and they present a critical number of open challenges like the definition and categorization of the different neural cyberattacks and their neuronal behavior, the impact of each cyberattack to the neuronal behavior, and their consequences in the brain and body.

To improve the previous challenges, the main contributions of this article are the following ones:

- The identification of cybersecurity vulnerabilities on emerging neurostimulation implants.
- To the best of our knowledge, the first description and implementation of neural cyberattacks focused on neuronal stimulation and affecting the activity of neural networks allocated in the human's brain. The proposed cyberattacks, *Neuronal Flooding* and *Neuronal Scanning*, are inspired by the behavior of current well-known cyberattacks in computer networks.
- The definition of three metrics to evaluate the impact of the two neural cyberattacks proposed: number of spikes, percentage of shifts, and dispersion of spikes.
- The implementation of the previous cyberattacks in a neuronal simulator to measure the impact produced by each one of them and the implications that they generate on the neuronal signaling. For that, we model a portion of a mouse's visual cortex based on the implementation of a CNN where the mouse is able to exit a maze.

The paper remainder is organized as follows. Section II gives an overview of the present state-of-the-art of current vulnerabilities, cyberattacks, and countermeasures affecting existing BCIs. After that, Section III illustrates emerging neurostimulation technologies and their cybersecurity concerns. Subsequently, Section IV offers a formal description of the cyberattacks proposed, while Section V describes the implemented use case. Section VI first presents the metrics

used to evaluate the impact of these cyberattacks, followed by the analysis of the results and impact that these cyberattacks generate. Finally, Section VII briefly discusses the outcomes and potential future works.

II. RELATED WORK

During the last five years, new concepts such as brain-hacking, or neurocrime have emerged to describe relevant aspects of cybersecurity in BCI [16], [17]. These works highlight that neuronal engineering devices, designed to stimulate targeted regions of the brain, would become a critical cybersecurity problem. In particular, they acknowledge that attackers may maliciously attempt to program the stimulation therapy, affecting the patient's safety. Furthermore, they emphasize that the cyberthreats do not need to be too sophisticated if they only want to cause harm. In this context, as indicated in this article, it is possible to have a high impact on the brain by taking advantage of neurostimulation implants and send malicious electrical signals to the brain. Despite the identification of these risks, there are no studies in the literature defining or implementing neural cyberattacks, where the evaluation of their impact over the brain remains unexplored. However, several vulnerabilities and attacks have been detected in BCI technologies performing neural data acquisition (e.g., EEG), which can serve as a starting point to perform neural cyberattacks. Section III offers additional considerations about vulnerabilities in BCI solutions.

Platforms and frameworks that enable the development of BCI applications also present cybersecurity concerns, as demonstrated in [18], [19]. In this context, the authors of [18] performed an analysis of the privacy concerns of BCI application stores, including Software Development Kits (SDKs), Application Programming Interfaces (APIs), and BCI applications. They discovered that most applications have unrestricted access to subjects' brainwave signals and can easily extract private information about their subjects. Moreover, Cody's Emokit project [17], managed to break the encryption of the Emotiv EPOC device (valid for all models before 2016), having access to all raw data transmitted. The authors of [19] proposed a mechanism to prevent side-channel extraction of subjects' private data, based on the anonymization of neural signals before their storage and transmission.

The majority of the existing BCI systems are oriented to acquire, or record, neural data. Specifically, EEG BCI devices have gained popularity in recent years, due to their low cost and versatility, influencing the number of existing cyberattacks exploiting BCI vulnerabilities. In this context, the authors of [20] studied and analyzed well-known BCI applications and their potential cybersecurity and privacy concerns. Martinovic *et al.* [14] were able to extract users' sensitive information, such as debit cards or PINs, by presenting particular visual stimuli to the users and analyzing their P300 potential response. Another attack, performed by Frank *et al.* [13], focused on presenting subliminal visual stimuli included within a video, aiming to affect the BCI

users' privacy. Finally, in our previous work [21], we studied the feasibility of performing cybersecurity attacks against the stages of the BCI cycle, considering different communication architectures, and highlighting their impact and possible countermeasures.

In conclusion, this section demonstrates that most of the related works are focused on presenting vulnerabilities and cyberattacks affecting the confidentiality, availability, and integrity of private data managed by BCIs. Nevertheless, there is a lack of solutions considering cyberattacks affecting the neuronal activity and, therefore, the subjects' safety. This article proposes two neural cyberattacks affecting the natural behavior of single and population of neurons.

III. CYBERSECURITY VULNERABILITIES OF EMERGING NEUROSTIMULATION IMPLANTS

This section introduces three promising BCI technologies capable of recording and stimulating neuronal activity with single-neuron resolution. For each scenario, we offer a description of its architecture, highlighting the cybersecurity vulnerabilities detected. Although these solutions are in an early stage, and they are still not commercial products, they are contemporary examples of how cybersecurity can affect existing and future implantable BCI solutions, and in particular for solutions that can target small neuron populations. These issues represent the starting point for the cyberattacks illustrated in the next sections of this article. It is important to note that the objective of this section is not to find vulnerabilities in BCI devices or architectures but to justify how the proposed cyberattacks could be performed in realistic BCI systems.

A. NEURALINK

Neuralink aims to record and stimulate the brain using new technologies, materials, and procedures to reduce the impact of implanting electrodes in the brain [11]. The first element of the Neuralink architecture are the *threads*, proposed as an alternative for traditional electrodes due to their biocompatibility, reduced size based on thin threads that are woven into the brain tissue, durability, and the number of electrodes per thread. Groups of threads connect to an *N1 sensor*, a sealed device in charge of receiving the neural recordings from the threads and sending them stimulation impulses. With a simple medical procedure, up to ten N1 implants can be placed in the brain cortex. These devices connect, using tiny wires tunneled under the scalp, to a *coil* implanted under the ear. The coil communicates wirelessly through the skin with a wearable device, or *link*, placed under the ear. The *link* contains a battery that represents the only power source in the architecture, deactivated if the user removes the link. FIGURE 1 represents this architecture.

Although the communication mechanisms between the coil and the link are not provided, the link is managed via Bluetooth from external devices, such as smartphones, using an application. In this sense, Neuralink users can manage and personalize their links, upgrade their firmware, and include

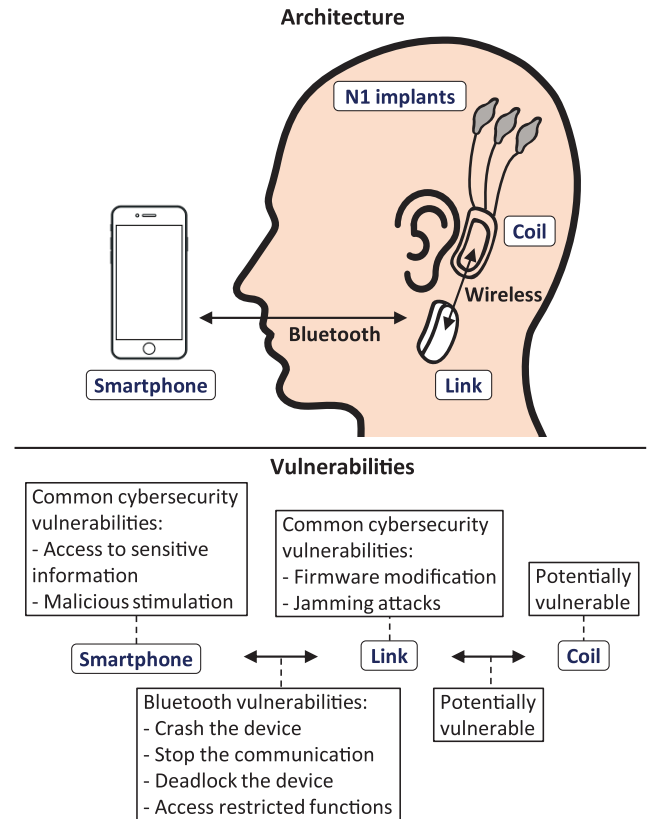


FIGURE 1. Architecture and vulnerabilities of Neuralink.

new security capabilities. We identify that this scenario can be potentially vulnerable as follows. First, the wireless mechanism used in the communication between the coil and the link could be vulnerable, depending on the protocol used [22]. Besides, the Bluetooth communication between the smartphone and the link can also be vulnerable, according to the version used [23], [24]. As an example, we identify Sweyn-Tooth, a set of 12 vulnerabilities affecting a large number of devices using Bluetooth Low Energy (BLE) technologies. Based on them, an attacker could crash the device and stop its communications [25], deadlock the device [26], or access functions only available for authorized users [27].

Moreover, the external device manages the logic of both acquisition and stimulation processes, including into these scenarios its inherent risks, and becoming one of the most sensitive elements of the architecture. In particular, Li et al. [20] detected that attackers could take total control of a smartphone running a BCI application, getting access to sensitive information, or performing malicious stimulation actions. Furthermore, the *link* is a critical element of the architecture, where attackers can modify the firmware of the device to have a malicious behavior, as identified by [28] for brain implants or to perform jamming attacks to disrupt the communication between devices, described by [29] for wireless networks.

B. NEURAL DUST

This architecture is composed of millions of resource-constrained nanoscale implantable devices, also known as

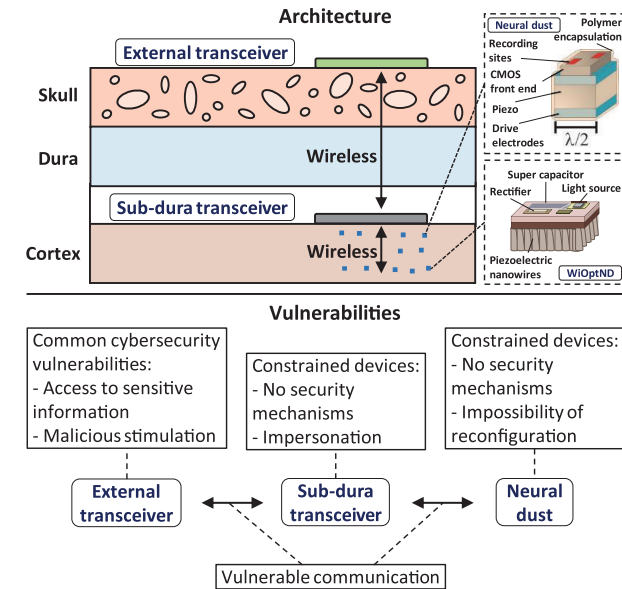


FIGURE 2. Architecture and vulnerabilities of Neural dust.

neural dust, floating in the cortex, able to monitor neural electrophysiological activity precisely [9]. These devices communicate with the sub-dura transceiver, a miniature device (constructed from components that are built from nanomaterials) placed beneath the skull and below the dura mater. This device uses two different transceivers to: (1) power and establish communication links with the neural dust, (2) communicate with external devices. During neural recording, the sub-dura transceiver performs both spatial and frequency discrimination with sufficient bandwidth to power and interrogate each neural dust. The external transceiver is a device without computational and storage restrictions, allocated outside of the patient's head. Wearables, smartphones, or PCs are examples of this device. The main task of the external transceiver is to power and communicate with the sub-dura transceiver and to receive the neuronal behavior from the sensing by the neural dust. FIGURE 2 presents the architecture of this solution, as well as the potential vulnerabilities that it presents.

Nevertheless, this technology has not been conceived following the principle of security and privacy by design. As a consequence, these devices do not implement authentication mechanisms to prevent malicious users from collecting neural sensing data from the neural dust, and they do not protect the transmitted data. In particular, the neural dust are resource-constrained devices without computational and storage capabilities to execute security functionalities like authentication protocols, ciphered communications, or data encryption. In this sense, external attackers could power and communicate to the implants to monitor private neural data. Finally, the sub-dura and external transceivers do not implement authentication protocols nor security mechanisms. An attacker could impersonate the external transceiver to communicate with the sub-dura device, and obtain sensitive neuronal signaling.

C. WIRELESS OPTOGENETIC NANONETWORKS

The Wireless Optogenetic Nanonetworking device (WiOptND) [10] is an extension from the neural dust [9] but with the capability of optogenetically stimulating the neurons. Optogenetic stimulation uses light to stimulate neurons genetically engineered with specific genes that are sensitive to signals at a particular wavelength. This in turn provides targeted stimulation of very small population of neurons that have been engineered, enabling precise targeting of neural circuits within the micro-columns. Similar to the architecture of the neural dust, the WiOptND also receives power that is emitted from the sub-dura, which in turn communicates to the external transceiver. However, since the WiOptND is responsible for stimulating the neurons, the external transceiver will communicate the sequence of firing the neurons to the sub-dura transceiver to synchronize the charging and communication of the WiOptND implants. This is achieved by sending the firing sequence, in the form of a raster plot, to the external transceiver. This opens up new opportunities for attackers to send malicious firing patterns into the external transceiver, which will produce a new sequence of firing patterns for neural stimulation, resulting in detrimental consequences for the brain. Finally, the architecture and vulnerabilities described in FIGURE 2 also apply for WiOptND.

In conclusion, the previous vulnerabilities raise different concerns affecting the integrity, confidentiality and availability of subject's neural data. These vulnerabilities motivate different attack vectors to perform the neural cyberattacks described in subsequent sections.

IV. DEFINITION OF NEURAL CYBERATTACKS

Once demonstrated the feasibility of stimulating individual neurons by attacking different technological solutions, we formally describe two cyberattacks, *Neuronal Scanning* and *Neuronal Flooding*, aiming to maliciously affect the natural activity of neurons during neurostimulation procedures. They are inspired by the behavior and goals of some of the most well-known and dangerous cyberattacks affecting computer networks.

To formalize both cyberattacks, we denote $\mathbb{NE} \subset \mathbb{N}$ as a subset of neurons from the brain, where $n \in \mathbb{NE}$ expresses every single neuron. The voltage of a single neuron in a specific instant of time is denoted as $v_n \in \mathbb{R}$, whereas $v_n \in \mathbb{R}$ indicates the voltage increase used to overstimulate a neuron n . Moreover, t^{win} represents a temporal window in which the cyberattack is performed, equivalent to the duration of the simulation in Section VI. t^{attk} is the time instant when the cyberattack starts, and Δt the amount of time between evaluations during the process. In the implementation of the cyberattacks, it represents the duration of the steps of the simulation.

1) NEURONAL FLOODING

In the cyberworld, a flooding cyberattack is designed to bring a network or service down by collapsing it with large

amounts of network traffic. Traffic is usually generated by many attackers and forwarded to one or more victims. Extrapolating this network cyberthreat to the brain, a Neuronal Flooding (FLO) cyberattack consists in stimulating multiple neurons in a particular instant of time, changing the normal behavior of the stimulation process and generating an overstimulation impact. The execution of this cyberattack does not require prior knowledge of the status of the affected neurons since the attacker only has to decide what neurons to stimulate and when. This fact makes this cyberattack less complex than other cyberattacks that require prior knowledge of the neuronal behavior.

In particular, FLO performs the overstimulation action at t^{attk} . In that precise moment, a subset of neurons $\mathbb{AN} \subseteq \mathbb{NE}$ is attacked. This cyberattack is formally described in Algorithm 1.

Algorithm 1 FLO Cyberattack Execution

```

t = 0
while t < twin do
  if t == tattk then
    for all n ∈ AN do
      vn ← vn + vi
    end for
  end if
  t ← t + Δt
end while
  
```

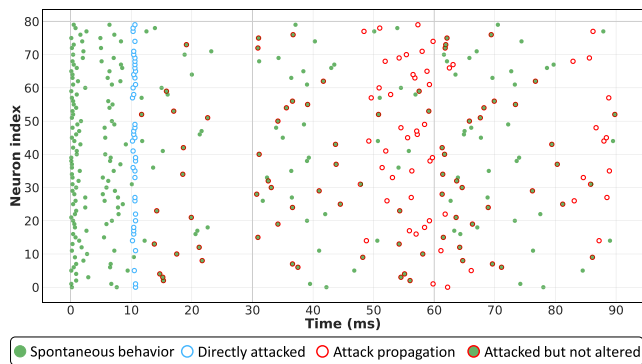


FIGURE 3. Raster plot of a FLO cyberattack when the attack is performed at 10ms.

FIGURE 3 represents an example to appreciate graphically the behavior of a FLO cyberattack, where the details of the neuronal network used in the simulation are not relevant at this point (addressed in Section V). In particular, it represents the comparison of the FLO cyberattack with the spontaneous behavior for a simulation of 80 neurons, a duration of 90ms, and 42 neurons attacked in the instant 10ms. Green dots represent the neuronal spontaneous behavior, blue circles indicate the instant when the neurons are attacked, red circles highlight the propagation of the cyberattack in time, and those dots with a green color and red outline represent spikes common to both spontaneous and under attack situations. In this figure, we can see that all the attacked neurons alter

their behavior, having spikes in different moments compared to the spontaneous activity.

2) NEURONAL SCANNING

Port scanning is another well-known cybersecurity technique performed by attackers to discover vulnerabilities in operating systems, programs, and protocols using network communications. In particular, it aims to test every networking port of a machine, checking if it is open and discovering the protocol or service available in that end-point. In the brain context, a Neuronal Scanning (SCA) cyberattack stimulates neurons sequentially, impacting only one neuron per instant of time. Based on that, it is essential to note that attackers do not require prior knowledge of the neuronal state to perform neuronal scanning cyberattacks. This fact, together with the stimulation of one neuron per instant of time, makes a low attack complexity.

Considering the notation previously defined, Algorithm 2 describes an SCA cyberattack. In particular, it sequentially overstimulates all the neurons included in the set of neurons \mathbb{NE} , without repetitions. For each neuron n , its voltage v_n increases by v_{i_n} . It is essential to indicate that the conditional clause limits the instants in which an attack can be performed, where t^{attk} represents the attack over the first neuron of the set, and $t^{attk} + |\mathbb{NE}|\Delta t$ the attack over the last neuron.

Algorithm 2 SCA Cyberattack Execution

```

t = 0
while t < twin do
  if t ∈ [tattk, tattk + |NE|Δt] then
    n ← (t - tattk) / Δt
    vn ← vn + vi
  end if
  t ← t + Δt
end while
  
```

Finally, FIGURE 4 shows, in a visual way, the behavior of an SCA cyberattack. We simulate 80 neurons during 90s, and sequentially attack all neurons, starting in the instant 10ms. The color code followed is the same as in FIGURE 3. As can be seen, the sequential attack of the neurons generates a diagonal line in the spikes. All spikes over the line remain unaltered since those neurons have not yet been affected by the attack. On the contrary, the spikes under the diagonal are affected by the attack.

V. EXPLOITING VULNERABILITIES DUE TO CYBERATTACKS

This section introduces the use case used to implement the cyberattacks defined in Section IV. We present the scenario and the experimental setup implemented to create the neuronal topology required to test the cyberattacks.

A. USE CASE AND EXPERIMENTAL SETUP

The knowledge of precise neocortical synaptic connections in mammalian is nowadays an open challenge [30]. Based on

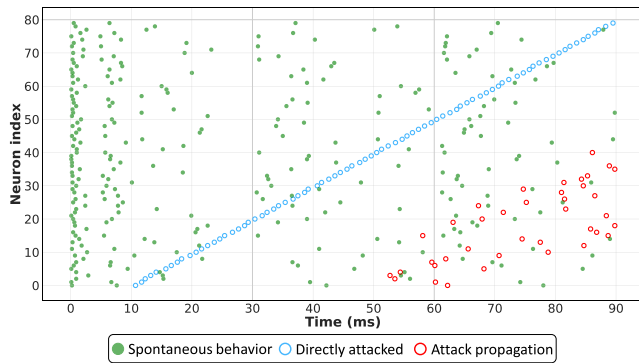


FIGURE 4. Raster plot of an SCA cyberattack, from the instant 10ms to 90ms.

this absence of realistic neuronal topologies, we have studied the primary visual cortex of mice and replicated a portion of it, modeled using a Convolutional Neural Network (CNN) [31]. This CNN was trained by means of reinforcement learning [32] to represent a simple system able to make decisions based on a maze and find its exit. As indicated by Kuzovkin *et al.* [33], CNNs, and biological neuronal networks present certain similarities. First, lower layers of a CNN explain gamma-band signals from earlier visual areas, whereas higher layers explain later visual regions. Furthermore, early visual areas are mapped to convolutional layers, where the fully connected layers match the activity of higher visual areas. That is to say, the visual recognition process in both networks is incremental and move from simple to abstract. At this point, it is essential to note that we cannot compare the topology and functionality of a CNN to the complexity of the neuronal connections of a real brain. We only used this technique to provide a simple topology that is then implemented in a neuronal simulator to evaluate how attacks over a simplistic but realistic environment can affect the activity of simulated neurons, as indicated in Section V-C.

In this context, we designed a simple proof of concept based on the idea of a mouse that has to solve the problem of finding the exit of a particular maze, inspired in the code from [34]. The mouse must find the exit with the smallest number of movements and starting from any position. We define a maze of 7×7 coordinates, as represented in FIGURE 5. It contains one starting position identified with “1”, while the exit is labeled with “27”. Moreover, the positions colored in gray represent obstacles, and those in white are accessible positions through which the mouse can move. In this scenario, the mouse can move in all four 2D directions: up, down, left, and right. The numbering from 1 to 27 defines the optimal path determined by the trained CNN to reach the exit position, considering the lowest number of steps. Finally, it is essential to define the concept of *visible position*. From each particular cell of the maze, the mouse can visualize a square of 3×3 adjacent positions, including those that represent obstacles. This situation is highlighted in FIGURE 5 with a red square, indicating the visible positions from the cell 15 of the optimal path.




FIGURE 5. Maze used in our use case to model the movement of the mouse, including the optimal path between the starting and final cells. There are nine visible positions from the cell 15, highlighted within a red square.

TABLE 1. Summary of the layers of the CNN.

Layer	Type	Filters	Input size	Output size	Kernel size	Stride	Activation function	Nodes
1	Conv2D	8	$7 \times 7 \times 1$	$5 \times 5 \times 8$	3×3	1	ReLU	200
2	Conv2D	8	$5 \times 5 \times 8$	$3 \times 3 \times 8$	3×3	1	ReLU	72
3	Dense	-	$3 \times 3 \times 8$	4	-	-	ReLU	4

B. CONVOLUTIONAL NEURAL NETWORK

Our objective was to generate a CNN able to exit the maze from any position. We also aimed to define a topology with a reduced number of nodes to be compatible with resource-constrained neuronal simulators since we aim to evaluate this topology in multiple simulators. Nevertheless, for simplicity, this work includes details of the implementation in only one simulator, as described in Section V-C. To solve our maze problem, we implemented a CNN composed of two convolution layers and a dense layer. The ensemble of these three layers defines a complete CNN of 276 neurons, representing a small portion of a mouse primary visual cortex, summarized in Table 1. We implemented this CNN using Keras on top of TensorFlow [35].

FIGURE 6 depicts the architecture of the implemented CNN which is also described in Table 1. In particular, we have included a first 2D convolution layer with a 3×3 kernel. This layer takes as input the current status of the maze, focusing each neuron on a square of 9 (3×3) adjacent positions. In our experiments we determined that 8 filters of size 3×3 in each layer were sufficiently expressive. To represent the maze, each position contains a 1 value if the position is accessible, a 0 value if it is an obstacle, or a 0.5 value in the position of the mouse.

During the training, each filter of the first layer specializes on a particular aspect of the maze. For example, a filter could focus on detecting vertical walls, while another could detect corners. The filters of the second layer can detect more complex scenarios by composing the output of these initial detectors. Since the input is a 7×7 maze, and the kernel is 3×3 , the first convolution process requires 25 neurons (5×5 kernel outputs) to cover the new 5×5 subset of the maze on the next layer. Since we use 8 different filters, the total number of neurons required to produce the first layer’s output of the CNN is 200 ($5 \times 5 \times 8$). This is illustrated in FIGURE 6, where each group of neurons has a different color




FIGURE 6. Visual representation of the implemented CNN. It introduces a simplifications of the whole topology, indicating how the convolution process is performed and how nodes connect between layers. The color of each node matches the color of its associated filter.

that matches the color of its filter. Therefore, since the first layer generates an output of size $5 \times 5 \times 8$, the application of the 3×3 kernels of the second convolutional layer requires a total of 72 ($3 \times 3 \times 8$) neurons. Finally, this new output is sent through a last dense layer of 4 neurons, one for each possible movement direction on the maze (left, up, right, down). Each output is an estimation of the probability of success with each movement, being selected the direction with the greatest score.

In order to understand Section V-C and Section VI, it is necessary to explain the mapping between the sequential number of each neuron and its position in its associated filter output. FIGURE 6 shows this mapping. Each neuron have associated a 3-dimensional vector, where the third coordinate is its filter and the two first coordinates, the position in that filter output. The order is as follows: the first neuron has the coordinates $[0,0,0]$, corresponding to the first neuron in the first filter output; the eighth neuron corresponds to $[0,0,7]$; the ninth one is $[0,1,0]$, and so on until the 200th neuron, with coordinates $[4,4,7]$.

C. BIOLOGICAL NEURONAL SIMULATION

After training the CNN, we represented its resulting topology in Brian2, a lightweight neuronal simulator [36]. We selected

TABLE 2. Relationship of parameters between artificial and biological networks.

	CNN	Simulation
Number of neurons		276
Number of layers		3
Neuronal topology	200 (Layer 1), 72 (Layer 2), 4 (Layer 3)	
Input data	Maze	
Types of neurons	Artificial	Pyramidal neuron from primary visual cortex
Connection weights	Filter weights	Synaptic weights

Brian2 because it is adequate to run neuronal models in user-grade computers, without the requirement of using multiple machines, or even supercomputers. It also presents a good behavior in the implementation of neuron models with simplified and discontinuous dynamics (such as Leaky Integrate-and-Fire or Izhikevich) [37]. Other alternatives, such as NEURON, present complex solutions to model neurons with fine granularity, offering distributed computation capabilities for high demanding simulations. Nevertheless, this functionality is unnecessary in our particular study.

We maintain in the biological simulation the exact number of layers, the number of neurons per layer, and the topological connections between neurons. However, there is a crucial difference between the implementation of these two approaches. In the CNN, a filter weight represents the importance that a connection between two neurons of different layers have on the topology and, thus, over the solution. In the biological simulation, we transform the CNN weights to synaptic weights, representing the increase of the voltage induced during an action potential. Table 2 summarizes these similarities and differences between both networks.

To represent the behavior of each neuron, we decided to use the Izhikevich neuronal model since it is computationally inexpensive, and it allows us to precisely model different types of neurons within different regions of the brain [38]. This model represents an abstraction of how cortical neurons behave in the brain. In particular, the following set of equations describes the Izhikevich model, whose parameters are indicated in Table 3. This model allows multiple configurations to mimic different regions of the brain. In our scenario, we assigned particular values to the previous parameters to implement a regular spiking signaling from the cerebral cortex, as indicated in [38]. Specifically, we aimed to model pyramidal neurons from the primary visual cortex of a mouse, which correspond to excitatory neurons typically present in the biological visual layers L2/3, L5, and L6 [39]. For simplicity, during the analysis of the results of the simulation, we will refer to these layers in subsequent sections as first layer (L2/3), second layer (L5) and third layer (L6).

$$v' = 0.04v^2 + 5v + 140 + u + I \tag{1}$$

$$u' = a(bv - u) \tag{2}$$


$$\text{if } v \geq 30mV, \text{ then } \begin{cases} v, & \leftarrow c \\ u, & \leftarrow u + d \end{cases} \tag{3}$$

TABLE 3. Parameters used in the Izhikevich model.


Parameter	Description	Values
v	Membrane potential of a neuron	$[-65, 30]$ mV
u	Membrane recovery variable providing negative feedback to v	$(-16, 2)$ mV/ms
a	Time scale of u	0.02/ms
b	Sensitivity of u to the sub-threshold fluctuations of v	0.2/ms
c	After-spike reset value of v	-65mV
d	After-spike reset value of u	8mV/ms
I	Injected synaptic currents	$\{10, 15\}$ mV/ms

To create our neuronal topology, we used the weights of the trained CNN as post-synaptic voltage values, normalized within the range between $5mV$ and $10mV$. We selected this range because these values constitute a conservative voltage raise within the range of values of v , indicated in Table 3. At the beginning of the simulation, we assigned the initial voltage of each neuron from a previously generated random list in the range $[-65mV, 0mV]$. This initial value for each neuron is constant between executions to allow their comparison. To define a more realistic use case, we represented in our simulation the movement of the mouse inside the maze (see FIGURE 5), staying one second in each position of the optimal path. To understand this, it is essential to introduce the concept of *intervening neurons*, which defines the set of neurons managing all the visible positions of the mouse when it is placed in a particular position of the maze. FIGURE 7a illustrates the relationship between the position 13 of the optimal path and its intervening neurons, not considering its related visible positions for simplicity. For this position, we define nine 3×3 squares within the surface delimited by the red square, where we represent only the first two squares to improve the legibility of the figure. Focusing on the first square, colored in blue, it comprises eight neurons indexes (49 to 56), obtained from the translation between 3-dimension coordinates previously commented in this section. The second one, highlighted in orange, associates eight different neurons. After applying all nine squares, we obtain the complete list of intervening neurons related to the position 13. This single process is repeated for every visible position from the position 13 (indicated in FIGURE 7a with red dots), obtaining the complete set of intervening neurons. This set of intervening neurons is presented in Table 4, where each visible position from the position 13 is identified by its maze coordinate for simplicity. The last row of the table presents the complete set of intervening neurons for the position 13, obtained as the union of all individual sets of neurons.

The movement of the mouse was implemented by providing external stimuli to the simulation via the I parameter, where a value of $15mV$ was assigned to all intervening neurons from the current location of the mouse. For all non-intervening neurons in a specific instant, we assigned a value of $10mV$. These values align with the range defined in [38]. This information was extracted from the topology of the CNN, which contains the relationship between the neurons of the first layer and the positions of the maze. We took into consideration these aspects in the experimental analysis performed in Section VI. Based on that, we modeled with a



(a) Relationship between visible positions from the current location of the mouse and their intervening neurons. In this example, the mouse is placed in position 13. Red dots indicate the visible positions from position 13.



(b) Impact of the I parameter on both intervening ($15mV$) and non intervening ($10mV$) neurons.

FIGURE 7. Relationship between positions of the maze and its implication in the modulation of neuronal signaling.

TABLE 4. List of intervening neurons associated to the position 13 of the optimal path of the maze.

Coordinate	List of intervening neurons
(2,2)	[1,24], [41,64], [81,104]
(2,3)	[9,32], [49, 72], [89, 112]
(2,4)	[17,40], [57,80], [97, 120]
(3,2)	[41,64], [81, 104], [121,144]
(3,3)	[49,72], [89,112], [129,152]
(3,4)	[57,80], [97,120], [137,160]
(4,2)	[81,104], [121,144], [161,184]
(4,3)	[89,112], [129,152], [169,192]
(4,4)	[97,120], [137,160], [177,200]
Position 13	[1,200]

higher value of I those intervening neurons, transmitting a more potent visual stimulus to those neurons related to adjacent positions from the current location. Based on Equation 1, an increase in the I parameter will produce a voltage rise in these intervening neurons, generating a raise in the amplitude of the electrical signal. This behavior was modeled taking into consideration the study performed in [40], which indicates that a known visual stimulus generates a voltage amplitude increase. FIGURE 7b graphically compares these differences between values of the I parameter. It highlights that intervening neurons present a higher number of spikes during a particular temporal window, which is interpreted by the brain as the reconnaissance of accessible cells in the maze from the current position.




FIGURE 8. Summary of our use case, indicating the translation between the topology of the CNN implemented and the biological network simulated.

Finally, FIGURE 8 introduces a graphical summary of the current use case. It depicts a mouse with a miniature brain implant solution in its primary visual cortex, such as Neuralink or Neural dust. To simulate its biological neuronal network, and based on a lack of realistic cortical topologies, a trained CNN provides the number of nodes and distribution in layers for the biological network. In particular, we modeled pyramidal neurons from visual layers L2/3, L5, and L6, using the Izhikevich model with a regular spiking signaling. Based on this scenario, an external attacker takes advantage of contemporary vulnerabilities in these implantable solutions to alter the behavior of the spontaneous activity of the biological neuronal network.

VI. RESULTS ANALYSIS BASED ON METRICS

In this section, we evaluate the impact that FLO and SCA cyberattacks have on spontaneous neuronal activity of the neuronal topology presented in Section V. To analyze the evolution of the cyberattacks impact while the mouse is moving across the maze, we consider the following three metrics:

- Number of spikes: determine if a cyberattack either increases or reduces the quantity of spikes compared to the spontaneous neuronal signaling.
- Percentage of shifts, being a shift the delay of a spike in time (forward or backward) compared to the spontaneous behavior: study if a cyberattack generates significant delays in the normal activity of the neurons.
- Dispersion of spikes in both dimensions of time and number of spikes: analyze the spiking patterns under attack, aiming to detect if the cyberattack causes a modification on the distribution of the spikes.

For each layer of the topology, and combining all of them, we measured and analyzed the number of spikes and

percentage of shifts. Finally, the dispersion of spikes is computed for each position of the optimal path and grouping all layers. Finally, we compared the impact generated by both cyberattacks.

To better understand the impact of FLO and SCA cyberattacks, FIGURE 9 compares the evolution of neuronal spikes for the spontaneous activity, a FLO cyberattack and an SCA cyberattack. We selected three positions of the optimal path to analyze in detail the spiking evolution along with the simulation, presenting only the first 100ms of each position. It is essential to note that this simplification is only for this figure, and all the results subsequently presented consider the complete duration of each position. As can be seen, in the spontaneous signaling, there is a certain natural dispersion caused by the behavior of the neuronal model used, and the movement of the mouse (due to the the modification of the associated I parameter). Specifically, each time the mouse changes from one position to another, the I parameter changes according to the intervening neurons, where a higher value of I is translated to a higher spike rate (see Algorithm 1). Since the mouse periodically changes its position, it modifies the spiking rate of the neurons, generating a natural dispersion in the absence of attacks. Looking at the first position of both spontaneous and FLO, in the instant 50ms, there is a clear difference between them, since we executed the attack in that exact instant. The set of attacked neurons generates spikes before it was intended due to the voltage rise produced by the attack. Consequently, we can see that the dispersion over the following positions (13 and 27) augments, altering the natural pattern of the neurons. Regarding the SCA cyberattack, it also starts in the instant 50ms but, its impact it is not yet present in the first 100ms of the initial position. If we check the subsequent positions, the attack gradually propagates,




FIGURE 9. Raster plots indicating the evolution of the spontaneous signaling and both FLO and SCA cyberattacks for three positions of the optimal path of the maze.

generating characteristic ascending patterns. Subsequent subsections analyze, in a more detailed way, the information contained in FIGURE 9, extending the analysis to all the positions of the optimal path and using the previous three metrics.

A. NEURONAL FLOODING

In this subsection, we aim to simultaneously attack multiple neurons and analyze its impact using the metrics previously indicated at the beginning of the section. The implementation of this cyberattack is based on the general description indicated in Algorithm 1. We decided to perform only the attacks over the first layer of the topology, from where each target neuron is randomly selected, to evaluate the propagation to deeper layers. Furthermore, we tested a combination of two additional parameters. The first one represents the number of simultaneously attacked neurons, $k \in \{5, 15, \dots, 95, 105\}$. \mathbb{AN} will contain k neurons randomly selected from \mathbb{NE} , the set of neurons in the first layer. It is worthy to note that we reached to attack simultaneously more than half of the neurons of the first layer, which represents a fairly aggressive portion of the neurons. The second parameter of the attack, $\mathbb{VI} = \{20, 40, 60\}$, indicates the different voltage increases in mV used to stimulate the neurons in \mathbb{AN} . Its maximum level, $60mV$, approximately represents two-thirds of the voltage range defined by the Izhikevich model. We have executed each combination of parameters 10 times, denoted as $exec = 10$, to ensure that the random selection of neurons performed is representative. The value of t^{sim} is 27s (one second per position of the optimal path), and t^{attk} , is 50ms. Table 5 summarizes the previously indicated parameters.

TABLE 5. Configuration of the implemented FLO cyberattack.

Parameter	Description	Values
t^{sim}	Duration of the simulation	27s
t^{attk}	Instant of the attack	50ms
\mathbb{NE}	Set containing all neurons of the first layer	$\{1, 2, \dots, 199, 200\}$
$ \mathbb{AN} $	Number of simultaneously attacked neurons	$\{5, 15, \dots, 95, 105\}$
\mathbb{VI}	Set containing the voltage used to attack random neurons	$\{20, 40, 60\}mV$
$exec$	Number of executions per combination of parameters	10




FIGURE 10. Number of intervening neurons related to visible positions from each position of the optimal path.

1) NUMBER OF SPIKES METRIC

To better understand the analysis of this metric, it is necessary to introduce FIGURE 10, which shows, for each position of the optimal path of the maze, the number of intervening neurons involved in the decision-making process of the mouse. Since these intervening neurons are dependent on the number of visible positions from a particular location of the maze, the number of intervening neurons is higher in central cells of the maze compared to those placed near the borders. Moreover, intervening neurons are dependent on the topology




FIGURE 11. Total number of spikes for all neurons of the topology per position of the optimal path, attacking different number of neurons (105 and 55 simultaneous neurons).

used and the convolution process of the CNN, as depicted in FIGURE 6.

FIGURE 11 compares, for the spontaneous signaling and two different configurations of FLO, the total number of spikes per position of the optimal path. In particular, the graph plots two different amounts of neurons in \mathbb{AN} (55 and 105 neurons) for all *exec* simulations. In this figure, we fixed v_i to a value of $40mV$ to improve its visualization. As can be seen, both figures share a common tendency, indicating that the higher the number of intervening neurons from a position, the higher the number of spikes. This is a consequence of how the mouse moves across the maze and how neurons and positions are related based on our particular topology. Comparing both figures, FIGURE 10 reaches its highest peaks one position before, since this change of intervening neurons needs to be propagated in time, affecting the number of spikes of its following position.

In FIGURE 11, we can see that, in general, FLO cyberattacks reduce the number of spikes compared to the spontaneous activity, increasing this reduction when the mouse progresses in the maze. Furthermore, increasing the impact of the attack, in terms of the number of attacked neurons, reduces the number of spikes. These aspects are aligned with the results later presented in Section VI-A3, where this reduction is caused by an increase of the dispersion in the attacked neurons. However, it is worth noticing the high number of spikes produced in the first position. The Izhikevich neuronal model for regular spiking generates a quick burst of spikes in a short time, and, after that, it stabilizes its spike rate, explaining this behavior. When we apply a FLO cyberattack, the attacked neurons anticipate their spikes, producing either a raise of spikes if the number of attacked neurons is not so elevated (low dispersion in time), or a reduction of spikes if most of the neurons are attacked (high dispersion). Moreover, the evolution of the simulation after the attack does not tend to come back to the spontaneous signaling, in terms of the number of spikes. In fact, these distances augment over time, reaching a difference of around 700 spikes in position 27, with some variability between both FLO configurations. Based on that, these results indicate that the effect of attacking neurons in a particular instant propagates until the end of the simulation.




FIGURE 12. Evolution of the mean of spikes with different number of attacked neurons and voltage increases, aggregating all positions of the optimal path.

After this analysis, we considered relevant to evaluate how the mean of spikes evolved through the three layers of the topology with different configurations of the FLO cyberattack. In particular, we tested different amounts of attacked neurons and voltage increase, with *exec* different executions for each combination of the previous parameters. Using *exec* executions introduces variability in terms of the randomly selected neurons for each execution. We present these results in FIGURE 12, which represents an aggregation of the number of spikes produced during the optimal path of the maze. It indicates that increasing the number of attacked neurons derives in a higher reduction in the number of spikes, while the application of different voltages does not produce a high impact. The dimmed colors surrounding the main lines of the figure indicate the fluctuations between the *exec* simulations. As can be seen, the difference in the mean of spikes compared to the spontaneous signaling grows when the number of attacked neurons raises, having a difference of around 60 spikes for 110 attacked neurons (half of the first layer). These results align with those presented in FIGURE 11 for the positions of the optimal path, where both figures present a clear descending trend when the number of attacked neurons augments. Finally, the use of different increases of voltage during the experiments did not generate a considerable impact on the number of spikes.

To expand the focus on this analysis and to determine whether this descending trend is exclusive to only certain layers, FIGURE 13 analyzes the same parameters but differentiating between the three layers of the topology and focusing only on the last position of the optimal path of the maze. We can see that the variation of the mean of spikes is more significant in deeper layers (2nd and 3rd). This variation is due to the distribution of our topology and the normal behavior of the brain, where initial layers propagate their behavior to subsequent layers, magnifying their activity via synapses. The y-axis range considerably differs between layers, being the difference with the spontaneous signaling of less than one spike in the first layer. The second layer offers a broader range of around 8 spikes in the most damaging situation, whereas the third layer has an approximate separation of between 10 to 25 spikes.

In summary, the previous figures indicate that, under attack, the mean of spikes decreases compared to the

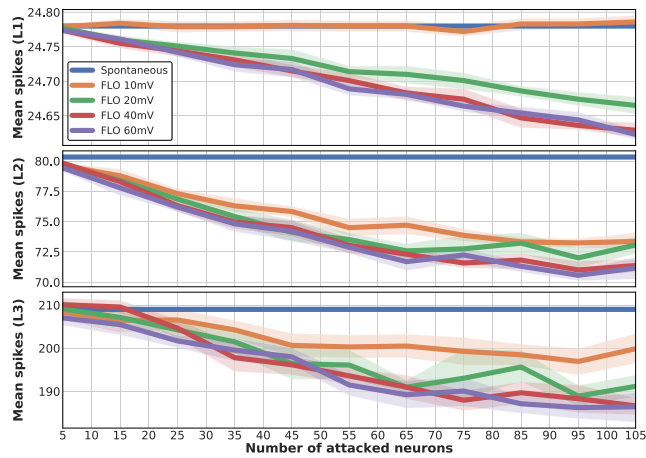


FIGURE 13. Mean of spikes for each layer of the topology, focusing on the last position of the optimal path.

spontaneous behavior. In particular, we highlight that increasing the number of attacked neurons derives in a higher impact in the mean of spikes. Nevertheless, there are no significant differences in the variation of the voltage used to attack the neurons. Finally, the number of intervening neurons from the visible positions of the optimal path of the maze strongly influences the mean of spikes.

2) PERCENTAGE OF SHIFTS METRIC

For this metric, we first evaluated the percentage of delayed shifts for an aggregation of all three layers. After that, we analyzed the same but combining all the positions of the optimal path of the maze. In this test, we included a different number of attacked neurons and voltage raises. FIGURE 14 describes this situation, where attacking a higher number of neurons produces a higher percentage of shifts. This ascending trend is aligned with the dispersion metric, since an enlargement in the parameters of the attack produces a growth of shifts. As a consequence, it generates a higher dispersion in time and number of spikes.

If we focus on each layer of the topology, FIGURE 15 represents a FLO cyberattack for the last position of the optimal path, where each color line indicates a voltage raise. Focusing on the first layer, we can see a linear growth when we augment the number of attacked neurons since only those neurons shift in the layer. Moving to subsequent layers, we can observe that the growth tendency is more prominent in the second layer. This indicates that, when we advance to the third layer, the effect of the attack gets slightly attenuated.

In conclusion, this metric indicates that attacking more neurons derives in a higher percentage of shifts. Additionally, and similarly to the metric studying the number of spikes, voltage increases have not a high impact on our scenario.

3) DISPERSION METRIC

We first focus on the spike dispersion over time caused by the different number of attacked neurons for each position of the optimal path. This means that, for each position of the




FIGURE 14. Shift percentage mean for an aggregation of all topological layers and positions of the optimal path.




FIGURE 15. Shift percentage mean for each layer of the topology, for the last position of the optimal path.

maze, we obtain the number of time instants with recorded spikes, independently of the number of spikes. If we take into account that each position of the maze corresponds to one second and that the sampling rate of Brian2, by default, is $0.1ms$, we have a total number of 10 000 instants per position. If a position presents a higher dispersion value than other positions, it indicates that there are more instants with spikes in the former one. We focus on a voltage raise value of $40mV$, since previous analysis indicated that this parameter has a low impact on our scenario.

In FIGURE 16 we can observe that the spontaneous signaling presents some similarities with the trend existing in FIGURE 11 and, specifically, in those positions with the most significant peaks. If a position presents a raise in the number of spikes, the probability of having spikes in FIGURE 16 for a longer period of time also increases. However, the natural dispersion of the simulation attenuates these peaks, where the I parameter changes according to the visible positions of the maze. Considering both FLO configurations, we can appreciate an enlargement in the temporal dispersion compared to the spontaneous behavior. FLO cyberattacks anticipate the spikes of the attacked neurons in a given moment, generating a higher dispersion as the simulation progresses. Specifically, the difference with the spontaneous signaling augments over




FIGURE 16. Spike dispersion over time for each position of the optimal path.

time, induced by the natural variability of the mouse's movements. Although the attack with 55 neurons presents a higher impact until position 17, from that position until the end, the attack with 105 neurons has a higher impact from this metric. A higher impact over the temporal dispersion when we attack more neurons simultaneously aligns with the results presented in FIGURE 13 for the number of spikes. These results are also related to those presented in Section VI-A2 for the percentage of shifts, where an intensification in these shifts derives in a dispersion growth.

We can also consider this dispersion from the perspective of the number of spikes. For each position of the optimal path, we evaluate the distribution of the number of spikes, setting the voltage increase to a value of $40mV$ and the number of simultaneous attacked neurons to 105. FIGURE 17 illustrates this distribution, where each position contains a violin plot for both the spontaneous and under attack behaviors. It is essential to highlight that this figure represents only one of the *exec* simulations performed for the complete set of experiments to ease the visualization. We can appreciate that the attack in position one reaches a peak of 110 spikes due to the increase of spikes induced by the attack performed at that particular moment. Focusing on the distribution indicated by each violin, the variance progressively reduces when the mouse progresses in the maze, concentrating the distribution of number of spikes around one. That means that in the last positions there are more instants where only one spike occurs, indicating that the attack increases the spike dispersion as the simulation progresses.

This situation aligns with the results presented in FIGURE 11, where a higher number of spikes influence this upper threshold. Nevertheless, it is worth considering the exception in position 13, where this threshold is considerably reduced. To understand this situation, we also have to consider FIGURE 16, which indicates that this position presents the highest percentage of dispersion, with more than 50% of spikes shifted. This position indicates the relationship between these two dispersion approaches, where a high temporal dispersion generates a reduction in the dispersion focused on the number of spikes.

In conclusion, FLO cyberattacks generate a large impact on the spontaneous neuronal activity. In particular, the previous figures highlight how the mouse's natural movement

induces particular natural dispersion, both in time and number of spikes. Performing FLO cyberattacks also produces an enlargement in the temporal dispersion, where the neuronal activity is more scattered. This can also be analyzed from the dispersion focused on the number of spikes since this reduction on the aggregation causes the spikes to tend to a low number. It means that there are more instants with a fewer number of spikes compared to the spontaneous behavior.

The previous analysis, based on the number of spikes, percentage of shifts, and dispersion, highlights the impact that FLO cyberattacks can generate over the spontaneous neuronal activity. We subsequently analyze these metrics together since they are strongly dependent between them. In particular, the application of a FLO cyberattack generates a decrease in the number of spikes, where these differences are more prominent in deeper layers of the topology. These results can be explained based on the dispersion induced by the attack, where a growth on the dispersion reduces the probability of multiple action potentials in the first layer. Consequently, the post-synaptic voltage raises arrive at subsequent layers in a more dispersed way, delaying the spikes. The metric focused on the percentage of shifts over the spontaneous signaling is closely related to the dispersion metric. An increase in the percentage of shifts entails a modification in the natural periodicity of the spikes. This change is directly translated to a higher dispersion rate, both in time and number of spikes. Finally, it is essential to note that this behavior and results are dependent on our particular topology. Nevertheless, they can serve as an example of how performing a FLO cyberattack can affect neuronal activity in a particular scenario.

B. NEURONAL SCANNING

This section details the implementation of an SCA cyberattack on our topology, based on the general description of the attack represented by Algorithm 2. For this particular implementation, we have sequentially attacked the 200 neurons that compose the first layer of the topology. We denote as $\mathbb{V} = \{5, 10, \dots, 60, 65\}$ the set of voltage raises, in mV , applied separately in each SCA cyberattack. As previously indicated for the FLO cyberattack, the duration of the simulation, t^{sim} , is 27s, staying the mouse one second in each position of the optimal path of the maze. Additionally, the attack initiates in the instant 50ms, represented by t^{atk} . To model the periodicity of attacking the neurons, Δt indicates the temporal separation between two attacks over two consecutive neurons, being 134ms in our particular implementation. Each combination of parameters is executed only once (*exec* = 1) since there is no variability in the selection of neurons, as it is the case of a FLO cyberattack. Finally, Table 6 indicates a summary of the parameters used in the implementation of SCA cyberattacks.

1) NUMBER OF SPIKES METRIC

FIGURE 18 compares the number of spikes per position of the optimal path between the spontaneous neuronal signaling and an SCA cyberattack. In particular, the SCA cyberattack




FIGURE 17. Spike dispersion over the number of spikes for each position of the optimal path.

TABLE 6. Configuration of the implemented SCA cyberattack.

Parameter	Description	Values
t^{sim}	Duration of the simulation	27s
t^{attk}	Start of the attack	50ms
t^{step}	Temporal duration between attacking two neurons	134ms
NE	Set containing all neurons of the first layer	{1, 2, ..., 199, 200}
VI	Set containing the voltage used to attack the neurons	{5, 10, ..., 60, 65} mV




FIGURE 18. Total number of spikes for all neurons of the topology, per position of the optimal path.

establishes a value of $40mV$ from the VI set and defines an aggregation of all three layers of the neuronal topology. We can appreciate the same trend observed in FIGURE 10 for the intervening neurons from each of the studied positions. The most prominent peaks are, as previously documented for FLO cyberattacks, delayed one position due to the time required to generate an impact over the neurons. These results can be explained based on the sequential behavior of an SCA cyberattack since the number of attacked neurons raises along time. In addition, this progressive reduction in the number of




FIGURE 19. Evolution of the spikes mean with different number of attacked neurons and voltage raises, for an aggregation of all positions of the optimal path.

spikes caused by the attack aligns with the results that will be presented in Section VI-B3 for the dispersion metric.

After this analysis, we evaluated in FIGURE 19 the mean of the spikes for the different voltage increases defined in VI, for an aggregation of the three layers of the topology and the positions of the optimal path. We can appreciate that increasing the voltage used to overstimulate the neurons produces a reduction in the number of spikes. It should be noticed that rises higher than $20mV$ do not significantly influence the impact of the attack. Performing an SCA cyberattack with a voltage of $60mV$, the most damaging situation considered, reaches the highest difference in the number of spikes, around 70 spikes compared to the spontaneous behavior.

FIGURE 20 presents a differentiation per layer of the topology for the last position of the optimal path. We can appreciate that, in the first layer, the variation in the number of spikes between different voltage increases is negligible,




FIGURE 20. Spikes mean for each layer of the topology, focusing on the last position of the optimal path.

being in all cases 24 spikes. Until $15mV$, it presents a small growth of spikes compared to the spontaneous signaling, which benefits of the anticipation of the spikes in time. In more aggressive voltages, the number of spikes gets more reduced than the spontaneous behavior. Moving to the second layer, these differences become more significant, with a number of spikes ranging between 2 and 14 spikes according to the voltage used. This layer presents a general descending trend, reaching the most damaging peak with $20mV$. This trend is common to the third layer, although the range in the number of voltages becomes broader, with a higher difference of 40 spikes compared to the spontaneous signaling. It is interesting to highlight the proliferation of spikes in the third layer when using $5mV$, based on the slight anticipation of spikes in time from the previous layers.

Comparing these results to those presented in FIGURE 19, we can appreciate in the latter specific differences in the evolution of the impact. In this figure, the most damaging voltage is $60mV$, compared to the $20mV$ highlighted for the second and third layers presented in FIGURE 20. This situation is explained by the fact that the analysis focused on differentiating the layers only considers one position and, because of that, some minor differences can arise.

In conclusion, the previous results indicate that performing an SCA cyberattack generates a reduction in the number of spikes, aggravated when the mouse moves across the maze. Increasing the voltage used to overstimulate the neurons does not produce a significant impact with voltages higher than $20mV$. Finally, the number of intervening neurons from each position of the optimal path influences this metric.

2) PERCENTAGE OF SHIFTS METRIC

FIGURE 21 first presents the results concerning the percentage of shifts for different voltage raises. These results represent an aggregation of the three layers and all the positions of the optimal path. In particular, this figure indicates that the percentage of shifts increases when we raise the voltage used to attack the neurons. We can see that an




FIGURE 21. Shift percentage mean for an aggregation of all topological layers, aggregating all positions of the optimal path.

overstimulation of $5mV$ generates an approximate 58% of shifts. Slightly increasing this voltage generates considerable impacts, between the range of $5mV$ and $20mV$, reaching a close percentage of 68%. Finally, increasing the stimulation with voltages higher than $20mV$ does not significantly enlarge the percentage of shifts. These thresholds align with those presented in FIGURE 19 for the aggregated number of spikes.

To further explore this metric, we have represented in FIGURE 22, a differentiation of each layer of the topology for just the last position of the optimal path. We can observe that the range of shifts is lower in the first layer compared to deeper layers, based on the influence that the first layer has on the latter due to the transmitted action potentials. Besides, the growth trend existing in the first layer is more prominent, being similar to the one shown in FIGURE 21 for the aggregated analysis of shifts. When we go deeper into the number of layers, we can see that the growth trend is not that aggressive using low voltages, which indicates that the attack progressively loses its effectiveness. It is important to highlight that the ranges shown in FIGURE 21 for the percentage of shifts are much higher than those presented in FIGURE 22. To understand this situation, it is worthy of reflecting on the behavior of SCA cyberattacks. In the first positions of the optimal path, only specific neurons are attacked. When the attack progresses along time, the number of neurons affected by the attack continues increasing. Based on that situation, this last figure focused on the layers presents higher ranges, since they correspond to the last position of the optimal path and, thus, all 200 neurons of the first layer have been affected.

In conclusion, performing an SCA cyberattack generates a raise in the percentage of shifts. This impact becomes more damaging when the mouse moves across the maze since the number of attacked neurons is more abundant along time. Besides, we can observe a degradation of the impact of the attack in deeper layers, where higher voltages are needed to cause a similar impact in terms of shifts.

3) DISPERSION METRIC

Focusing on the temporal dispersion caused by an SCA cyberattack, FIGURE 23 presents its analysis for each position of the optimal path and the aggregation of all the neurons of the




FIGURE 22. Shift percentage mean for each layer of the topology, only focusing on last position of the optimal path.




FIGURE 23. Spike dispersion over time for each position of the optimal path.

topology. We can observe that performing an SCA cyberattack progressively augments the temporal dispersion, based on the incremental number of attacked neurons over time. In particular, this dispersion is not significant in the first five positions of the optimal path, due to the number of attacked neurons until that moment and the specific connections of our topology.

After that, we analyze in FIGURE 24 the dispersion from the perspective of the number of spikes. In particular, we represent, for each position of the optimal path, a violin distribution of how the spikes behave. We can observe that, in the first five positions, there are no significant visual differences in the distributions, although the median of the distribution start to slightly decrease. This is justified by the reduced number of neurons affected by the attack until that instant. After that position, the differences with the spontaneous behavior progressively augment, both in the peaks in the number of spikes and the shape of the violins. Focusing on the number of spikes, the maximum number of simultaneous spikes presents a reduction, particularly in the last positions. The shape of the violins progressively changes, due to a reduction in their variance, where the number of spikes concentrates at the value of one only spike. That is to say, the majority of the instants in the last positions had only one spike. These results

are aligned to those presented in FIGURE 23 for the analysis of the temporal spike dispersion, since both figures indicate that this dispersion increases when the mouse progresses in the maze.

In summary, this metric indicates that performing an SCA cyberattack disrupts the normal neuronal spiking frequency, inducing dispersion in both temporal and number of spikes dimensions. These differences aggravate when the mouse progresses in the maze, based on the sequential functioning of SCA cyberattacks.

The previous three metrics highlight how SCA cyberattacks can affect the spontaneous neuronal activity on our particular topology. We should consider them as different perspectives to analyze a common issue. As previously indicated, an SCA cyberattack progressively induced a decrease in the number of spikes over time, aggravated in deeper layers of the topology. This decrease is strongly related to both dispersion metrics. The attack generates an alteration in the frequency of spikes in time, producing more instants with spikes in the simulation. Specifically, the previous results indicate that in the last positions of the maze, most of the instant only have one spike, which generates a clear difference with the spontaneous activity. The dispersion metric is strongly related to the percentage of shifts since this dispersion will cause a displacement of the spikes in time. In terms of shifts, the attack gets attenuated in deeper layers.

C. IMPACT COMPARATIVE BETWEEN NEURONAL FLOODING AND SCANNING

This last section compares the results previously discussed for FLO and SCA cyberattacks. Focusing on the total number of spikes (FIGURE 11 and FIGURE 18), we can observe that an SCA cyberattack generates a more impacting reduction in the number of spikes than the most aggressive FLO configuration. The last positions particularly highlight these differences.

When we analyze the number of spikes aggregating both positions and layers (FIGURE 12 and FIGURE 19), we can appreciate one of the main differences between the attacks. In FLO cyberattacks, we can define as parameters of the attack the number of neurons and the voltage used to attack those neurons. In SCA cyberattacks, we can only specify the voltage, since our implementation affects all neurons of the first layer. Based on that, there is not an immediate comparison between these figures in terms of their trend. Nevertheless, we can compare the most aggressive configuration for each attack to determine which produces the highest reduction of spikes. We can see that SCA presents a slightly higher impact than FLO.

Focusing on the distribution of spikes per layer (FIGURE 13 and FIGURE 20), we can observe that there are no significant changes between the attacks. In the second one, SCA presents a slightly lower number of spikes. Finally, the third layer amplifies these differences, where SCA has a more significant reduction of spikes.




FIGURE 24. Spike dispersion over the number of spikes for each position of the optimal path.

In terms of the percentage of shifts (FIGURE 14 and FIGURE 21), FLO presents a higher impact on this metric. Extending this comparison for each layer of the topology (FIGURE 15 and FIGURE 22), we can see that the main difference lies in the first layer, where SCA duplicates its impact since subsequent layers present similar results. Based on that, we can conclude that FLO presents a higher impact on this metric, although the difference in percentages is slight.

There is a clear difference between both attacks in terms of the temporal dispersion metric (FIGURE 16 and FIGURE 23). FLO has a higher dispersion in the first five positions of the optimal path since the targeted neurons neurons are all attacked in the same instant. After that, SCA evolves in a more damaging way. Focusing on the dispersion based on the number of spikes (FIGURE 17 and FIGURE 24), we can observe that FLO is more effective in the first positions.

This comparative highlights that the inner mechanisms of each attack generates different behaviors in the neuronal activity. FLO is adequate for attacks aiming to disrupt the neuronal activity in a short period of time, affecting multiple neurons in the same instant of time. On the contrary, SCA is a more effective attack for long-term effects, requiring a certain amount of time to reach a significant impact on the neurons. From that threshold, the impact caused on the neurons is more concerning.

VII. CONCLUSION

This work first presents security vulnerabilities of micron-scale BCI to cyberattacks, particularly for implants that can do single-cell or small population sensing and stimulation. Taking these vulnerabilities as a starting point, we describe

two novel neural cyberattacks focused on the alteration of neuronal signaling. In particular, we investigated the Neuronal Flooding (FLO) and Neuronal Scanning (SCA), inspired by well-known approaches found in the cybersecurity field. Our investigation is based on a case study of a mouse that learns its navigation within a maze trained by a Convolutional Neural Network (CNN). The CNN was converted into a biological neuronal simulation model representing the workings and functions of real neurons within the brain. The two attacks were applied to the mouse as it migrated through the maze. To evaluate the impact of these attacks on neuronal activity, we proposed three metrics: number of spikes, percentage of shifts, and dispersion of spikes, both over time and number of spikes.

A number of experiments have demonstrated that both attacks can alter the spontaneous neuronal signaling, where the behavior of these attacks generates distinct differences. FLO attacks all targeted neurons in the same instant of time, while SCA presents an incremental behavior, which requires more time to affect the neuronal activity. Focusing on the results, SCA presents a more damaging impact in terms of the number of spikes, which generates a higher reduction than FLO. In terms of shifts, FLO causes more spikes to differ in time than SCA, although these differences are not very significant. Finally, SCA presents a higher impact on the dispersion of the neurons, both in time and number of spikes. These results are highly dependent on the topology used, the neuronal model utilized to represent the neurons, and the types of neurons used (pyramidal from the primary visual cortex). Because of that, this work should be considered as a first step in the study of cyberattacks affecting spontaneous neuronal signaling.

As future work, we plan to define a taxonomy of neuronal cyberattacks affecting not only overstimulation but also neuronal activity inhibition. We aim to explore how neural cyberattacks can affect realistic neuronal tissues and, in particular, various neural circuits within the cortex. Our research lays the groundwork for security countermeasures to also be integrated into BCI systems that utilize miniature implants for small neuronal population stimulation that can have a tremendous effect on the brain.

ACKNOWLEDGMENT

The authors would like to thank Luigi Petrucco, Ethan Tyler, and SciDraw for their publicly-available scientific images [41], [42].

REFERENCES

- [1] M. A. Lebedev and M. A. L. Nicolelis, "Brain-machine interfaces: From basic science to neuroprostheses and neurorehabilitation," *Physiol. Rev.*, vol. 97, no. 2, pp. 767–837, Apr. 2017.
- [2] L. Yao, X. Sheng, N. Mrachacz-Kersting, X. Zhu, D. Farina, and N. Jiang, "Sensory stimulation training for BCI system based on somatosensory attentional orientation," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 3, pp. 640–646, Mar. 2019.
- [3] J. E. O'Doherty, M. A. Lebedev, P. J. Ifft, K. Z. Zhuang, S. Shokur, H. Bleuler, and M. A. L. Nicolelis, "Active tactile exploration enabled by a brain-machine-brain interface," *Nature*, vol. 479, pp. 228–231, 2011.
- [4] T. Kaufmann, S. M. Schulz, A. Köblitz, G. Renner, C. Wessig, and A. Käbler, "Face stimuli effectively prevent brain–computer interface inefficiency in patients with neurodegenerative disease," *Clin. Neurophysiol.*, vol. 124, no. 5, pp. 893–900, May 2013.
- [5] M. Pais-Vieira, G. Chiuffa, M. Lebedev, A. Yadav, and M. A. L. Nicolelis, "Building an organic computing device with multiple interconnected brains," *Sci. Rep.*, vol. 5, no. 1, Dec. 2015, Art. no. 11869.
- [6] D. Sempredoni and L. Viganá, "Privacy, security and trust in the Internet of neurons," 2018, *arXiv:1807.06077*. [Online]. Available: <http://arxiv.org/abs/1807.06077>
- [7] R. A. Ramadan and A. V. Vasilakos, "Brain computer interface: Control signals review," *Neurocomputing*, vol. 223, pp. 26–44, Feb. 2017.
- [8] A. Thomson, S. Tielens, T. Schuhmann, T. De Graaf, G. Kenis, B. Rutten, and A. Sack, "The effect of transcranial magnetic stimulation on living human neurons," *Brain Stimulation*, vol. 12, no. 2, p. 524, Mar. 2019.
- [9] D. Seo, J. M. Carmena, J. M. Rabaey, E. Alon, and M. M. Maharbiz, "Neural dust: An ultrasonic, low power solution for chronic brain-machine interfaces," 2013, *arXiv:1307.2196*. [Online]. Available: <https://arxiv.org/abs/1307.2196>
- [10] S. A. Wirdatmadja, M. T. Barros, Y. Koucheryavy, J. M. Jornet, and S. Balasubramaniam, "Wireless optogenetic nanonetworks for brain stimulation: Device model and charging protocols," *IEEE Trans. Nanobiosci.*, vol. 16, no. 8, pp. 859–872, Dec. 2017.
- [11] E. M. Neuralink. (2019). *An Integrated Brain-Machine Interface Platform With Thousands of Channels*. [Online]. Available: <https://www.biorxiv.org/content/early/2019/08/02/703801>
- [12] P. Ballarin Usieto and J. Minguez. (2018). *Avoiding Brain Hacking—Challenges of Cybersecurity and Privacy in Brain Computer Interfaces*. [Online]. Available: <https://www.bitbrain.com/blog/cybersecurity-brain-computer-interface>
- [13] M. Frank, T. Hwu, S. Jain, R. T. Knight, I. Martinovic, P. Mittal, D. Perito, I. Sluganovic, and D. Song, "Using EEG-based BCI devices to subliminally probe for private information," in *Proc. Workshop Privacy Electron. Soc.*, 2017, pp. 133–136.
- [14] I. Martinovic, D. Davies, and M. Frank, "On the feasibility of side-channel attacks with brain-computer interfaces," in *Proc. 21st USENIX Secur. Symp.*, 2012, pp. 143–158.
- [15] K. Sundararajan, "Privacy and security issues in brain computer interface," M.S thesis, Auckland Univ. Technol., Auckland, New Zealand, 2017. [Online]. Available: <http://orap.auckland.ac.nz/bitstream/handle/10292/11449/SundararajanK.pdf>
- [16] M. Ienca, "Neuroprivacy, neurosecurity and brain-hacking: Emerging issues in neural engineering," *Bioethica Forum*, vol. 8, no. 2, pp. 51–53, 2015.
- [17] M. Ienca and P. Haselager, "Hacking the brain: Brain–computer interfacing technology and the ethics of neurosecurity," *Ethics Inf. Technol.*, vol. 18, no. 2, pp. 117–129, Jun. 2016.
- [18] H. Takabi, A. Bhalotiya, and M. Alohalay, "Brain computer interface (BCI) applications: Privacy threats and countermeasures," in *Proc. IEEE 2nd Int. Conf. Collaboration Internet Comput. (CIC)*, Nov. 2016, pp. 102–111.
- [19] T. Bonaci, R. Calo, and H. J. Chizeck, "App stores for the brain: Privacy and security in brain-computer interfaces," *IEEE Technol. Soc. Mag.*, vol. 34, no. 2, pp. 32–39, Jun. 2015.
- [20] Q. Li, D. Ding, and M. Conti, "Brain-computer interface applications: Security and privacy challenges," in *Proc. IEEE Conf. Commun. Netw. Secur. (CNS)*, Sep. 2015, pp. 663–666.
- [21] S. L. Bernal, A. H. Celdrán, G. M. Pérez, M. T. Barros, and S. Balasubramaniam, "Cybersecurity in brain-computer interfaces: State-of-the-art, opportunities, and future challenges," 2019, *arXiv:1908.03536*. [Online]. Available: <http://arxiv.org/abs/1908.03536>
- [22] S. F. Lorenzo, J. A. Benito, P. G. Cardarelli, J. A. Garaia, and S. A. Juaristi, "A comprehensive review of RFID and Bluetooth security: Practical analysis," *Technologies*, vol. 7, no. 1, p. 15, Jan. 2019, doi: [10.3390/technologies7010015](https://doi.org/10.3390/technologies7010015).
- [23] S. Sevier and A. Tekeoglu, "Analyzing the security of Bluetooth low energy," in *Proc. Int. Conf. Electron., Inf., Commun. (ICEIC)*, Jan. 2019, pp. 1–5.
- [24] M. Zubair, D. Unal, A. Al-Ali, and A. Shikfa, "Exploiting Bluetooth vulnerabilities in e-Health IoT devices," in *Proc. 3rd Int. Conf. Future Neww. Distrib. Syst.*, 2019, pp. 1–7, doi: [10.1145/3341325.3342000](https://doi.org/10.1145/3341325.3342000).
- [25] NIST. (2019). *CVE-2019-16336*. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-16336>
- [26] NIST. (2019). *CVE-2019-19192*. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-19192>
- [27] NIST. (2019). *CVE-2019-19194*. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-19194>
- [28] L. Pycroft, S. G. Boccard, S. L. F. Owen, J. F. Stein, J. J. Fitzgerald, A. L. Green, and T. Z. Aziz, "Brainjacking: Implant security issues in invasive neuromodulation," *World Neurosurg.*, vol. 92, pp. 454–462, Aug. 2016.
- [29] S. Vaddamani, B. Eksioğlu, H. Medal, and A. Nandi, "Jamming attacks on wireless networks: A taxonomic survey," *Int. J. Prod. Econ.*, vol. 172, pp. 76–94, Feb. 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S092552731500451X>
- [30] E. Gal, M. London, A. Globerson, S. Ramaswamy, M. W. Reimann, E. Muller, H. Markram, and I. Segev, "Rich cell-type-specific network topology in neocortical microcircuitry," *Nature Neurosci.*, vol. 20, no. 7, pp. 1004–1013, Jul. 2017, doi: [10.1038/nn.4576](https://doi.org/10.1038/nn.4576).
- [31] A. Gáron, *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. Newton, MA, USA: O'Reilly Media, 2019.
- [32] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
- [33] I. Kuzovkin, R. Vicente, M. Petton, J.-P. Lachaux, M. Baciú, P. Kahane, S. Rheims, J. R. Vidal, and J. Aru, "Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex," *Commun. Biol.*, vol. 1, no. 1, p. 107, Aug. 2018, doi: [10.1038/s42003-018-0110-y](https://doi.org/10.1038/s42003-018-0110-y).
- [34] S. Zafrany. (2013). *Deep Reinforcement Learning for Maze Solving*. [Online]. Available: <https://www.samyza.com/ML/rl/qmaze.html>
- [35] F. Chollet. (2015). *Keras*. [Online]. Available: <https://keras.io>
- [36] M. Stimberg, R. Brette, and D. F. Goodman, "Brian 2, an intuitive and efficient neural simulator," *eLife*, vol. 8, Aug. 2019, Art. no. e47314.
- [37] R. A. Tikidji-Hamburyan, V. Narayana, Z. Bozkus, and T. A. El-Ghazawi, "Software for brain network simulations: A comparative study," *Frontiers Neuroinform.*, vol. 11, p. 46, Jul. 2017. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fninf.2017.00046>
- [38] E. M. Izhikevich, "Simple model of spiking neurons," *IEEE Trans. Neural Netw.*, vol. 14, no. 6, pp. 1569–1572, Nov. 2003.
- [39] L. Bachatene, V. Bharmuria, and S. Molotchnikoff, "Adaptation and neuronal network in visual cortex," in *Vision Cortex*, S. Molotchnikoff and J. Rouat, Eds. Rijeka, Croatia: IntechOpen, 2012, ch. 15, doi: [10.5772/46011](https://doi.org/10.5772/46011).

- [40] A. M. Norcia, L. G. Appelbaum, J. M. Ales, B. R. Cottreau, and B. Rossion, "The steady-state visual evoked potential in vision research: A review," *J. Vis.*, vol. 15, no. 6, p. 4, May 2015, doi: [10.1167/15.6.4](https://doi.org/10.1167/15.6.4).
- [41] L. Petrucco. *Mouse Head Schema*. Accessed: Jul. 17, 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.3925903>
- [42] E. Tyler. *Mouse*. Accessed: Jul. 17, 2020. [Online]. Available: <https://zenodo.org/record/3925901>



SERGIO LÓPEZ BERNAL received the B.Sc. and M.Sc. degrees in computer engineering from the University of Murcia, and the M.Sc. degree in architecture and engineering for the IoT from IMT Atlantique, France. He is currently pursuing the Ph.D. degree with the University of Murcia. His research interests include ICT security on brain-computer interfaces and network and information security.



ALBERTO HUERTAS CELDRÁN received the M.Sc. and Ph.D. degrees in computer science from the University of Murcia, Spain. He is currently an Irish Research Council Government of Ireland Postdoctoral Research Fellow Associated with the TSSG, Waterford Institute of Technology, Ireland. His scientific interests include cybersecurity, privacy, brain-computer interfaces (BCI), continuous authentication, and computer networks.



LORENZO FERNÁNDEZ MAIMÓ received the M.Sc. and Ph.D. degrees in computer science from the University of Murcia, Spain. He is currently an Associate Professor with the Department of Computer Engineering, University of Murcia. His research interests include machine learning and deep learning applied to cybersecurity and computer vision.



MICHAEL TAYNNAN BARROS (Member, IEEE) received the B.Tech. degree in telematics from the Federal Institute of Education, Science, and Technology of Paraíba, in 2011, the M.Sc. degree in computer science from the Federal University of Campina Grande, in 2012, and the Ph.D. degree in telecommunication software from WIT, in 2016. He is currently a Lecturer with the School of Computer Science and Electronic Engineering, University of Essex, U.K. He was a recipient of the Marie Skłodowska Curie Individual Fellowship (MSCA-IF).



SASITHARAN BALASUBRAMANIAM (Senior Member, IEE) received the B.E. degree in electrical and electronic engineering from the University of Queensland, Brisbane, QLD, Australia, in 1998, the M.E. degree in computer and communication engineering from the Queensland University of Technology, Brisbane, in 1999, and the Ph.D. degree from the University of Queensland, in 2005. He is currently the Director of Research with the TSSG, Waterford Institute of Technology, Ireland. His current research interests include molecular communications, the Internet of (Bio) NanoThings, and terahertz wireless communications.



GREGORIO MARTÍNEZ PÉREZ (Member, IEEE) received the M.Sc. and Ph.D. degrees in computer science from the University of Murcia, Spain. He is currently a Full Professor with the Department of Information and Communications Engineering, University of Murcia. His research interests include cybersecurity, privacy, and networking, working on different national and European IST research projects on these topics.

...