Politecnico di Torino

Politecnico di Torino

Master's Degree in Biomedical Engineering

Adaptive Neurostimulation through Binaural Beats for Stress Modulation

Supervisor:

Prof. Luca Mesin

Candidates:

Emanuele Della Corina

Leonardo Iaconinoto

Academic Year 2025/2026

Abstract

Stress is a complex physiological response to conditions perceived as perturbing the organism's internal balance and involves the autonomic nervous system, resulting in a wide range of physiological responses, including cardiovascular and electrodermal changes. The aim of this work is to design and experimentally evaluate an adaptive neuromodulation system based on Binaural Beats (BB) for the modulation of cognitive stress. The study is organized into three phases: development of a protocol for stress induction, design of a system for real-time detection of the physiological state, and implementation of an algorithm for modulating the beat frequency according to the subject's response. The stress detection system is based on a One-Class Support Vector Machine (OC-SVM) model, which uses features extracted from electrocardiogram (ECG) and electrodermal activity (EDA) signals, acquired using MotemaSens (OT Bioelettronica), to estimate physiological deviation from the resting state. The experimental protocol includes a Baseline phase, used to train the OC-SVM, and three conditions based on a timed mathematical task: No stimulation, Sham (fixed stimulation), and Adaptive BB (ABB). Since the OC-SVM Score tended to saturate during high-deviation conditions, BB modulation was driven by ScoreBB, a variable derived from the normalized Mahalanobis distance (MD) between the current feature vector and the individual baseline distribution. Through sigmoid mapping and temporal smoothing, ScoreBB was used to dynamically adjust the beat frequency within the alpha band, shifting it toward relaxation-related values when physiological deviation increased. The protocol was completed by 14 healthy subjects. Statistical analysis showed that the mathematical task induced a physiological response compatible with cognitive stress, as indicated by significant increases ($p < 0.01$) in heart rate (HR) and MD from Baseline to No stimulation. In the comparison among task conditions, HR was significantly lower ($p < 0.001$) in both Sham and ABB than in No stimulation, with mean reductions of 4.01 bpm and 4.61 bpm, respectively. MD was also significantly reduced ($p < 0.01$) in both stimulation conditions, with mean reductions of 1.65 and 4.35, indicating that the ECG–EDA feature vector was closer to the individual baseline. The classifier Score showed a significant difference ($p < 0.05$) only between No stimulation and ABB, with a mean increase of 0.362. However, no significant differences were found between Sham and ABB, suggesting that the specific contribution of adaptive binaural modulation could not be clearly separated from the general acoustic effect. An exploratory analysis of stimulus delivery order was also performed by dividing the sample into two balanced subgroups: Sham-Adaptive BB (S-ABB) and Adaptive BB-Sham (ABB-S). In the S-ABB group, the reduction in deviation from baseline was more evident during the ABB condition. Therefore, the sequence of stimulus delivery could rep-

resent a relevant element in the evaluation of adaptive stimulation effectiveness. In conclusion, the obtained results support the validity of the stress induction protocol and suggest a potential benefit of acoustic stimulation, with some evidence in favour of adaptive stimulation. Future developments should include a larger sample and a more detailed analysis of electrodermal components. In addition, the task used to induce stress could be made more variable and interactive across conditions to limit learning effects.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Stress Overview

Stress is an intrinsic part of human life and involves several dimensions of the individual. A stressful condition triggers a series of physical, physiological, and behavioural reactions aimed at maintaining the organism's internal balance, known as homeostasis [1].

Stressors that disrupt this internal equilibrium can range from acute events, such as the fear experienced during a potentially life-threatening accident, to chronic conditions, such as prolonged emotional tension. Stress can also arise in response to different types of stimuli, including physical injury, biochemical imbalance, or exposure to pathogenic agents [2].

1.1.1 Historical Background

The scientific study of stress has deep roots in physiology, with important contributions dating back to the late 19th century. The French biologist Claude Bernard was among the first to formalise the concept of internal equilibrium, introducing the idea that organisms continuously work to maintain a stable internal environment. This principle was later defined as homeostasis [3].

Building on Bernard's ideas, Walter Cannon further highlighted the role of the autonomic nervous system (ANS) in maintaining homeostasis. The ANS is divided into sympathetic nervous system (SNS) and parasympathetic nervous system (PNS). He stated that any tendency toward internal change is counteracted by opposing physiological mechanisms. Cannon also extended the concept beyond purely physical stressors by linking his theory of emotion with his theory of homeostasis. In this context, he proposed that the release of adrenaline into the bloodstream has several adaptive functions, allowing the organism to respond to

acute stress through the so-called fight-or-flight response. This represented one of the first descriptions of a coordinated physiological stress reaction [3].

While Cannon mainly focused on acute stress responses, Hans Selye shifted attention toward chronic stress. In 1956, Selye introduced the term "stress" and defined it as any threat that disrupts homeostasis. He identified several hormones involved in the stress response, with particular attention to glucocorticoids, and showed that both humans and animals present similar physiological reactions to stressors. These reactions include behavioural changes and activation of the endocrine and neuroendocrine systems. His work on the effects of chronic stress on the hypothalamic-pituitary-adrenal (HPA) axis contributed to later studies investigating the link between chronic stress and HPA axis hyperactivity [3, 4].

1.1.2 Physiological Responses to Stress

The physiological responses triggered by stressors involve a complex interaction between neurological, endocrine, and immune processes [5].

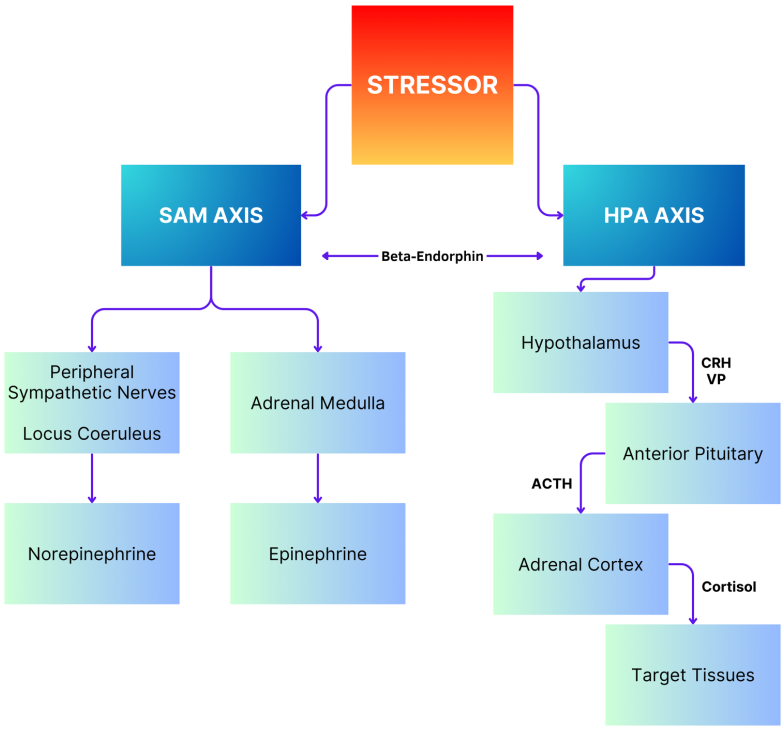

Figure 1.1: Schematic representation of stress physiology. Note: ACTH: adrenocorticotropic hormone; CRH: corticotropin-releasing hormone; HPA: hypothalamic-pituitary-adrenal; SAM: sympathetic-adrenal-medullary; VP: vasopressin.

At the physiological level, the stress response consists of two main components: a fast response mediated by the sympathetic-adrenal-medullary (SAM) axis and a slower response mediated by the HPA axis [5].

Activation of the SAM system causes the immediate release of epinephrine and norepinephrine into the bloodstream, together with norepinephrine release from sympathetic nerve endings [6]. Epinephrine and norepinephrine interact with α- and β-adrenergic receptors in the central nervous system, in smooth muscle cells, and in several organs throughout the body. Activation of these receptors produces a wide range of physiological effects. In the cardiovascular system, heart rate, blood pressure, and cardiac output increase. Blood glucose levels also rise, supporting greater oxygen consumption and energy production. At the same time, peripheral vasoconstriction occurs, while blood flow to skeletal muscles increases [7].

Unlike the SAM axis, the HPA axis is responsible for the slower component of the stress response. This pathway begins with the release of corticotropin-releasing hormone (CRH) from the hypothalamus. In response to CRH, the anterior pituitary gland releases adrenocorticotropic hormone (ACTH) into the bloodstream. ACTH then stimulates the adrenal glands to produce and secrete glucocorticoids, mainly cortisol [7].

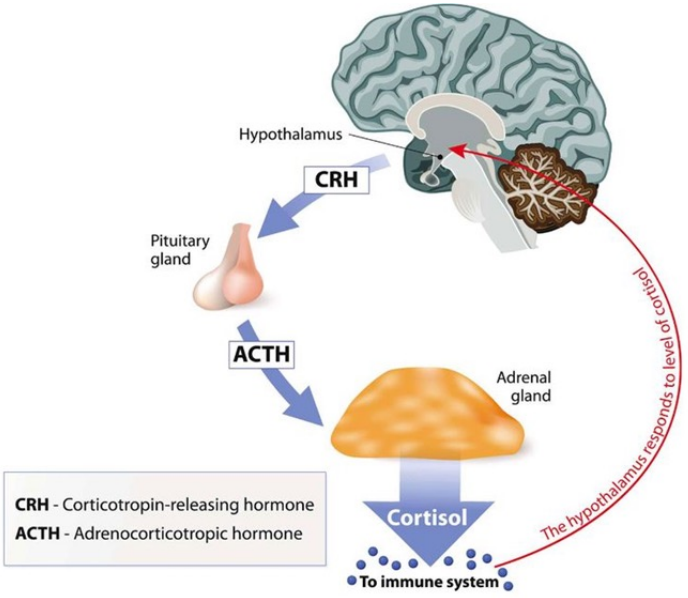

Figure 1.2: Illustration of HPA axis activation during stress.

1.2 Heart

1.2.1 Overview of the Cardiovascular System

The cardiovascular system is an organ system whose primary function is to meet the metabolic demands of the entire organism by transporting essential substances throughout the body. It is composed of the heart, blood, and blood vessels [8]. The heart acts as a pump that propels blood through the vascular network, delivering oxygen and nutrients to body tissues while removing carbon dioxide and metabolic waste products. In addition to its circulatory role, the cardiovascular system also performs sensory and endocrine functions, working in coordination with blood vessels to help regulate blood pressure and maintain homeostasis. Blood serves not only as a medium for the transport of nutrients and waste materials, but also as a carrier of hormones, facilitating communication between different organs and systems of the body [8].

From a functional perspective, the cardiovascular system is divided into two main circuits [8]:

- Pulmonary circuit, which consists of the blood vessels connecting the lungs and the heart;
- Systemic circuit, which comprises all remaining vessels that distribute blood to the rest of the body.

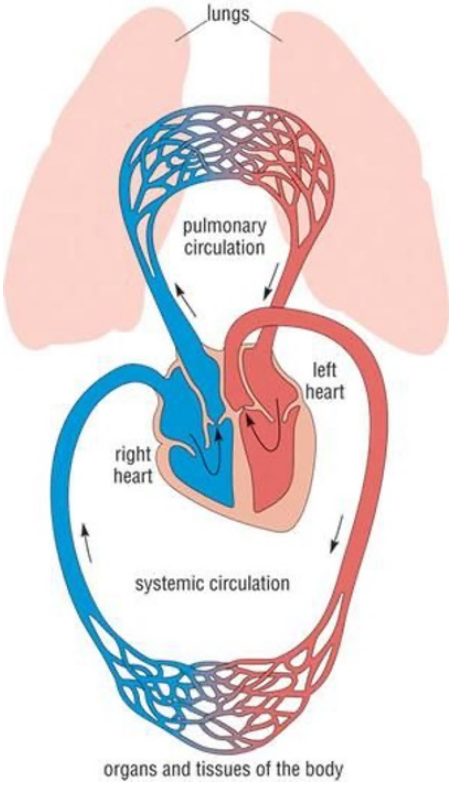

Figure 1.3: Anatomy of the Cardiovascular System.

1.2.2 Anatomy of the heart

The heart is a muscle located in the center of the thoracic cavity, just above the diaphragm. Its weight ranges between 250 and 350 grams, depending on the individual's sex and body mass [8]. The cardiac muscle generates the force required to propel blood throughout the body through rhythmic contractions and relaxations. Each contraction creates the pressure necessary to pump blood into the blood vessels.

The heart is composed of four chambers: two atria and two ventricles [8]. The ventricles have thicker walls because they must pump blood over greater distances, whereas the atria only push blood into the adjacent ventricles. The left ventricle pumps blood to the entire body, while the right ventricle sends it only to the lungs. The atria contract first, pushing blood into the ventricles, which then contract to distribute blood to the rest of the body. It is essential that blood flow proceeds in only one direction; to ensure this, the heart is equipped with four cardiac valves that maintain unidirectional flow [8].

The atria and ventricles are connected by atrioventricular (AV) valves, which

allow blood to flow from the atria to the ventricles but prevent backflow. The AV valves open and close according to the pressure difference between the chambers: when atrial pressure exceeds ventricular pressure, the valves open; when ventricular pressure becomes greater, they close.

In addition to the AV valves, there are the semilunar valves, located between the ventricles and the major arteries. These valves perform a similar function by allowing blood to flow in a single direction. The pulmonary valve is located between the right ventricle and the pulmonary trunk, while the aortic valve lies between the left ventricle and the aorta [8].

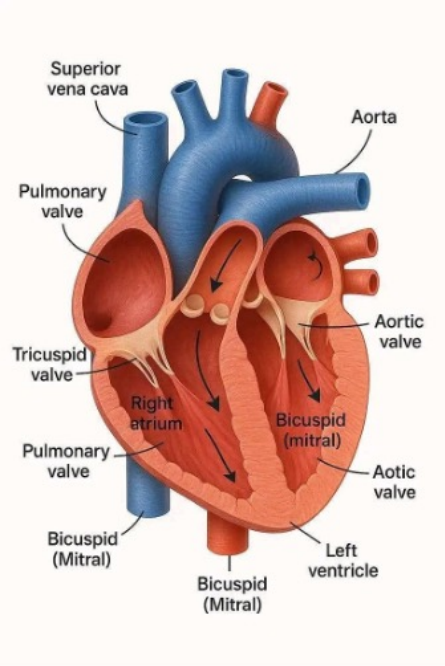

Figure 1.4: Anatomy of the heart.

1.2.3 Electrical Activity of the Heart

For the heart to function properly, it must contract in a synchronized manner: the atria contract first, followed by the ventricles [8]. Cardiac contractions are generated by signals that originate within the cardiac muscle itself; therefore, the contractile activity of the heart is described as myogenic.

The autorhythmicity of the heart is maintained by two specialized types of cells [8]:

- Pacemaker cells, which generate action potentials and determine the heart's rhythm;

- Conduction fibers, which ensure the organized propagation of electrical impulses across the cardiac tissue.

Pacemaker cells are primarily located in the sinoatrial (SA) node and the atrioventricular (AV) node. The propagation of the action potential through the heart produces a depolarization, which is subsequently followed by muscle contraction. The electrical impulse originates in the SA node, which depolarizes the atria completely before reaching the AV node, where conduction is delayed. This delay at the AV node is essential, as it allows the atria to contract before the ventricles, ensuring efficient ventricular filling prior to contraction [8].

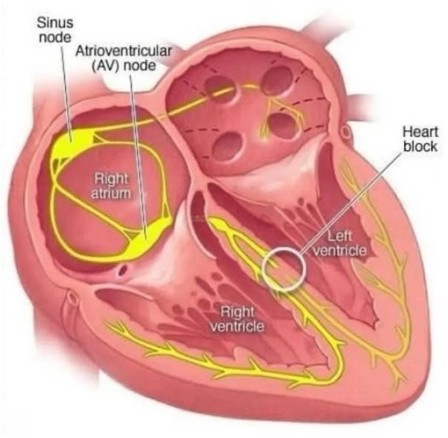

Figure 1.5: Electrical activity of the heart.

1.3 Electrocardiography (ECG)

Electrocardiography is today an essential component of the initial evaluation of patients presenting with cardiac symptoms. Specifically, it plays an important role as a non-invasive and cost-effective diagnostic tool for assessing arrhythmias and ischemic heart disease [9].

An ECG records the electrical currents that flow through the heart during a cardiac cycle, using electrodes placed on the skin [8, 9]. The standard procedure for recording an ECG is based on Einthoven's triangle. This technique allows the electrical activity of the heart to be observed from three different directions [9]. The model forms a triangle around the heart, with its vertices located at the right arm, left arm, and left leg. An electrode is placed at each vertex, and these electrodes are connected in pairs to a voltage recorder. Each pair of electrodes forms what is known as a lead. Thus, there are three standard limb leads:

- Lead I measures the potential difference between the left arm and right arm;

- Lead II measures the potential difference between the left leg and right arm;
- Lead III measures the potential difference between the left leg and left arm [8].

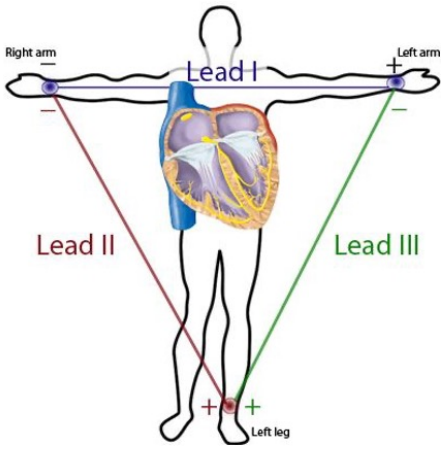

Figure 1.6: Einthoven triangle.

A standard 12-lead electrocardiogram includes electrodes placed on both the limbs and the chest wall. Each lead provides a different electrical perspective of cardiac activity [9].

In the ECG tracing, the PQRST complex represents different phases of the cardiac cycle and provides key information about heart function [10].

- The P wave corresponds to atrial depolarization;
- The QRS complex reflects ventricular depolarization;
- The T wave represents ventricular repolarization.

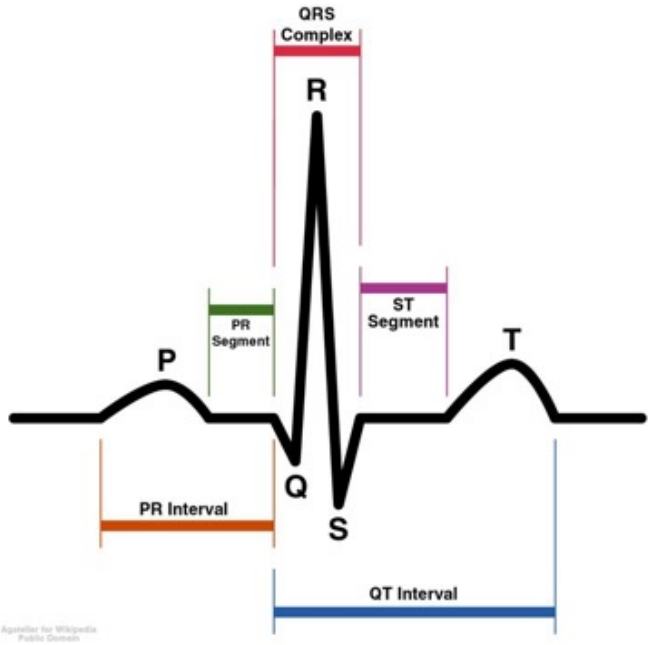

Figure 1.7: ECG of a healthy subject.

Atrial repolarization is not visible on the ECG because it is masked by the QRS complex [10]. In addition to the waves, the intervals on the ECG also provide valuable diagnostic information.

- The P–Q (or P–R) interval represents the conduction time through the AV node;
- The Q–T interval corresponds to the duration of ventricular contraction, known as ventricular systole, while the T–Q interval (or diastolic interval) represents the ventricular relaxation phase;
- The R–R interval estimates the time between two consecutive heartbeats [10].

1.4 Electrodermal Activity (EDA)

Electrodermal activity (EDA) is a physiological signal that reflects changes in the electrical conductance of the skin. These changes are determined by the activity of the eccrine sweat glands, which are innervated exclusively by SNS. Sympathetic activation increases sweat secretion, even when there is no visible sweating. This reduces skin resistance and leads to a measurable increase in conductance through electrodes placed on the skin surface, typically on the fingers or on the palm of

the hand [11,12]. For this reason, EDA is one of the few physiological signals that provides a direct measure of peripheral sympathetic activity.

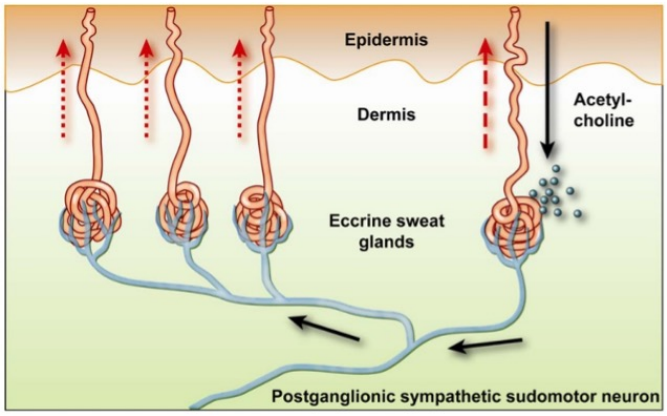

Figure 1.8: Physiological mechanism of EDA.

In terms of signal structure, EDA is traditionally described as the sum of two main components: a tonic component, called Skin Conductance Level (SCL), and a phasic component, represented by Skin Conductance Responses (SCRs). The tonic component represents the baseline level of skin conductance and changes slowly over time, reflecting the general activation level of the ANS. The phasic component, instead, is characterized by rapid transient responses associated with discrete events, such as sensory, cognitive, or emotional stimuli, and appears as peaks superimposed on the tonic level [11–13].

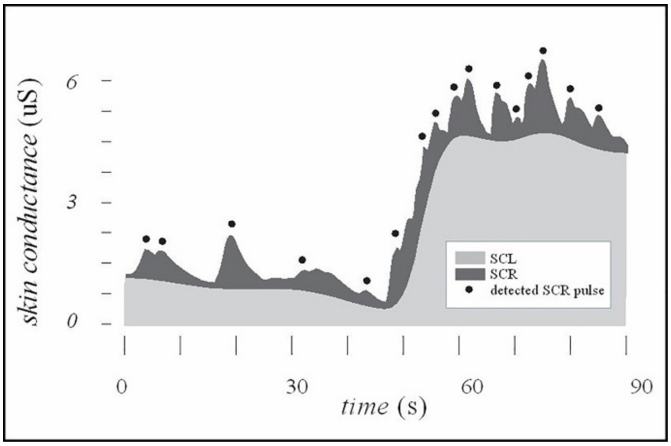

Figure 1.9: Representation of the EDA signal with the tonic component (SCL) and phasic responses (SCRs).

The EDA signal has a relatively low amplitude, generally expressed in microsiemens (µS). Typical SCL values range from about 1 to 20 µS in healthy subjects, depending on physiological, environmental, and individual conditions. SCRs have lower amplitudes, often between 0.01 and 3 µS, and last for a few seconds, with a rapid rising phase followed by a slower decay [12, 13].

The amplitude and frequency of phasic responses are influenced by the intensity and relevance of the stimulus, as well as by the psychophysiological state of the subject. In terms of spectral content, EDA is a very low-frequency signal, with most of its energy concentrated below 0.5 Hz. The tonic component is associated with very slow variations, below 0.05 Hz, while SCRs mainly contribute to the frequency band between 0.05 and 0.5 Hz. This spectral distribution makes EDA particularly suitable for time-domain analysis and decomposition techniques that separate tonic and phasic components, rather than for a purely spectral analysis [11, 14].

The relationship between EDA and stress is well established in psychophysiological literature. Stress, whether psychological or cognitive, is associated with SNS activation, which results in an increase in skin conductance. Several studies have shown that stressful conditions induce significant increases in SCL, as well as an increase in the frequency and amplitude of SCRs compared with resting or relaxation conditions [13–15].

As a result, EDA is widely used as an indicator of SNS activation and as a non-invasive biomarker of stress. However, it is important to highlight that EDA is not a specific indicator of stress. Instead, it more generally reflects the subject's level of sympathetic activation and arousal. For this reason, changes in the signal should be interpreted by considering the experimental context and, when possible, integrated with other physiological signals [11, 13, 15].

In recent years, interest in EDA has further increased thanks to the spread of wearable devices, which allow continuous signal acquisition in ecological and daily-life contexts. Systematic reviews in the literature show that features derived from EDA are informative for monitoring perceived stress, especially when integrated into multimodal systems together with other physiological signals, such as heart rate and skin temperature [14, 15].

In this context, integrating EDA with cardiovascular signals, such as those derived from ECG, makes it possible to obtain a more complete representation of the stress response. This combination brings together information related to peripheral sympathetic activation with indicators of cardiac activity and autonomic regulation [15].

The combination of a solid physiological basis, non-invasive measurement, and high sensitivity to changes in psychological state makes EDA one of the reference signals in the study of stress and emotional responses. Despite these ad-

vantages, the EDA signal is sensitive to several external and individual factors, including environmental temperature, skin hydration, movement, electrode position, and contact pressure. These aspects make appropriate signal acquisition and pre-processing procedures necessary.

1.5 Adaptive Neurostimulation

Adaptive neurostimulation is an advanced approach within the field of neuromodulation, in which the parameters of the stimulus are dynamically adjusted according to the physiological or neural state of the subject. Unlike traditional stimulation techniques, which are based on static and predefined protocols, this approach introduces an automatic regulation mechanism based on the use of biological signals as indicators of the state of the nervous system [16, 17].

Traditionally, neurostimulation has been implemented using an open-loop approach, in which the stimulation parameters are set at the beginning and remain constant over time. As a result, the stimulus is delivered independently of the subject's response. Although this strategy has been applied in several clinical contexts, such as the treatment of movement disorders, epilepsy, and chronic pain, it presents some limitations because it cannot adapt to the temporal and interindividual variability of physiological processes. The closed-loop approach, on which adaptive neurostimulation is based, makes it possible to overcome these limitations, since stimulation is modulated in real time according to physiological signals acquired from the subject [18, 19].

In adaptive neurostimulation systems, the key principle is the presence of a feedback mechanism, through which the information extracted from biological signals is used to continuously update the parameters of the stimulus. This allows a form of automatic control of the nervous system to be implemented, where stimulation is no longer static but changes over time according to the subject's internal state [16, 17].

From a functional point of view, a closed-loop system consists of several phases: signal acquisition, extraction of relevant features, estimation of the physiological state, and modulation of stimulation. The signals used may reflect neural activity, muscle activity, or indices of ANS activity. The analysis of these signals is based on advanced processing techniques and machine learning algorithms, which make it possible to model the relationship between the observed features and the functional states of interest [17, 20].

A crucial aspect of adaptive neurostimulation is its ability to operate in dynamic biological systems. The nervous system is characterized by strong variability over time and is highly dependent on the physiological and behavioral context. The same stimulus can produce completely different effects depending on the subject's

state. Closed-loop systems make it possible to take this variability into account by adapting stimulation according to the current conditions, potentially improving the effectiveness of the intervention [17].

Interest in this approach has further increased because of its possible application in the modulation of cognitive and emotional processes. The use of closed-loop systems has opened new perspectives in the study and regulation of affective states and anxiety, where the ability to adapt the stimulus according to the subject's state represents a key element [21]. Recent studies have also shown that closed-loop systems may offer important advantages over open-loop approaches, including greater efficiency, fewer side effects, and better personalization of stimulation [19, 20].

In the context of stress regulation, adaptive neurostimulation represents a particularly interesting approach, since stress is not a static condition but a dynamic state that changes over time depending on individual, environmental, and cognitive factors. For this reason, a system able to modulate stimulation according to the subject's physiological state may be more suitable than fixed and predefined protocols. This perspective is especially relevant for the development of non-invasive stimulation strategies aimed at promoting psychophysiological states that are more compatible with relaxation and reduced arousal [17, 21].

The implementation of adaptive neurostimulation systems offers several promising opportunities, but it still presents some challenges. These include the need to identify reliable biomarkers, manage noise in physiological signals, and design robust control strategies. However, advances in acquisition technologies and data analysis methods are helping to make these systems increasingly effective and applicable, including in non-invasive contexts [17, 20].

1.6 Binaural Beats

Binaural beats (BB) are a specific form of auditory stimulation that can modulate brain activity through central processing mechanisms of the auditory system. They are generated by presenting two pure tones with slightly different frequencies, usually below 1000 Hz, simultaneously to each ear through stereo headphones. Under these conditions, the auditory system does not perceive two separate sounds, but rather a single rhythmic oscillation whose frequency corresponds to the difference between the two presented frequencies. This produces the auditory sensation of a beat. This phenomenon is not physically present in the acoustic stimulus itself, but arises from central neural processing [22].

Figure 1.10: Generation of a 10 Hz BB.

From a neurophysiological point of view, BB perception occurs at the level of the brainstem, mainly in the superior olivary complex, a structure involved in binaural integration and spatial sound localization. This feature distinguishes BB from other forms of rhythmic auditory stimulation, such as monaural beats, in which the interference between frequencies occurs at the peripheral level. Because of the central nature of this phenomenon, BB may be useful as a tool for modulating brain activity [22].

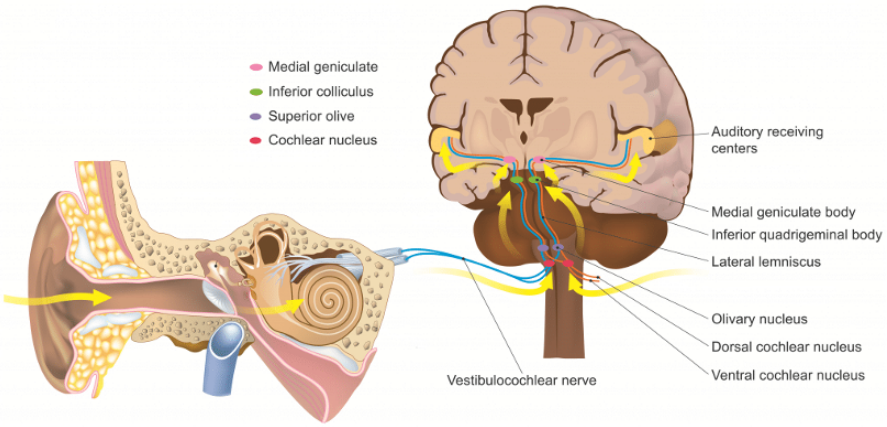

Figure 1.11: Auditory pathways.

Exposure to BB may be associated with a modulation of brain oscillations, which can be observed through electroencephalography (EEG) recordings. This suggests the possible presence of neural entrainment, namely the synchronization of neural activity with the frequency of the perceived beat. However, this effect

depends on several factors, including the beat frequency, the duration of stimulation, the experimental context, and the individual characteristics of the subject. BB-induced entrainment generally appears to be weaker and less stable than the entrainment obtained with other forms of rhythmic stimulation, although it can still be detected under specific experimental conditions [23].

BB can be used to modulate different cognitive and emotional states, including attention, relaxation, alertness, and anxiety. The frequency bands commonly used in binaural stimulation reflect the classical classification of EEG oscillations, with the aim of promoting mental states that are consistent with the functional characteristics of each band. However, the observed effects are generally modest and show high interindividual variability, so the results should be interpreted with caution [24].

BB also have an important application in stress modulation. Their potential use has mainly been linked to their ability to indirectly influence the balance between SNS and PNS, through the modulation of arousal states, understood as the activation levels of the ANS, and brain oscillations associated with relaxation. In particular, binaural stimulation at frequencies within the theta and alpha bands has been associated with reduced alertness, mental relaxation, and decreased perceived anxiety. These conditions are physiologically compatible with a reduction in sympathetic activation. Stress is characterized by an increase in physiological arousal and SNS activity, leading to changes in the cardiovascular, electrodermal, and neuroendocrine systems. The possibility of promoting relaxation through non-invasive auditory stimulation suggests that BB may act as indirect modulators of stress responses, reducing the intensity of the associated physiological reactions. However, the effects of BB should not be interpreted as a direct intervention on the ANS, but rather as a central modulation of mental states that, in turn, influence peripheral physiological responses [24, 25].

The use of BB as a brain stimulation tool also presents some methodological limitations and critical issues. The results reported in the literature are highly variable due to several factors, including the lack of standardized protocols, difficulties in controlling placebo effects, differences in stimulation duration, the frequency used, and the dependence on experimental conditions. In addition, the response to BB can vary significantly between subjects, making the interpretation of results more complex and questioning the effectiveness of BB-induced neural entrainment as a generalizable effect [25].

This variability makes BB particularly interesting in the context of adaptive neurostimulation. Unlike traditional stimulation, which is applied using fixed parameters and independently of the subject's current state, an adaptive system allows the stimulus to be modulated according to the detected psychophysiological condition. In the case of stress, this feature is particularly relevant, since individ-

ual responses can change over time and depend on personal, environmental, and contextual factors [26, 27].

1.7 Machine Learning

Machine Learning (ML) is a subset of artificial intelligence focused on the development of algorithms capable of learning automatically from data. ML models improve their performance through experience, without the need to be explicitly programmed for each specific task. In this context, learning refers to the ability of a model to identify structures, relationships, and regularities within data, in order to make predictions or decisions on new, previously unseen observations [28].

One of the main classifications of ML algorithms is based on the type of input data and learning strategy [29]:

- Supervised learning: the model is trained using a labelled dataset, in which each input example is associated with a correct output. This approach is commonly used in classification and regression problems, where the objective is to learn a function that correctly maps inputs to their corresponding labels;
- Unsupervised learning: the model operates on unlabelled data and aims to discover hidden structures or patterns, such as clusters or lower-dimensional representations of the original data;
- Reinforcement learning: the model learns through iterative interaction with an environment, receiving reward or penalty signals based on the actions taken. The objective is to develop a strategy that maximises cumulative reward over time.

A crucial aspect in the development of ML models is their ability to generalise, meaning the capacity to maintain good performance on data that were not included in the training set. An overly complex model may fit the training data too closely, also capturing noise and losing predictive ability on new data. This condition is known as overfitting. Conversely, a model that is too simple may fail to capture meaningful relationships in the data, leading to underfitting. Managing this bias-variance trade-off is one of the fundamental challenges of machine learning [30].

1.7.1 Support Vector Machine

Support Vector Machines (SVM) are supervised learning algorithms widely used for classification and regression problems. They are appreciated for their solid theoretical basis and good generalisation ability. SVMs were first introduced by

Cortes and Vapnik, who showed that they are based on the identification of an optimal decision boundary that separates classes by maximising the margin between samples belonging to different classes [31].

In the case of linearly separable data, the objective is to find a hyperplane that maximises the distance from the nearest points of each class, known as support vectors. This strategy produces robust classifiers that are less sensitive to data fluctuations and noise. In real-world settings, however, data often overlap and are not perfectly separable. To address this issue, the concept of soft margin is introduced, allowing a certain number of classification errors through the use of slack variables [31, 32].

The trade-off between margin maximisation and error penalisation is regulated by the regularisation parameter C, which controls the weight assigned to classification errors in the cost function. High values of C tend to favour correct classification of the training data, but may reduce model robustness. Lower values promote a wider margin and may improve generalisation. The appropriate choice of this parameter is therefore important for balancing model complexity and generalisation ability [32].

When data are not linearly separable in the original feature space, SVMs can use the kernel trick, a technique that implicitly projects data into a higher-dimensional space where linear separation becomes possible. Kernel functions, such as the polynomial kernel or the radial basis function (RBF) kernel, allow complex non-linear relationships to be modelled without explicitly computing the transformed feature space [33].

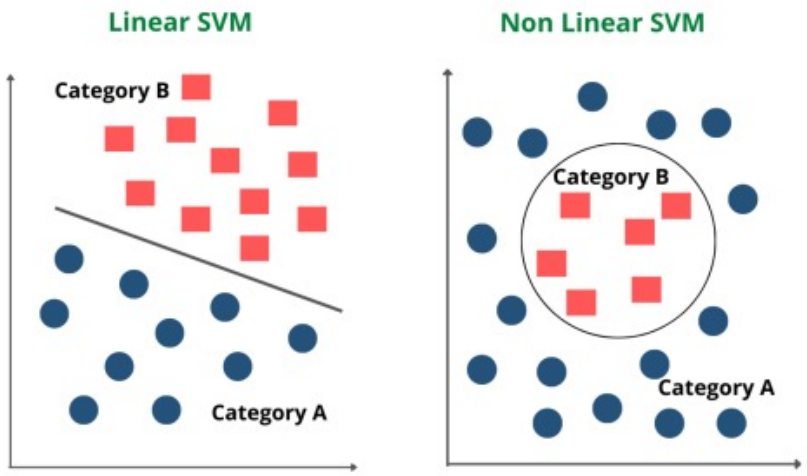

Figure 1.12: Comparison between linear and non-linear SVM.

In the case of Support Vector Regression (SVR), in addition to the parameter

C, the parameter ε is introduced. This parameter defines an insensitive region around the regression function, within which errors are not penalised. This allows more robust models to be obtained in the presence of noise and can reduce the number of support vectors. The parameter ε therefore plays an important role in controlling model complexity and its ability to adapt to noisy data [32, 33].

Thanks to these properties, SVMs are particularly effective in contexts characterised by a high number of features and relatively small datasets. The combination of margin maximisation, control of model complexity, and kernel functions makes SVMs a versatile and reliable tool in several application domains, especially when good accuracy and generalisation ability are required [30, 32].

1.7.2 One-Class Support Vector Machines

Traditional SVMs are supervised classification algorithms that usually require labelled data from two or more classes. However, in many real-world applications, data from the normal condition are available, whereas examples of anomalous conditions are rare, incomplete, or difficult to define. This scenario is commonly referred to as novelty detection or one-class classification.

Schölkopf et al. developed One-Class SVMs (OC-SVM) as an extension of SVMs to this type of problem [34]. The main objective of the OC-SVM is to model the distribution of data belonging to the normal class by constructing a boundary that encloses the region of the feature space where these data are most concentrated. After training, new observations are compared with this region and classified as normal if they fall within the identified boundary, or as anomalous if they fall outside it.

Formally, the OC-SVM seeks to determine a hyperplane that separates the training data from the origin in the kernel-transformed feature space. The optimisation process aims to maximise the distance between this hyperplane and the origin, while ensuring that most of the training samples lie on the same side of the decision boundary.

Similarly to standard SVMs, OC-SVMs rely on the use of kernel functions to model non-linear data distributions. In addition, the parameter ν plays a central role in controlling the behaviour of the model. It defines an upper bound on the fraction of training samples that can be considered outliers and, at the same time, a lower bound on the fraction of support vectors used by the model. As ν approaches zero, the model becomes more conservative and tends to include a larger portion of the training data within the normal region [34].

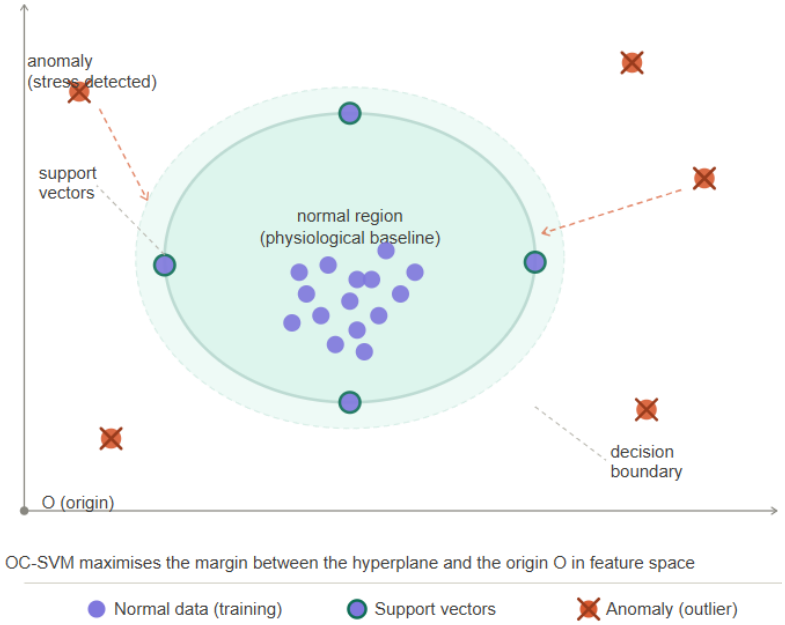

Figure 1.13: One-Class SVM decision boundary estimated from normal-class training data.

The OC-SVM is well suited to applications in which it is necessary to detect deviations from a reference condition. In the context of physiological monitoring, the normal class can correspond to the resting or baseline state of the subject. This subject-dependent approach makes it possible to reduce the limitations caused by inter-subject physiological variability, which can affect supervised classifiers trained on heterogeneous populations [35].

1.8 State of the Art

In recent years, automatic stress detection has received growing attention in biomedical engineering, psychophysiology, and wearable technologies. Traditionally, stress assessment has relied on subjective questionnaires, psychometric scales, or neuroendocrine measures, such as cortisol. These tools have provided useful results, but they also present several limitations, including low temporal resolution, dependence on self-assessment, and limited applicability in real time. For this reason, research has progressively moved toward the use of non-invasive physiological signals, which can continuously describe changes in ANS activity associated with stress [27].

Among the signals most commonly used for stress detection are cardiovascular

signals, particularly ECG and Heart Rate Variability (HRV) indices. HRV reflects the interaction between SNS and PNS in the regulation of cardiac rhythm and therefore represents a relevant indicator of autonomic state. Several studies have shown that psychological or cognitive stress conditions are associated with changes in heart rate and RR interval variability. In general, increased arousal and sympathetic activation tend to modify the balance of the ANS, making HR and HRV parameters frequently used in automatic stress recognition systems [36].

EDA also plays an important role as one of the most widely used indicators for stress assessment. As described in the section on electrodermal activity, EDA reflects changes in skin conductance related to the activity of eccrine sweat glands, which are mainly controlled by the SNS. This makes the signal particularly sensitive to changes in arousal. However, EDA is not a stress-specific indicator, since it can also increase in response to other factors, such as attention, emotion, sensory stimuli, or movement. For this reason, the literature highlights the importance of interpreting EDA within the experimental context and, when possible, integrating it with other physiological signals [11, 13, 15].

The multimodal approach, based on the integration of cardiovascular and electrodermal signals, represents one of the most established directions in the field of stress detection. Compared with the use of a single signal, the combination of ECG and EDA provides a more complete description of the psychophysiological response, since it combines information related to cardiac regulation and peripheral sympathetic activation. This type of approach is particularly useful in machine learning systems, where different physiological features are combined to classify or estimate the subject's state [15, 27].

Regarding the algorithms used, the literature reports the application of several machine learning methods for stress detection, including SVM, Random Forest, neural networks, and deep learning models. Supervised algorithms are often used when labelled data related to stress and relaxation conditions are available, whereas personalized or one-class approaches can be useful when an individual reference condition, such as baseline, is available. This strategy is particularly interesting in physiological monitoring, since it allows the resting state of each subject to be modelled and subsequent deviations from this profile to be identified, reducing as much as possible the impact of inter-subject variability [27, 35].

Despite the progress achieved, stress detection still presents several critical issues. A first limitation concerns the individual variability of the stress response: different subjects may show different physiological reactions to stress, even when exposed to the same stimulus. In addition, many studies are conducted in controlled environments and on relatively small samples, making model generalization to real-world contexts complex. Another critical aspect concerns the distinction between stress and general arousal: signals such as EDA and HRV may reflect an

increase in physiological activation, but they do not necessarily identify a stress condition in a specific way. For this reason, it is essential to accurately define the experimental context and the task used to induce stress [27].

A second research area concerns the development of techniques for stress mitigation. Among the most studied strategies in the literature are biofeedback, guided breathing, meditation, digital interventions through smartphones, wearable technologies, and different forms of sensory stimulation. Biofeedback is based on providing subjects with information about their physiological state, such as heart rate, HRV, respiration, or electrodermal activity, with the aim of promoting self-regulation. The available studies show promising results, but also strong heterogeneity in terms of protocols, intervention duration, signals used, and outcomes considered [37].

In recent years, the evolution of wearable technologies has promoted the development of adaptive and real-time interventions, in which support for stress regulation is delivered when the system detects an increase in physiological activation. This approach represents an evolution compared with traditional interventions, since it does not simply propose generic relaxation strategies, but aims to provide a personalized intervention according to the current state of the subject. Automatic stress detection therefore no longer represents only a monitoring tool, but becomes an integral part of a regulation system [38].

Among the techniques used to mitigate stress, closed-loop or adaptive neuromodulation also plays an important role, as described in the previous sections. In this approach, stimulation is modified in real time based on biological signals acquired from the subject, in order to take into account the dynamic nature of stress and the variability of individual responses. This opens the way to more personalized interventions compared with fixed-stimulation protocols [17, 19, 20].

Among the non-invasive stimulation techniques proposed to modulate cognitive and emotional states, BB represent a strategy of particular interest. As described in the section on binaural beats, they consist of the separate presentation of two pure tones with slightly different frequencies to the two ears, generating the central perception of a beat at the frequency corresponding to the difference between the two stimuli. Some studies have investigated the possible role of BB in the modulation of anxiety, attention, relaxation, and pain perception. In particular, frequencies belonging to the theta and alpha bands have often been associated with relaxation states and reduced arousal [24, 25].

Based on the studies reported in the literature, the effectiveness of BB remains under discussion. The available results are heterogeneous and suggest that the effects depend on several factors, including beat frequency, stimulation duration, target band, individual characteristics, type of experimental control, and administration context. In addition, brain entrainment induced by BB is not

always consistently demonstrated. Therefore, BB can be considered among the promising strategies for stress modulation, but their effectiveness is not yet fully established [39].

1.9 Aim of the Study

The main aim of this work is to design an adaptive neuromodulation system based on BB for stress modulation. To achieve this objective, the work is organized into three main phases:

1. Development of an experimental protocol able to induce cognitive stress in the participants;
2. Design of a system for real-time stress detection, based on the analysis of the acquired physiological signals;
3. Design of a BB-based stimulus modulation algorithm adapted to the estimated physiological state of the subject.

Finally, the proposed system is experimentally evaluated by comparing the different conditions included in the experimental protocol.

Chapter 2

Materials and Methods

2.1 Instrumentation and Software

2.1.1 MotemaSens and sensors

The ECG and EDA signals were recorded using the MotemaSens device, produced by OT Bioelettronica, Turin, Italy [40, 41]. MotemaSens is an amplifier with eight independent analog channels, identified as CH1–CH8, designed for the acquisition of biopotentials. The device is portable and can operate in both wired and wireless mode, allowing data transfer via USB connection, storage on an SD card, or Bluetooth (BT) transmission [40,41]. In this study, the USB mode was used to connect the amplifier to the PC. From a hardware perspective, the system integrates an STM32 microcontroller based on ARM Cortex-M4 architecture, responsible for managing signal acquisition and external communications, and a high-resolution multichannel analog-to-digital converter, ADS1298, specifically designed for biopotential applications and characterized by low noise and a wide dynamic range. The firmware, developed in C using the STM32CubeIDE environment, manages the configuration of the peripherals, including USB, SPI, UART, and I2C, as well as the synchronous sampling of the signals [40, 41]. When connected to a computer running Windows, MotemaSens is automatically recognized as a virtual serial port, COM. Communication is managed through a Matlab script provided by the manufacturer, which allows the acquisition parameters to be set, including sampling frequency, duration, and recording mode. The same script also allows start/stop commands to be sent and electrocardiographic traces to be displayed in real time. In addition to visualization, the script decodes the digital samples and saves the raw data for subsequent processing [40, 41].

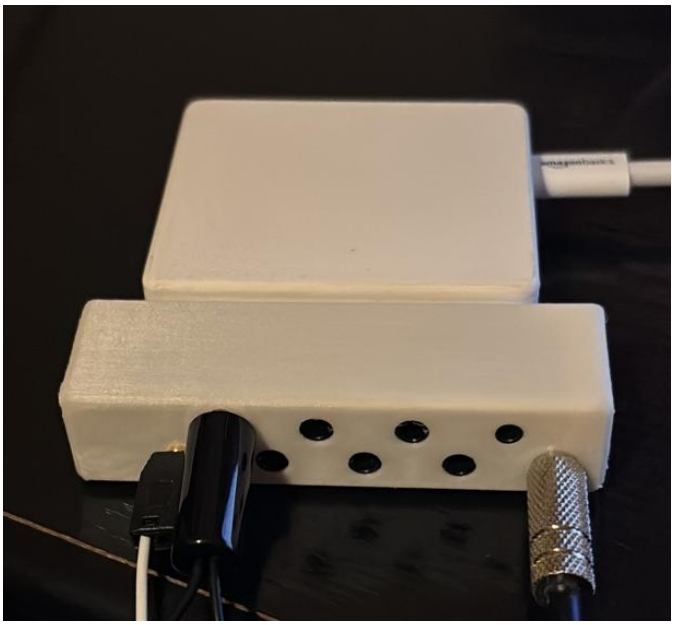

Figure 2.1: MotemaSens amplifier used for ECG and EDA signal acquisition, shown with the acquisition cables connected.

The amplifier includes two adapters with different electronic boards. Both adapters include two I2C ports, used to interface with external sensors, and a PAT REF connector for the reference input. The channels for biopotential acquisition consist of six 3.5 mm monopolar inputs, corresponding to channels CH1–CH6. Channels CH7 and CH8, instead, are associated with 2.5 mm mono jack inputs dedicated to microphones or other auxiliary sensors. The supplied adapters are:

- EMO adapter: channels CH1–CH5 allow ECG or electromyographic (EMG) signals to be acquired, while channels CH6–CH8 can be used to connect chest strap, temperature, and resistance sensors, respectively;
- EMG adapter: channels CH1–CH6 allow only EMG signals to be acquired, while channels CH7 and CH8 can be used to connect microphones.

This modular configuration allows the device to be adapted to different types of sensors and experimental protocols. Communication with the computer can take place either via USB or BT, since both modes use the same command protocol and data format, ensuring full functional equivalence between the two interfaces [40,41].

In order to acquire ECG and EDA signals, the EMO adapter was used in this study. The ECG signal was acquired using a single-channel bipolar configuration, CH1, with two bracelet electrodes positioned on the wrists and a reference electrode positioned on the left forearm, connected to the PAT REF connector. Before each

acquisition, the bracelet electrodes were wetted to reduce the electrode-skin contact impedance, thereby improving electrical coupling and the stability of the acquired signal.

The EDA signal was acquired using a GSR, Galvanic Skin Response, sensor based on a Grove interface, developed by the electronics company Seeed Studio. This interface provides a standardized modular connection for simplified integration of electronic devices. The GSR module used in this study integrates a skin resistance sensor with an analog output and can be easily interfaced through a four-pin connector. The sensor operates by applying a weak voltage between two electrodes placed on the skin surface and measuring the resulting voltage variation through a resistive divider circuit. The output is therefore an analog voltage signal proportional to skin resistance. The GSR sensor used includes two electrodes, which were positioned on the index and middle fingers of the left hand and fixed using elastic straps [42].

Figure 2.2: GSR sensor.

In addition to the MotemaSens device and the GSR sensor, the following equipment was used during the experimental protocol:

- Btootos Bluetooth noise-cancelling headphones, used for the auditory stimulation;
- MacBook Pro computer, used for the processing of the acquired signals and for the implementation of the experimental protocol;
- 22-inch LG television, used as an external screen for displaying the graphical interface of the mathematical task.

2.1.2 Matlab

A main Matlab script was developed to implement the different phases of the protocol. The script is divided into two sections, which manage the classifier training phase and the real-time acquisition phase, respectively.

In the first section, the initial parameters for training are set. These include the sampling frequency, which is the same as that used by the amplifier for signal acquisition, the length of the signal window to be analyzed, and the overlap between windows. Scrolling plots are then initialized to visualize the signals in real time during the baseline phase of the protocol, in order to check signal quality and identify any possible anomalies. Signal acquisition during this phase is managed by a while loop, within which the function that enables communication with the amplifier is called, updating both the signal buffers and their visualization. The number of loop iterations is defined according to the duration of the training phase, considering one iteration as corresponding to the time interval at which the device updates the data.

The second section of the script is dedicated to the real-time acquisition and processing of physiological signals. In this phase, the previously trained model is first loaded, together with the feature normalization parameters. Then, as in the first section, the parameters required for online processing are initialized, including the analysis window length and the overlap between consecutive windows. The acquired signals are stored in temporary buffers, which are progressively updated at each iteration of the acquisition loop. As in the training phase, a real-time visualization system for ECG and EDA signals was implemented using scrolling plots, allowing signal quality to be monitored during the experiment. In this case as well, data acquisition is managed by a while loop, within which the new samples coming from the acquisition device are read and concatenated to the buffers. When the buffer length reaches the size of the analysis window, signal processing is performed. For each window, features are extracted from the ECG and EDA signals and then normalized using the parameters calculated during the training phase. The normalized features are then provided as input to the classifier, which returns an estimate of the subject's state. An additional measure is then calculated based on the distance of the features from the baseline distribution, namely the Mahalanobis distance, which is used to obtain a continuous indicator of the stress state. To reduce the variability of estimates within a single measurement, the value used in the subsequent phases is obtained as the average over multiple consecutive windows.

In parallel, a Matlab timer was used to manage audio generation, allowing the stimulation to run asynchronously with respect to the main acquisition loop. At each timer activation, a callback function, audioCallback, is called. Its task is to update the audio signal according to the current state of the system. The callback

function retrieves the most recent value of the stress indicator and uses it as input for the BB generation function. At the end of the acquisition, the classification results and the extracted features are saved for subsequent offline analyses.

The function that manages communication with MotemaSens was developed starting from the script provided by the manufacturer. The time interval at which the data are updated is represented by the variable refresh, set to 200 ms, while the sampling frequency and the conversion factor in mV are set to 500 Hz and 0.000286, respectively. The ECG acquisition mode and USB connection are then selected for data transfer.

Model training is implemented using a script that manages signal segmentation into epochs and feature extraction. Feature extraction is performed through two functions dedicated to the ECG and EDA variables, respectively. EDA preprocessing is performed directly within the function that extracts its features, whereas, since ECG preprocessing is more complex, an additional function was implemented. In this one, the custom rico function, described in Section 2.3.1, was used to apply the 50 Hz notch filter. This function returns the filter coefficients based on the minimum attenuation required at 50 Hz and on the sampling interval. Training is performed using the built-in fitcsvm function, provided by the Matlab Statistics and Machine Learning Toolbox. In particular, this function was used to build the SVM model from the features extracted from the physiological signals, which had been previously normalized. Training was performed using a Gaussian kernel, RBF, specified through the KernelFunction parameter (KernelFunction = 'rbf'), while the KernelScale parameter was set automatically (KernelScale = 'auto'). The fraction of samples considered anomalous was controlled through the OutlierFraction parameter, which allows the model sensitivity to deviations from the baseline to be adjusted. Finally, after calculating the average Mahalanobis distance obtained for each window, the resulting model is returned as a Matlab object and saved to be integrated into the real-time processing pipeline.

BB generation was implemented through a function dedicated to the real-time synthesis and playback of stereo audio signals. The signal was built by generating two sinusoidal waves, one for each audio channel, with frequencies centered around a predefined carrier frequency. The two signals were then combined to form a stereo signal, which was sent to the output through the audioDeviceWriter interface, using a sampling frequency of 44.1 kHz. Signal synthesis was performed in frame-based mode, generating successive blocks of audio samples. To ensure signal continuity between consecutive frames, persistent variables were used to track the phase of the sinusoids. The function also includes mechanisms for initializing and releasing audio resources, allowing controlled start and stop of playback. In particular, a stop mode was implemented to interrupt the audio output and release the playback device. When the sham operating mode is implemented, the

two audio channels have the same frequency, eliminating the binaural component while maintaining an acoustic stimulus. The audio signal generated in this way is updated in real time during the execution of the experimental protocol and is integrated into the overall system pipeline.

2.1.3 Python

The mathematical task used to induce cognitive stress was implemented in Python through the development of a dedicated script, using the tkinter library to create the graphical user interface. The application was structured by defining a main class responsible for managing the entire task flow, including interface initialization, event handling, and dynamic content updating.

The application is based on an event-driven architecture, in which the execution of the different functions is guided by events generated by the user and by the system. This structure ensures interface responsiveness and avoids the use of blocking loops. The graphical interface consists of several elements: labels for displaying information, such as the number of correct and incorrect answers and the timer, an input field for entering the user's answers, and components for managing visual feedback. User interaction is managed through keyboard events, particularly by detecting the pressing of the Enter key.

The temporal management of the task is implemented through the timing functions provided by the tkinter library, in particular using the after method, which allows periodic updates to be performed without interrupting the main loop of the graphical interface. In this way, it is possible to simultaneously monitor the overall duration of the task and the time constraints associated with individual responses.

The interface is updated dynamically by modifying the textual content and graphical properties of its elements according to the user's actions and the current state of the application. At the end of execution, the system automatically disables the interactive elements and displays a task completion message.

2.2 Protocol Design

The protocol was divided into two main phases. The first phase aimed to collect ECG and EDA signals from participants in a resting state, in order to train the classifier. The second phase involved stress induction and auditory stimulation to assess the effects of BB. The entire procedure lasted approximately 30 minutes: about 22 minutes were devoted to the execution of the experimental protocol, while about 10 minutes were required for electrode setup and removal, as well as for explaining the tasks to the participant.

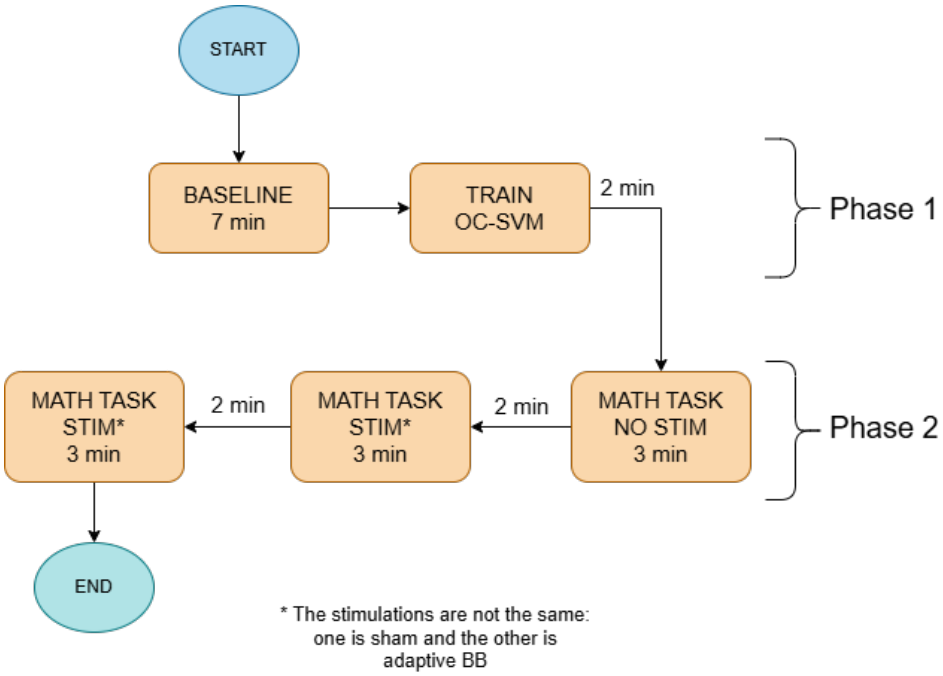

Figure 2.3: Experimental protocol flowchart.

2.2.1 Phase 1 — Baseline Acquisition

The first phase played a crucial role in the correct functioning of the entire protocol. Training the OC-SVM classifier with data that accurately reflected the participant's resting state was essential for the success of the experiment. To promote this condition, both the baseline phase and the entire protocol were conducted in a room free from noise and external distractions.

The baseline acquisition phase lasted 7 minutes. ECG and EDA signals were recorded using the MotemaSens amplifier, described in Section 2.1.1, together with the associated sensors described in the same section.

During baseline acquisition, participants were asked to close their eyes. At the same time, relaxing instrumental ambient music was played through the headphones described in Section 2.1.1. The selected track, titled "Bella musica rilassante per dormire per alleviare lo stress • Calma la mente, medita (volando)", was chosen because of its slow harmonic progression, absence of vocal content, and moderate tempo, which are features commonly associated with relaxation induction in the literature [43]. The track is publicly available on YouTube [44].

Throughout the entire recording phase, participants remained silent and were instructed to relax as much as possible. At the end of the 7-minute acquisition, the recorded data were processed by a Matlab script that trained the classifier in

offline mode. After this phase, the participant was given a 2-minute rest period.

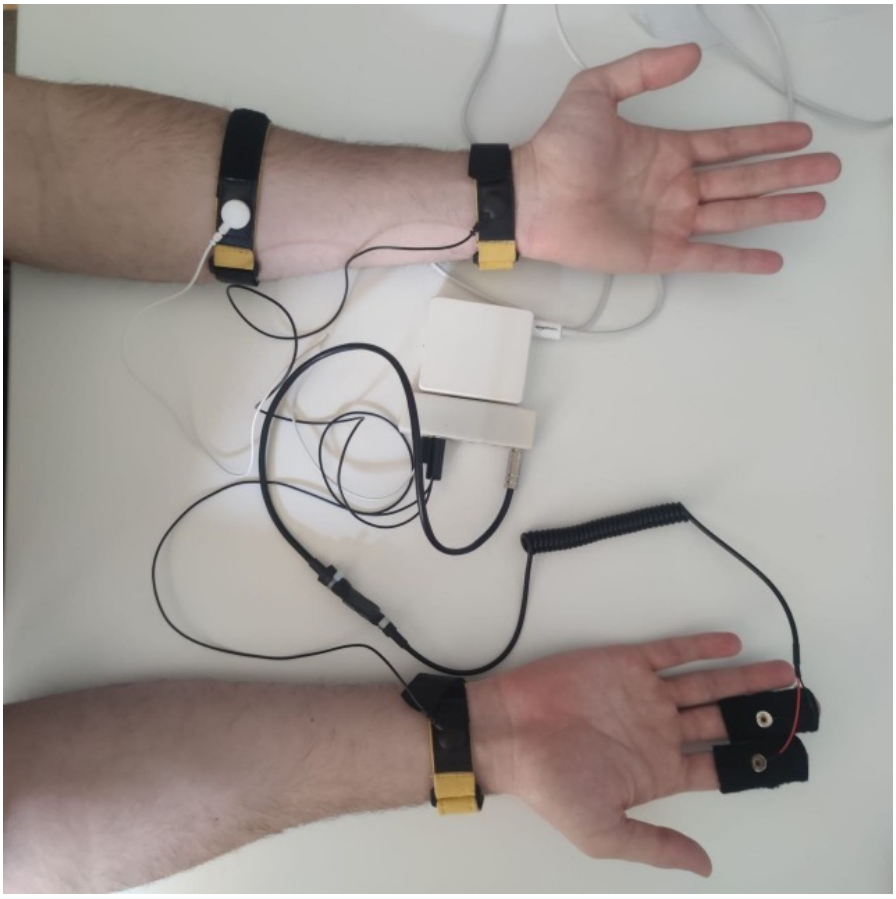

Figure 2.4: Participant with electrodes attached during the protocol.

2.2.2 Phase 2 — Stress Induction and Stimulation

The second phase consisted of three mathematical tasks, each separated by a 2-minute rest period. Each task lasted 3 minutes and was designed to induce cognitive stress in the participant.

The tasks were administered through a Python application. A number between 900 and 1200 appeared on the screen, and the participant had 7 seconds to subtract 7 from the displayed number and communicate the result to the operator. If the answer was correct and provided within the time limit, the resulting number appeared on the screen and the participant had to perform the subtraction again. This procedure was repeated until the end of the task.

The presence of the operator, who manually entered the result communicated by the participant, was necessary to prevent motion artifacts and EMG noise that

could have been induced by typing on a keyboard.

Figure 2.5: Graphical interface of the mathematical task.

During this second phase, the first mathematical task was administered without any sound through the headphones. In the second and third tasks, an auditory stimulus was delivered: in one case, a sham stimulus was used, while in the other, an adaptive BB stimulus was delivered according to the modulation described in Section 2.7.

The sham stimulus consisted of a pure tone at the carrier frequency of 400 Hz, delivered identically to both ears with no interaural frequency difference (df = 0 Hz). In this way, the participant perceived a continuous and stable sound that was perceptually similar to the carrier component of the BB stimulus, but without the interaural frequency difference that represents the active component of binaural stimulation. This design made it possible to control for the placebo effect by separating the generic acoustic effect from the specific effect of BB.

The assignment of the two stimuli to the second and third tasks followed a counterbalanced scheme. Each participant had a 50% probability of receiving the adaptive BB stimulus in the second task and the sham stimulus in the third task, or vice versa. The protocol was designed so that 50% of participants received the sham stimulus first, while the remaining 50% received the adaptive stimulus

first. This counterbalancing was introduced to evaluate whether adaptive BB could produce beneficial effects on stress mitigation even in sessions following their administration, while also controlling for possible order effects.

Before the beginning of this phase, participants were informed that, during the second and third mathematical tasks, an auditory stimulus would be delivered through the headphones with the aim of mitigating the stress accumulated during the tasks. In this way, participants were aware that a sound would be present in the headphones, but they were unable to determine its nature. This approach helped balance the placebo effect between the sham and adaptive BB stimuli, ensuring single-blind experimental conditions.

2.3 Signal Analysis

This section describes the pre-processing pipeline applied to the ECG and EDA signals. This phase is essential for the correct execution of the protocol: appropriate filtering reduces noise and motion artifacts, producing cleaner and more reliable signals. After pre-processing, the signals are ready for segmentation and feature extraction, which are then used to train the OC-SVM classifier.

The entire pre-processing and training pipeline was implemented in Matlab.

2.3.1 ECG

The ECG signal was acquired using the MotemaSens system at a sampling frequency of 500 Hz. A first filtering stage was performed directly by the amplifier: a notch filter can be configured on the dedicated ECG acquisition channel, and in this study a hardware notch filter with a cut-off frequency of 50 Hz was applied to suppress the fundamental power-line interference component.

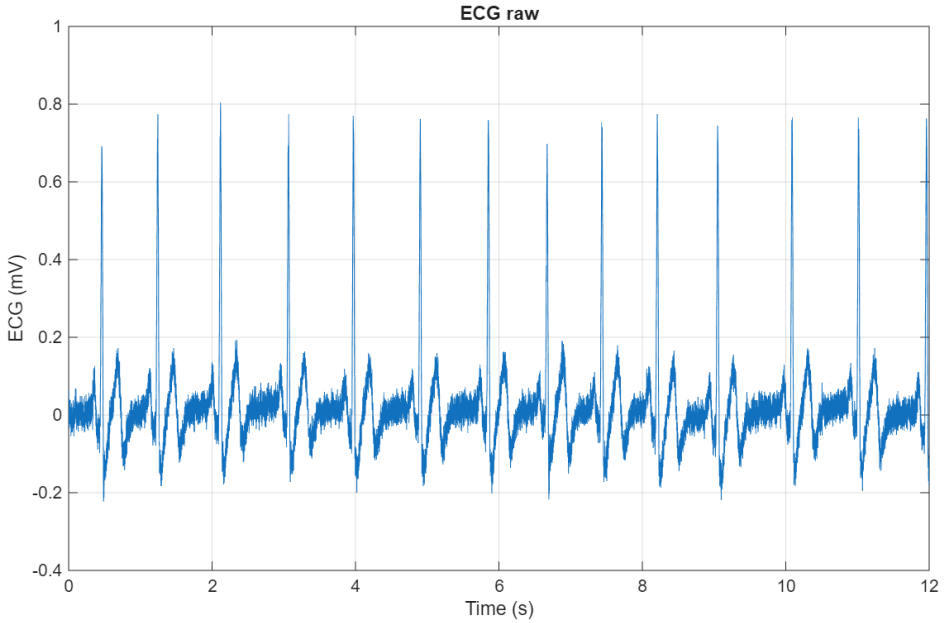

Figure 2.6: Raw ECG output from MotemaSens.

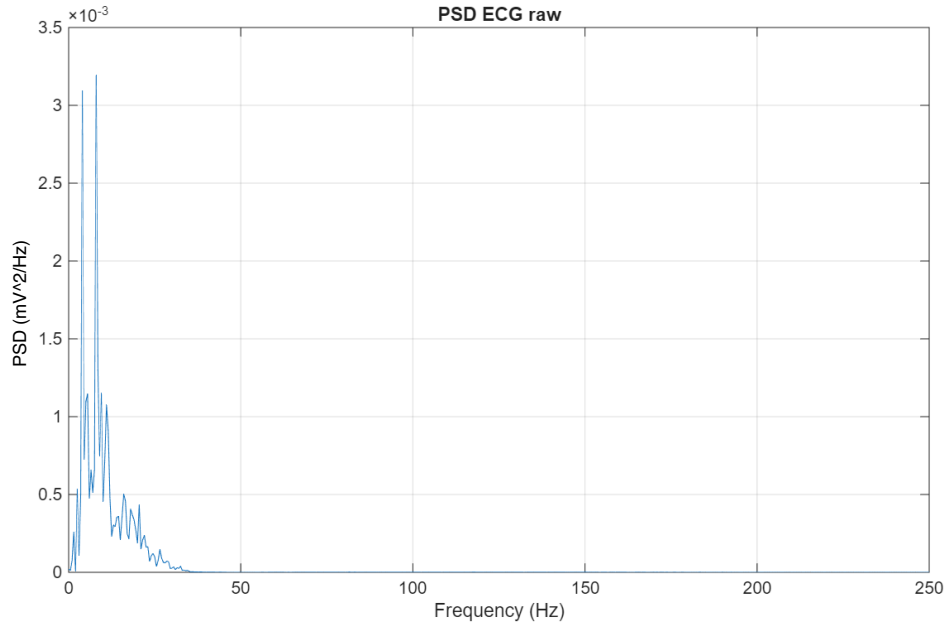

Figure 2.7: ECG frequency spectrum from MotemaSens.

The MotemaSens hardware notch filter operates at 50 Hz, effectively suppressing the fundamental power-line interference. However, in several recordings, a residual periodic interference was observed at 100 Hz, corresponding to the second harmonic of the power-line frequency. This interference may be caused by nonlinear distortion in the acquisition electronics or by external sources. Since this component is distinct from the 50 Hz fundamental frequency, it is not attenuated by the hardware filter. To address this issue, an additional recursive IIR notch filter was designed and applied in software, centred at 100 Hz with an attenuation bandwidth of 4 Hz. The filter was applied in Matlab using the causal $filter$ command.

The filter coefficients were computed analytically using the custom rico function, which receives as input the minimum attenuation parameter z, the attenuation bandwidth B, expressed in Hz, the centre frequency f_c, expressed in Hz, and the sampling interval $T = 1/f_s$. The filter is a second-order IIR structure whose transfer function in the z-domain is:

$$H(z) = \frac{1 + c_1 z^{-1} + c_2 z^{-2}}{1 - c_3 z^{-1} - c_4 z^{-2}} \tag{2.1}$$

The numerator, corresponding to the moving-average part, places a pair of complex conjugate zeros on the unit circle at f_c, producing exact attenuation at the target frequency. The denominator, corresponding to the autoregressive part, places a pair of poles just inside the unit circle at the same frequency, restoring the gain outside the notch and ensuring filter stability. The coefficients are derived as follows, with $b = \pi BT$ and $a = bz$:

$$c_1 = -2(1 - a)\cos(2\pi f_c T) \tag{2.2}$$

$$c_2 = (1 - a)^2 \tag{2.3}$$

$$c_3 = 2(1 - b)\cos(2\pi f_c T) \tag{2.4}$$

$$c_4 = -(1 - b)^2 \tag{2.5}$$

The parameter z controls the depth of attenuation at f_c: a smaller value of z deepens the notch. The parameter B controls the separation between zeros and poles and therefore the width of the attenuation band. In this implementation, $z = 0.01$ and $B = 4Hz$ were selected to obtain strong and selective attenuation at 100 Hz, while minimally affecting the surrounding spectral content.

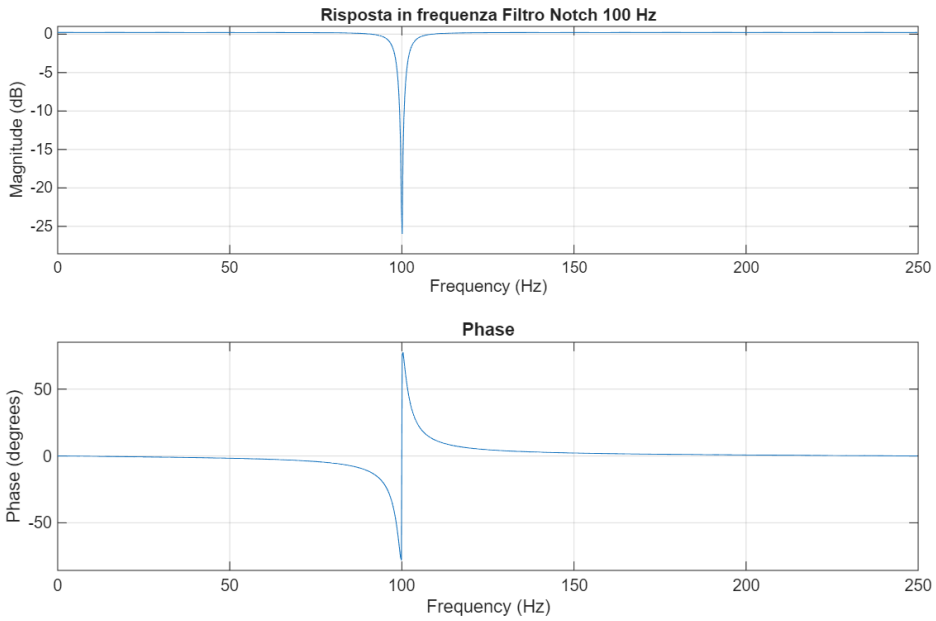

Figure 2.8: Pole-zero plot and frequency response of the 100 Hz Notch filter.

A first-order Butterworth bandpass filter was then applied. The passband ranged from 1 Hz to 30 Hz: the lower cut-off frequency removes low-frequency disturbances caused by baseline drift and motion artifacts, while the upper cut-off frequency attenuates high-frequency electrical noise and EMG-related artifacts. The decision to sacrifice part of the morphological information was deliberate and consistent with the features used to train the classifier. For stress assessment, accurate R-peak detection is the main requirement, even at the cost of losing information from other portions of the ECG waveform, such as the ST segment.

The bandpass filter was applied using Matlab's $filtfilt$ command, which performs zero-phase forward-backward filtering and removes the phase distortion introduced by IIR filters. Strictly speaking, zero-phase filtering requires the complete signal to be available and is therefore an offline operation. In the real-time classification pipeline, $filtfilt$ was applied independently to each 20-second sliding window. Since the window duration is sufficiently long compared with the filter transient response, and since a first-order Butterworth filter settles within a few samples, the edge artifacts introduced by this approach were considered negligible in practice. In addition, the first 200 samples of the initial window were explicitly discarded in the processing pipeline to reduce the filter startup transient.

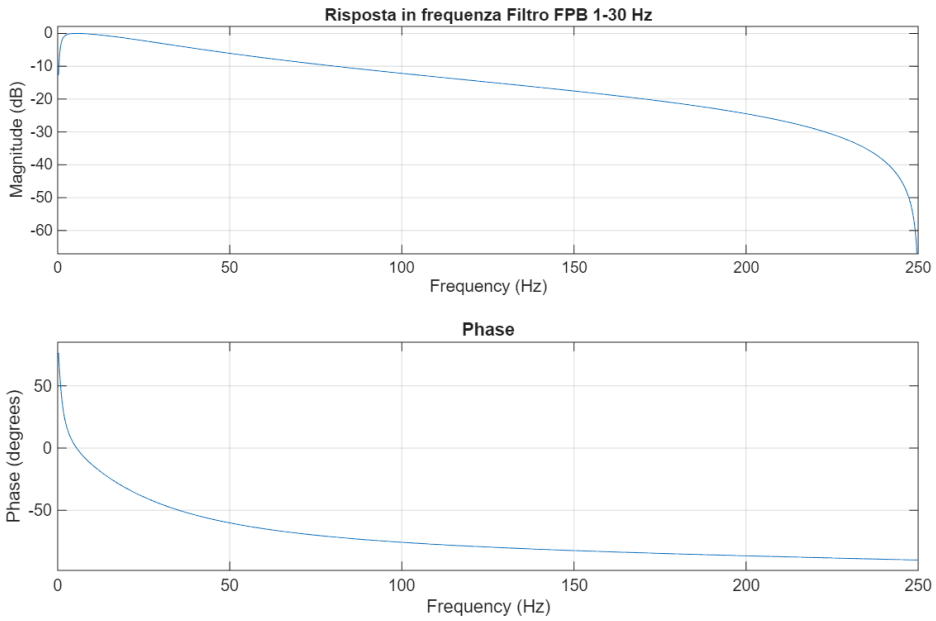

Figure 2.9: Butterworth filter: phase and frequency response.

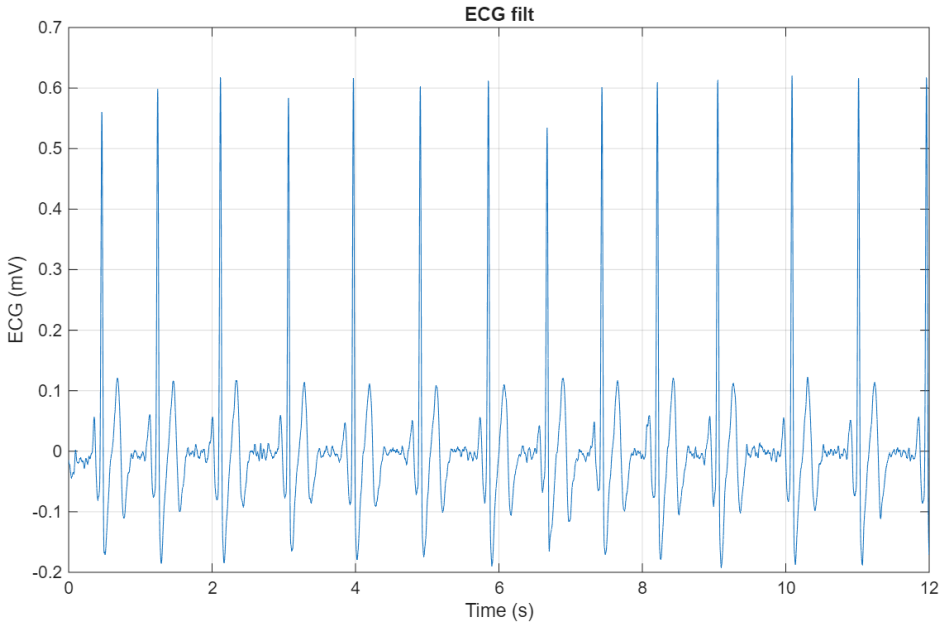

Figure 2.10: Filtered ECG signal.

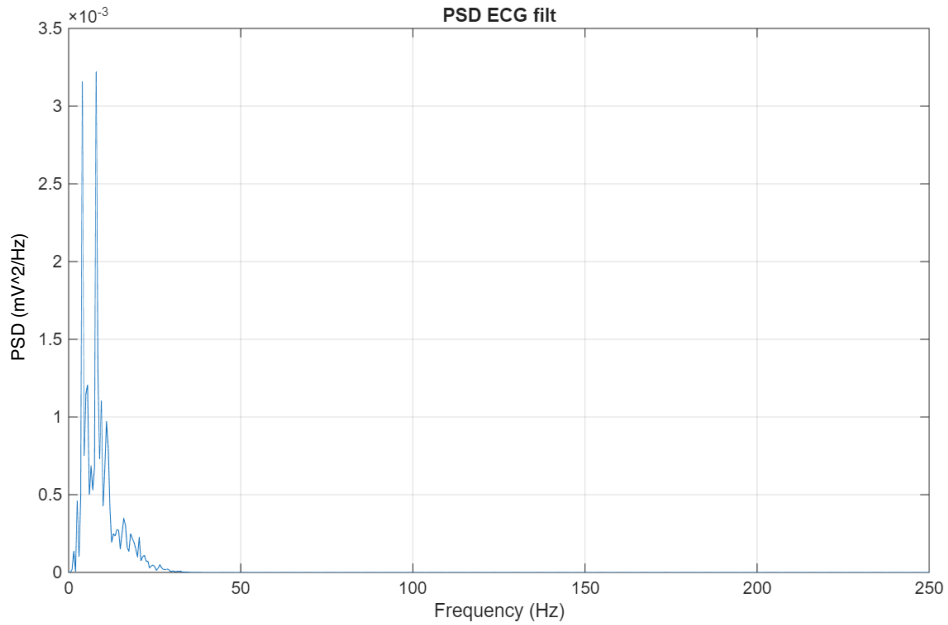

Figure 2.11: Filtered ECG frequency spectrum.

2.3.2 EDA

The EDA signal, like the ECG signal, was acquired using MotemaSens at a sampling frequency of 500 Hz and passed through the same hardware 50 Hz notch filter. As described in Section 1.4, EDA reflects changes in the electrical properties of the skin, which are modulated by sweat gland activity under the control of the SNS. It is conventionally described through two main components: the Skin Conductance Level (SCL), a slow tonic component related to general arousal, and the Skin Conductance Response (SCR), a faster phasic component associated with discrete sympathetic activations.

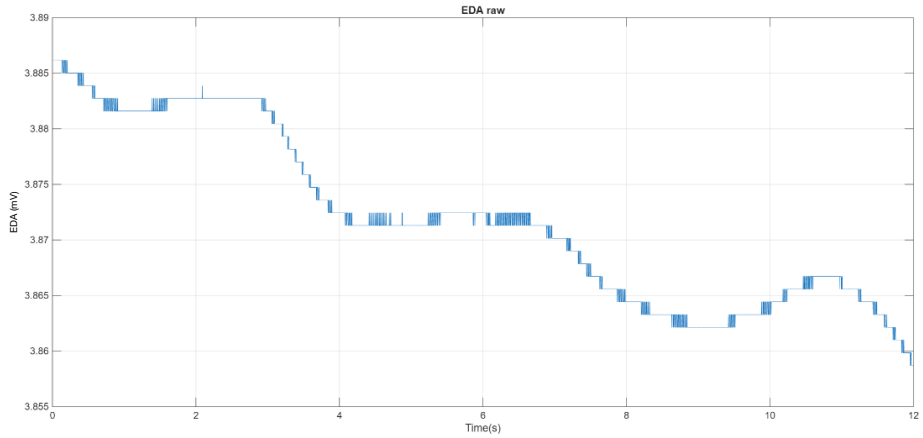

Figure 2.12: Raw EDA output from MotemaSens.

A one-dimensional non-linear median filter was applied in Matlab using the $medfilt1$ command, with a window of 0.2 seconds, corresponding to 100 samples at 500 Hz. This filter reduces noise and motion artifacts and, in particular, removes spurious transient spikes that may occur during signal acquisition, without distorting the underlying slow signal trend. After this step, a moving-average filter with a window of 0.5 seconds was applied using the *movmean* command. This linear smoothing stage further attenuates residual high-frequency fluctuations and isolates the slow component of the EDA signal that is most relevant for subsequent feature extraction.

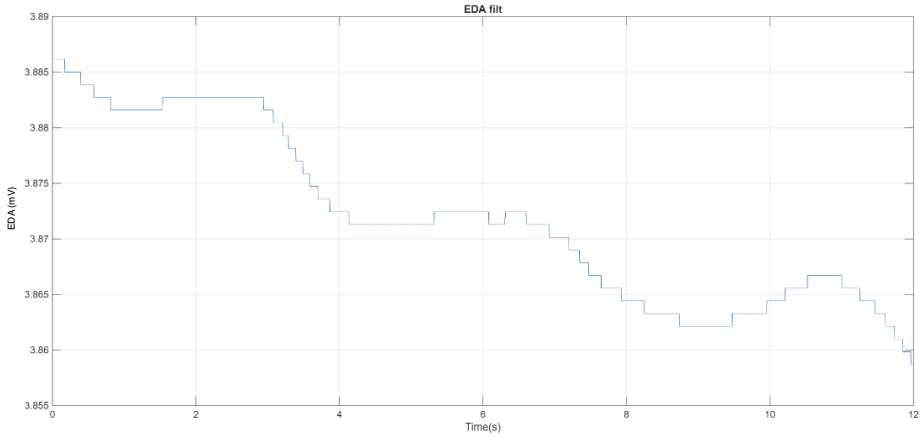

Figure 2.13: EDA signal after filtering.

2.4 Feature Extraction

After the pre-processing pipeline described in the previous section, the signals were ready for segmentation and feature extraction. This phase represents a critical step in the protocol: the selected features are those used to train the OC-SVM classifier, and their quality directly affects the reliability and generalisability of the classification output.

A key aspect in the design of this stage is the nature of the classifier itself. Since the OC-SVM is trained exclusively on data from a single class, namely the subject's resting baseline, the selected features must show good intra-subject stability under relaxed conditions, while still being sensitive enough to deviate from the baseline distribution under stress. This trade-off guided both the choice of segmentation parameters and the selection of the features.

In total, five features were extracted from the ECG signal and three from the EDA signal, producing an eight-dimensional feature vector used as input to the classifier. These features are described in detail in the following subsections.

2.4.1 Epoch Segmentation

Before feature extraction, both ECG and EDA signals were divided into partially overlapping epochs using a sliding-window approach. Each epoch had a duration of 20 seconds, corresponding to 10,000 samples at the acquisition frequency of 500 Hz. Consecutive epochs overlapped by 50%, meaning that the window advanced by 10 seconds at each step.

This choice represents a compromise between two competing requirements: a sufficiently long window to obtain stable Heart Rate Variability (HRV) estimates, which require an adequate number of RR intervals, and a sufficiently short step to allow the classifier to respond to stress-related changes with reasonable temporal resolution in the real-time pipeline. In addition, a 20-second window ensured good pre-processing quality, further improving the effectiveness of the extracted features. Tests were also performed using 30-second windows, but no significant differences emerged compared with the 20-second window; therefore, this configuration was maintained.

Segmentation was applied in the same way during both the offline training phase, where the complete 7-minute baseline recording was available, and the real-time classification phase, where epochs were extracted from a continuously updated buffer as new data became available.

One additional consideration concerns the first epoch of each acquisition: the initial 200 samples, corresponding to 0.4 seconds, were discarded before feature extraction to reduce the startup transient introduced by the filtering stage, as discussed in the signal analysis section. All subsequent epochs were processed

without this correction, since the filter had already settled by that point.

2.4.2 ECG Features

ECG-based features were extracted from the RR interval series derived from each epoch. R-peaks were detected on the filtered ECG using Matlab's *findpeaks* function, with a minimum peak height set to twice the standard deviation of the signal and a minimum inter-peak distance of 0.5 seconds, corresponding to a maximum heart rate of 120 bpm. The resulting RR intervals were then subjected to a two-stage artifact rejection procedure: intervals outside the physiological range of 0.3–2.0 seconds were discarded, and any remaining interval deviating by more than 30% from the epoch median was also removed. Five features were then computed from the cleaned RR series.

RRmean

RRmean represents the mean of the RR intervals between consecutive R-peaks within each epoch. It is directly related to mean heart rate (HR) and reflects the balance between the sympathetic and parasympathetic branches of the ANS. An increase in RRmean, corresponding to a slower heart rate, is associated with relaxation and reduced sympathetic activation, whereas a decrease indicates greater arousal or stress. It is computed as the arithmetic mean of the N valid RR intervals in the epoch:

$$RR_{mean} = \frac{1}{N} \sum_{i=1}^{N} RR_i \tag{2.6}$$

where RR_i denotes the i-th RR interval in seconds.

RRstd

RRstd is the standard deviation of the RR intervals within the epoch. It quantifies the overall variability of cardiac rhythm and reflects the combined activity of SNS and PNS. RRstd generally decreases with increasing stress, as sympathetic dominance reduces beat-to-beat variability. It is defined as:

$$RR_{std} = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (RR_i - RR_{mean})^2} \tag{2.7}$$

RMSSD

RMSSD (Root Mean Square of Successive Differences) is the square root of the mean of the squared differences between successive RR intervals. It is particularly sensitive to rapid changes in parasympathetic tone and therefore responds quickly to acute stress, which inhibits vagal activity. RMSSD usually decreases with the onset of stress and is considered one of the most reliable short-term HRV indices for autonomic assessment. It is defined as:

$$RMSSD = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N-1}(RR_{i+1} - RR_i)^2} \tag{2.8}$$

pNN20

pNN20 represents the percentage of successive RR interval pairs that differ by more than 20 ms. It is a count-based HRV index that mainly reflects parasympathetic activity and generally decreases with the onset of stress. The threshold of 20 ms, instead of the more common 50 ms used in pNN50, was selected because it captures autonomic changes more reliably in shorter windows, such as those used in this protocol, where the number of available RR intervals is limited. It is defined as:

$$pNN20 = \frac{1}{N-1}\sum_{i=1}^{N-1} I\left(|RR_{i+1} - RR_i| > 0.02\right) \cdot 100 \tag{2.9}$$

where $I(\cdot)$ is the indicator function, equal to 1 if the condition is satisfied and 0 otherwise.

CVRR

CVRR (Coefficient of Variation of RR intervals) is the ratio between RRstd and RRmean. By normalising the standard deviation by the mean RR interval, CVRR makes HRV more comparable across subjects with different baseline heart rates. This is particularly important in a one-class framework, where the model is trained on a single subject's baseline: CVRR helps distinguish the subject's intrinsic variability from stress-related changes, reducing the risk of confounding the two. It is defined as:

$$CVRR = \frac{RR_{std}}{RR_{mean}} \tag{2.10}$$

2.4.3 EDA Features

EDA-based features were extracted from the filtered and smoothed EDA signal described in Section 2.3.2. As described there, the pre-processing isolated the slow component of the signal, related to tonic electrodermal activity. Three features were computed from this signal within each epoch, capturing complementary aspects of electrodermal activity: signal variability, amplitude range, and response dynamics.

EDAstd

EDAstd is the standard deviation of the smoothed EDA signal within the epoch. It quantifies the slow variability of the signal, reflecting fluctuations in tonic sympathetic arousal. Under relaxed baseline conditions, the EDA signal is relatively stable, resulting in a low EDAstd. Stress-related increases in sympathetic activity produce larger and more irregular signal changes, increasing EDAstd. It is computed as:

$$EDA_{std} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(EDA_i - \overline{EDA})^2} \tag{2.11}$$

where EDA_i denotes the i-th sample of the smoothed signal within the epoch and $\overline{EDA}$ its mean.

EDArange

EDArange is the peak-to-peak amplitude of the smoothed EDA signal within the epoch, defined as the difference between its maximum and minimum values. It captures the overall excursion of the signal across the epoch and complements EDAstd by providing a measure of the absolute amplitude of signal changes rather than their dispersion. Large EDArange values indicate stronger electrodermal variations, whereas a small range is consistent with the stable, low-arousal state expected during baseline. It is defined as:

$$EDA_{range} = \max(EDA_i) - \min(EDA_i), \quad i = 1, \ldots, N \tag{2.12}$$

EDAder

EDAder is the mean absolute value of the first derivative of the smoothed EDA signal, computed as a discrete first-order difference scaled by the sampling frequency. It captures the rapidity of signal changes within the epoch: a high EDAder indicates fast-varying electrodermal activity, associated with phasic sympathetic

responses, whereas a low value is characteristic of the slow and stable tonic activity observed during relaxation. This feature therefore provides sensitivity to the phasic component without requiring explicit SCR detection. It is defined as:

$$EDA_{der} = \frac{f_s}{N-1} \sum_{i=1}^{N-1} |EDA_{i+1} - EDA_i| \tag{2.13}$$

where $f_s = 500Hz$ is the sampling frequency, and the scaling by f_s converts the discrete difference into an approximation of the continuous derivative, expressed in mV/s.

2.5 Model Training

The classifier model used in this study was trained using the physiological signals acquired during the baseline phase, in order to build a representation of the subject's non-stressed state.

As described in the previous section, training was based on a sliding-window approach, chosen to continuously update the estimate of the subject's state over time. The ECG and EDA signals were initially divided into temporal windows with a fixed length of 20 seconds and a 50% overlap. For each window, the ECG and EDA features were extracted, representing the physiological characteristics of the signals.

The extracted features were then combined into a single descriptive vector. Observations containing invalid values (NaN), caused by possible errors during feature extraction, were removed in order to ensure the quality of the dataset. Before implementing the training phase, the features were normalized using z-score standardization, according to the following formula:

$$z = \frac{x - \mu_B}{\sigma_B} \tag{2.14}$$

where x is the feature vector, μ_B is the mean value, and σ_B is the standard deviation.

To avoid numerical instabilities caused by very low variance values, a lower limit was imposed on the standard deviation of each feature. These limits were defined according to the order of magnitude of the variables considered and are reported in Table 2.1.

Table 2.1: Minimum values of σ for each feature.

Feature	**Minimum value of σ**
RRmean	0.01
RRstd	0.01
RMSSD	0.01
pNN20	1
CVRR	0.01
EDAstd	0.001
EDArange	0.001
EDAder	0.001

The classification model was trained using an OC-SVM, providing as input only the samples acquired during the baseline phase. This configuration allows the model to represent the subject's physiological reference condition. The biological responses associated with stress phases are highly variable and depend on several factors. As a result, the features may change considerably even when the same phase of the experimental protocol is repeated multiple times. For this reason, the use of a two-class classifier, in which the second class represents a stress condition, could have introduced critical issues both during training and during real-time predictions.

Training was implemented using RBF, chosen for its ability to model non-linear relationships in the feature space. This choice is particularly suitable for physiological signals, where the dynamics associated with the stress state are not linearly separable.

The RBF kernel measures the similarity between pairs of feature vectors based on their Euclidean distance, according to the following formula:

$$K(x_i, x_j) = \exp\left(-\gamma \cdot \|x_i - x_j\|^2\right) \tag{2.15}$$

where x_i and x_j represent the feature vectors for which the Euclidean distance is calculated. The similarity function assigns high values to similar samples and values close to zero to distant samples. This formulation allows the data to be implicitly projected into a high-dimensional space, making the non-linear separation of observations possible.

The parameter that controls the width of the similarity function is γ, defined as $\gamma = \frac{1}{2 \cdot \text{KernelScale}^2}$, and it determines the sensitivity of the model. High values may lead to overfitting, while values that are too low may lead to underfitting. The kernel scale parameter was set automatically, allowing the kernel width to adapt

to the data distribution without requiring manual selection. This choice provides a compromise between generalization ability and sensitivity to feature variations.

The outlier fraction was set to $\nu = 0.2$, corresponding to a maximum percentage of 20% of training samples that can be considered anomalous. This parameter, provided as input to the classification model, makes the model more robust to possible physiological variations or artifacts that may also be present during the baseline phase, avoiding an excessively rigid representation of the reference state. Training tests were also performed by setting $\nu = 0.1$, but during real-time prediction the model was too sensitive to the subject's physiological variations during the execution of the experimental protocol.

In parallel, the distribution of the features in the normalized space was also estimated by calculating the centroid and the covariance matrix. The latter was regularized by adding a shrinkage term equal to $\lambda = 1.2$, a value chosen after several tests in which it was increased iteratively, in order to improve numerical stability. Based on these parameters, the Mahalanobis distance (MD) was defined and used to quantify the deviation of the samples from the baseline condition, according to the following formula:

$$MD(x) = \sqrt{(x - c_B)^T \cdot \Sigma_B^{-1} \cdot (x - c_B)} \tag{2.16}$$

where x is the feature vector, c_B is the mean feature vector calculated during baseline, and Σ_B is the covariance matrix. Finally, the mean and standard deviation of this distance were calculated from the training data and then used during the real-time classification phase. The final model consists of the SVM classifier and the parameters required for feature normalization and MD calculation.

2.6 Real-Time Predictions

The classifier trained in the previous phase is used in real time to evaluate the subject's physiological state during the different phases of the experimental protocol, starting from the features extracted from the ECG and EDA signals.

For each analyzed temporal window, the feature vector is normalized using the mean and standard deviation of the baseline, estimated during training. This operation ensures consistency between the data used to train the model and the data acquired in real time.

The normalized vector is then provided as input to the OC-SVM classifier, which returns two variables:

- Label, a binary variable that represents the decision regarding whether the sample belongs to the baseline distribution;

- Score, a continuous variable that indicates how much the sample deviates from the baseline. Values lower than 0 indicate outliers; the more negative the score, the farther the sample is from the baseline.

In this context, the model is able to identify deviations from the physiological reference state, which are interpreted as stress conditions.

MD from the baseline distribution is then calculated and normalized to obtain a continuous indicator of the level of physiological deviation. In particular, the normalized distance is saturated at four standard deviations in order to avoid excessively extreme values and is then mapped to the interval $[-1, -3]$. This variable, called ScoreBB, is subsequently used as the input variable for the adaptive modulation of the BB frequency.

This measure is particularly relevant because, under conditions of high physiological deviation, the classifier score may saturate and lose discriminative ability. For this reason, BB frequency modulation is based on ScoreBB rather than on the classifier score.

2.7 BB Modulation

In adaptive mode, the BB frequency was dynamically modulated according to the subject's physiological state, represented by ScoreBB, computed during the previous real-time prediction phase. This continuous indicator, obtained as an average over multiple temporal windows, is first normalized with respect to a reference interval defined between the baseline condition and the maximum stress condition. In particular, a dimensionless variable is calculated according to the following relationship:

$$S = \frac{-z_{dec} - \|z_{baseline}\|}{\|z_{maxstress}\| - \|z_{baseline}\|} \tag{2.17}$$

where:

- z_{dec} is the continuous stress index;
- $z_{baseline} = -1$;
- $z_{maxstress} = -3$.

To ensure system stability, the resulting value is limited to the interval [0, 1] through a saturation operation, also referred to as clamping, thus avoiding values outside the expected range.

The normalized variable is then transformed using a sigmoid function, defined as:

$$\sigma = \frac{1}{1 + e^{-k \cdot (S - 0.5)}} \tag{2.18}$$

where $k = 7$ is a parameter that controls the slope of the curve.

The use of the sigmoid function introduces a non-linear mapping between the stress level and the BB frequency. In particular, the sigmoid reduces the sensitivity of the system at the extremes of the interval and increases the resolution in the central region, allowing a more gradual and physiologically plausible modulation.

The value obtained from this transformation is used to define the BB frequency through an inverse linear relationship. The frequency is constrained within the 8–12 Hz interval, belonging to the alpha band, according to the following expression:

$$df_{target} = 8 + 4 \cdot (1 - \sigma) \tag{2.19}$$

Under high-stress conditions, df_{target} tends to converge towards 8 Hz, at the lower boundary of the alpha band and close to the theta range. In the absence of stress, it tends to stabilise towards 12 Hz, within the upper alpha range. This choice was based on the assumption that lower alpha frequencies may favour reduced arousal and greater relaxation.

To ensure a gradual frequency variation and avoid abrupt transitions, a temporal smoothing mechanism based on an exponential average was introduced:

$$df_{current} = (1 - \alpha) \cdot df_{previous} + \alpha \cdot df_{target} \tag{2.20}$$

where $f_{current}$ is the current frequency, $f_{previous}$ is the frequency calculated in the previous cycle, and α controls the adaptation speed of the frequency. An asymmetric dynamic was applied, using different values of this parameter according to the trend of the stress state, as shown in Table 2.2.

Table 2.2: Values of α according to the stress state.

Condition	α	**Feedback**
$f_{target} < f_{previous}$	0.1	Increased stress, fast response
$f_{target} > f_{previous}$	0.03	Relaxation, slow response

Once the value of $df_{current}$ was obtained, the two tones used to generate the BB were defined around a fixed carrier frequency, set to $f_c = 400$ Hz. In particular, the frequency delivered to the left ear and the frequency delivered to the right ear were calculated as:

$$f_{left} = f_c - \frac{df_{current}}{2} \tag{2.21}$$

$$f_{right} = f_c + \frac{df_{current}}{2} \tag{2.22}$$

In this way, the mean frequency of the two tones remained constant and equal to the carrier frequency, while their difference was equal to the desired BB frequency. Therefore, the adaptive algorithm modified only the interaural frequency difference, while keeping the acoustic stimulus centred around the same carrier frequency.

To implement the sham phase, the frequency difference between the two audio channels is set to zero, while keeping the carrier frequency unchanged at 400 Hz. In this way, the adaptive BB component is removed while the acoustic stimulus is maintained. This makes it possible to evaluate the specific effect of BB modulation compared with simple acoustic stimulation.

Chapter 3

Results and Discussion

3.1 Sample Description

The experimental protocol was completed by 14 healthy subjects. The main characteristics of the participants are reported in Table 3.1. In particular, sex, age (25.2 ± 3.4 years old), hours of sleep during the night before the experiment(7.0 ± 1.1 hours), education level, and stimulation order were considered.

The sample included 10 male and 4 female participants, with an age range between 18 and 30 years. Regarding education level, 8 participants had a Bachelor's degree, while 6 had a high school diploma. The order of stimulation was balanced across subjects, with 7 participants following the Sham–Adaptive BB order and 7 following the Adaptive BB–Sham order.

Table 3.1: Main characteristics of the participants included in the study. EL indicates Education Level, where HS denotes High School and B denotes Bachelor's degree. SO indicates Stimulation Order, where S–ABB denotes Sham–Adaptive BB and ABB–S denotes Adaptive BB–Sham.

Participants' characteristics					
Subject	**Age**	**Sex**	**Hours of sleep**	**EL**	**SO**
S01	25	M	7 h 45 min	B	S–ABB
S02	24	M	5 h	B	ABB–S
S03	24	F	8 h	B	S–ABB
S04	22	M	6 h 30 min	HS	ABB–S
S05	28	M	6 h 30 min	B	ABB–S

Continued on next page

Table 3.1 – continued from previous page

Participants' characteristics					
Subject	**Age**	**Sex**	**Hours of sleep**	**EL**	**SO**
S06	24	M	8 h	B	S–ABB
S07	25	F	6 h	B	ABB–S
S08	28	M	9 h	B	ABB–S
S09	18	F	7 h 30 min	HS	S–ABB
S10	21	F	8 h	HS	ABB–S
S11	28	M	6 h	HS	S–ABB
S12	28	M	6 h 30 min	HS	S–ABB
S13	28	M	6 h 30 min	B	S–ABB
S14	30	M	7 h	HS	ABB–S

3.2 Verification of Cognitive Stress Induction

To verify whether the mathematical task induced a physiological response compatible with cognitive stress in the participants, the values obtained during the Baseline phase were compared with those recorded in the No stimulation condition, in which the task was performed without acoustic stimulation. For each participant and each condition, each variable was represented by the mean value calculated across all temporal windows acquired during the corresponding phase. The analysis considered all features extracted from the ECG and EDA signals, together with MD, using the non-parametric Wilcoxon signed-rank test for paired samples. This test was selected because of the limited number of participants available for analysis and because it does not require the assumption of a normal data distribution, which cannot be guaranteed in the present study. Since multiple variables were compared, the obtained p-values were corrected using the Holm–Bonferroni procedure, taking into account the number of tests performed simultaneously under the same conditions.

The comparison between Baseline and No stimulation conditions showed statistically significant differences in HR, CVRR, EDAstd, EDAder, and MD. In particular, higher values of HR ($p_{\text{Holm}} = 0.0011$), CVRR ($p_{\text{Holm}} = 0.0336$), EDAstd ($p_{\text{Holm}} = 0.0234$), EDAder ($p_{\text{Holm}} = 0.0085$), and MD ($p_{\text{Holm}} = 0.0020$) were observed in the No stimulation condition compared with Baseline. The remaining features did not show statistically significant differences after correction for multiple comparisons.

Figure 3.1 reports the comparison between Baseline and No stimulation condi-

tions for the physiological features that showed statistically significant differences after Holm–Bonferroni correction: HR, CVRR, EDAstd, and EDAder. The first row shows the cardiovascular-related variables, HR and CVRR, while the second row shows the electrodermal-related variables, EDAstd and EDAder. All four variables showed higher values in the No stimulation condition than in Baseline.

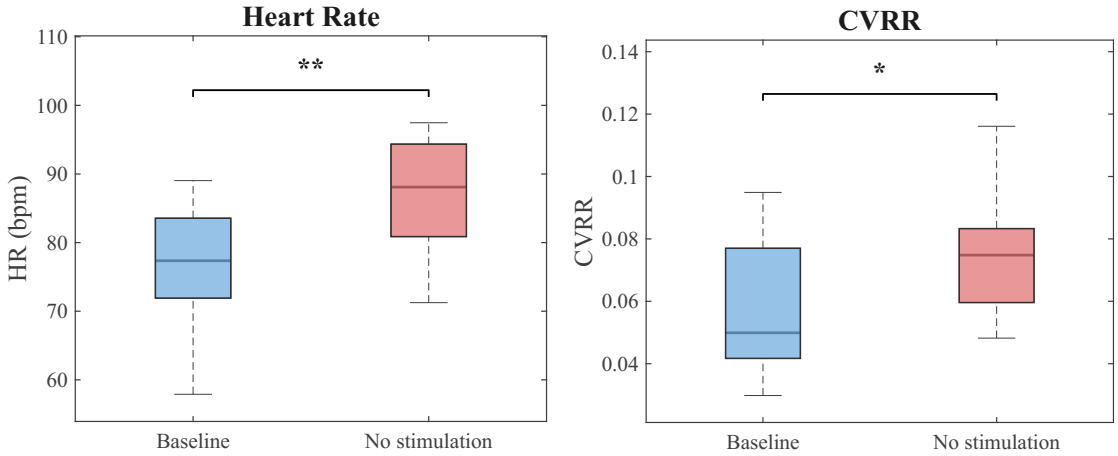

(a) HR

(b) CVRR

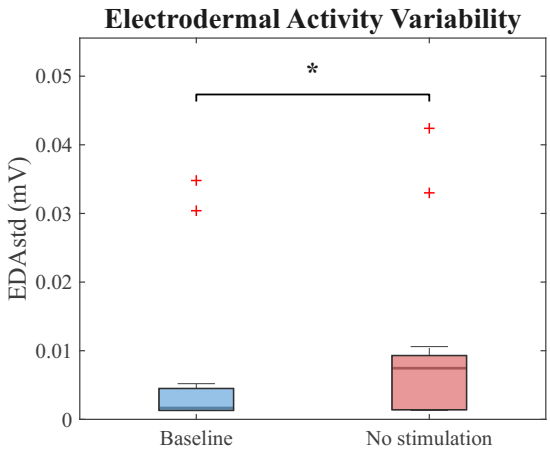

(c) EDAstd

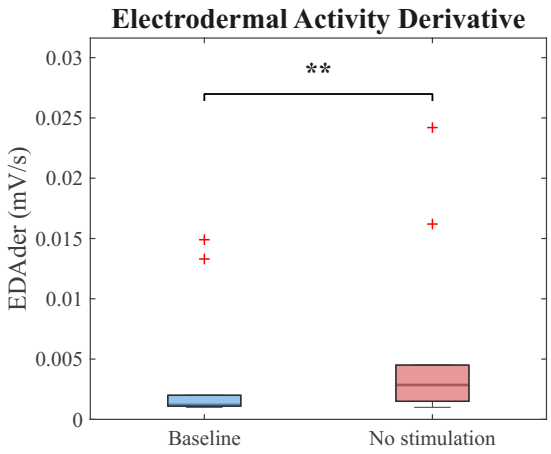

(d) EDAder

Figure 3.1: Comparison between Baseline phase and No stimulation condition for: (a) HR, expressed in bpm; (b) CVRR, dimensionless; (c) EDAstd, expressed in mV; (d) EDAder, expressed in mV/s. Differences between conditions were evaluated using the Wilcoxon signed-rank test for paired samples, with Holm–Bonferroni correction for multiple comparisons. Asterisks indicate statistical significance after correction: * $p_{\text{Holm}} < 0.05$; ** $p_{\text{Holm}} < 0.01$; *** $p_{\text{Holm}} < 0.001$.

The analysis of the physiological features is essential to evaluate whether the

mathematical task effectively induced a response compatible with cognitive stress in the participants. The significant increase in HR suggests greater cardiovascular activation during task execution, consistent with increased involvement of the SNS [7]. Among the significant variables, HR showed the clearest change, with a mean increase of 9.7 bpm from Baseline to No stimulation.

The significant increase in CVRR also indicates a change in RR interval dynamics during the task. Since CVRR is defined as the ratio between the standard deviation of RR intervals and their mean value, its increase may reflect both an increase in RR interval variability and a reduction in their mean duration, consistent with the increase in HR described above. This result therefore supports the presence of a modified cardiovascular response during the task, compatible with a condition of greater physiological activation [36].

At the same time, the significant increase in EDAstd and EDAder indicates, respectively, greater variability and a faster temporal variation of the EDA signal. These results support the presence of increased sympathetic activation during task execution [11].

MD is reported separately in Figure 3.2, since it represents a synthetic index of the deviation of the overall physiological feature vector from the individual baseline profile. The low values observed during the Baseline phase are expected because of the definition of the index, since the baseline profile represents the reference used to calculate the distance.

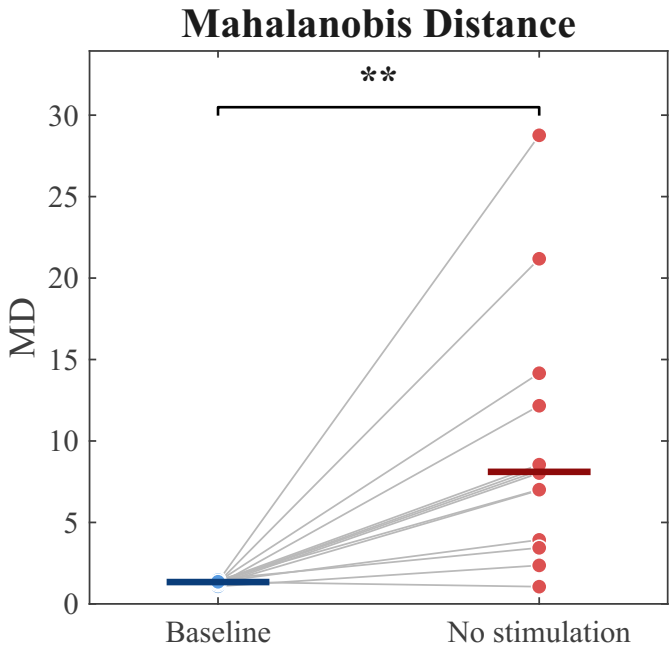

Figure 3.2: Comparison of MD between Baseline and No stimulation condition. Points represent individual participant values, while lines connect paired observations from the same participant in the two conditions; horizontal segments indicate the median of each distribution. Differences between conditions were evaluated using the Wilcoxon signed-rank test for paired samples, with Holm–Bonferroni correction for multiple comparisons. Asterisks indicate statistical significance after correction: * $p_{\text{Holm}} < 0.05$; ** $p_{\text{Holm}} < 0.01$; *** $p_{\text{Holm}} < 0.001$.

MD increased by a mean value of 8.27 from Baseline to No stimulation, further confirming that, during the task performed without stimulation, the overall vector of physiological features deviated from the individual profile identified during the Baseline phase.

Therefore, the combined increase in cardiovascular variables, electrodermal variables, and the index derived from the classification model supports the validity of the mathematical task as a tool capable of inducing cognitive stress, which can then be used to evaluate the effect of stimulation.

3.3 Effect of Auditory Stimulation Conditions

After comparing the Baseline and No stimulation conditions, the effect of adaptive binaural stimulation, together with the possible placebo effect, was evaluated during the execution of the mathematical task. For this purpose, the three experimental task conditions were compared: No stimulation, Sham, and Adaptive BB

(ABB). The No stimulation condition represents task execution in the absence of acoustic stimulation; the Sham condition corresponds to the playback of a sound at a constant frequency; the ABB condition involves the playback of BB, modulated according to the estimated physiological state of the subject.

The analysis was conducted considering all features extracted from the ECG and EDA signals, MD, and the Score produced by the classifier. The Score was analyzed only during the task conditions, since the baseline represents the reference condition used to train the classifier.

To analyze more than two data distributions for each variable, a global comparison among the No stimulation, Sham, and ABB conditions was performed using the non-parametric Friedman test for repeated measures. This test was selected because the three conditions were acquired from the same subjects and the sample size was limited. For variables that showed a significant effect of the condition, pairwise post-hoc comparisons were then performed using the Wilcoxon signed-rank test for paired samples. The p-values of the post-hoc comparisons were corrected using the Holm–Bonferroni procedure.

The Friedman test showed a significant effect of condition for HR ($p = 0.00001$), EDAder ($p = 0.01555$), MD ($p = 0.00471$), and Score ($p = 0.03243$).

Figure 3.3 reports the boxplots of the selected physiological features across the three task conditions. For HR, the post-hoc comparisons showed significantly lower values both in the Sham condition compared with No stimulation ($p_{\text{Holm}} = 0.00037$) and in the ABB condition compared with No stimulation ($p_{\text{Holm}} = 0.00037$), with mean reductions of 4.01 bpm and 4.61 bpm, respectively. The comparison between Sham and ABB did not show statistically significant differences after correction for multiple comparisons. For EDAder, although the Friedman test was significant, none of the post-hoc comparisons remained significant after Holm–Bonferroni correction. CVRR and EDAstd did not show a significant global effect of condition.

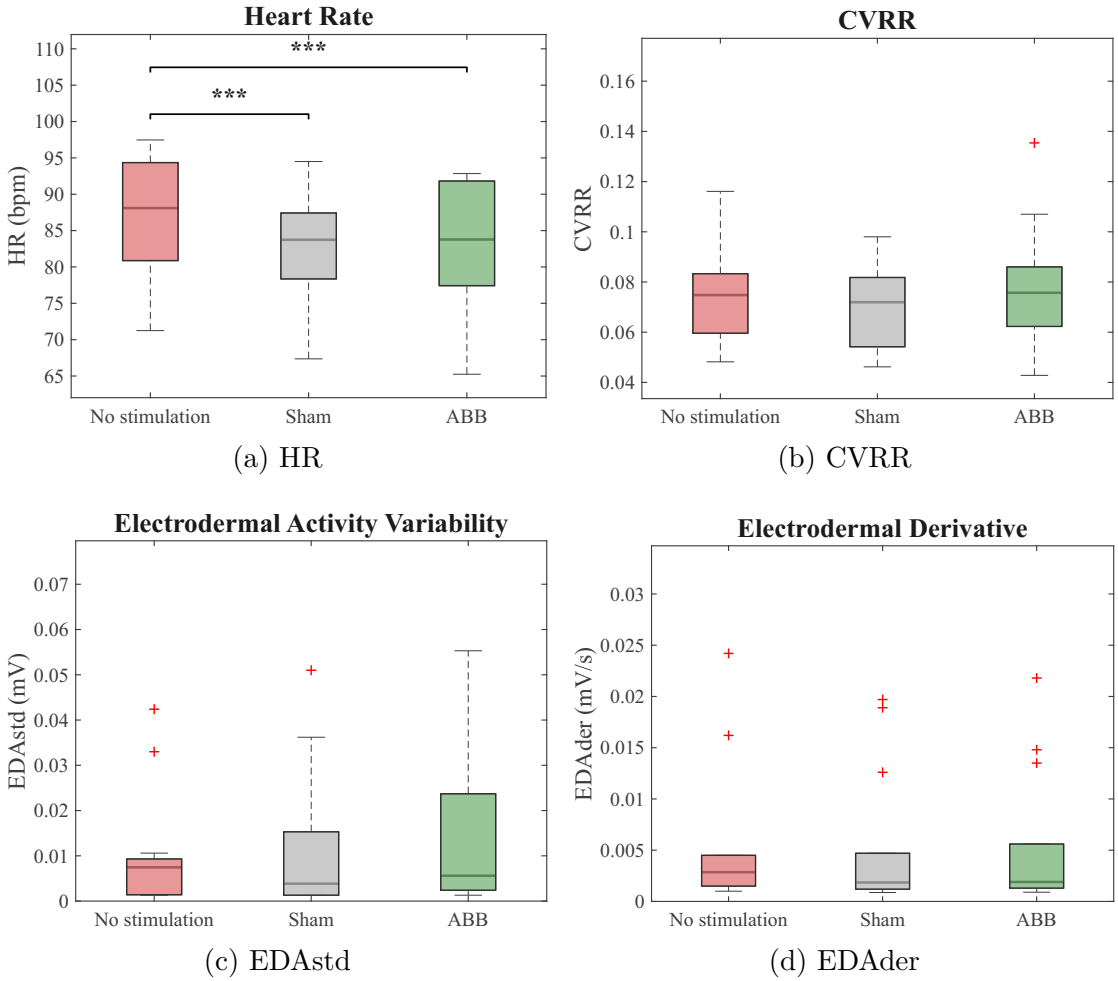

(a) HR (b) CVRR

(c) EDAstd (d) EDAder

Figure 3.3: Comparison among the three experimental task conditions for: (a) HR, expressed in bpm; (b) CVRR, dimensionless; (c) EDAstd, expressed in mV; (d) EDAder, expressed in mV/s. Global differences among conditions were evaluated using the Friedman test. Pairwise post-hoc comparisons were performed using the Wilcoxon signed-rank test for paired samples, with Holm–Bonferroni correction. Asterisks indicate significant comparisons after correction: * $p_{\text{Holm}} < 0.05$; ** $p_{\text{Holm}} < 0.01$; *** $p_{\text{Holm}} < 0.001$.

Figure 3.4a reports the boxplot of MD across the three task conditions. The post-hoc comparisons showed significantly lower values in the Sham condition compared with No stimulation ($p_{\text{Holm}} = 0.03320$) and in the ABB condition compared with No stimulation ($p_{\text{Holm}} = 0.00256$), with mean reductions of 1.65 and 4.35, respectively. The comparison between Sham and ABB was not statistically significant after correction.

Figure 3.4b reports the behaviour of the classifier Score across the three task

conditions. For this variable, the post-hoc comparison showed a significant difference between No stimulation and ABB ($p_{\text{Holm}} = 0.02563$), with a mean increase of 0.362, indicating values closer to the normality region defined during the Baseline phase. The comparison between No stimulation and Sham showed nominal significance before correction, but did not remain significant after correction for multiple comparisons. The comparison between Sham and ABB also did not show statistically significant differences after correction.

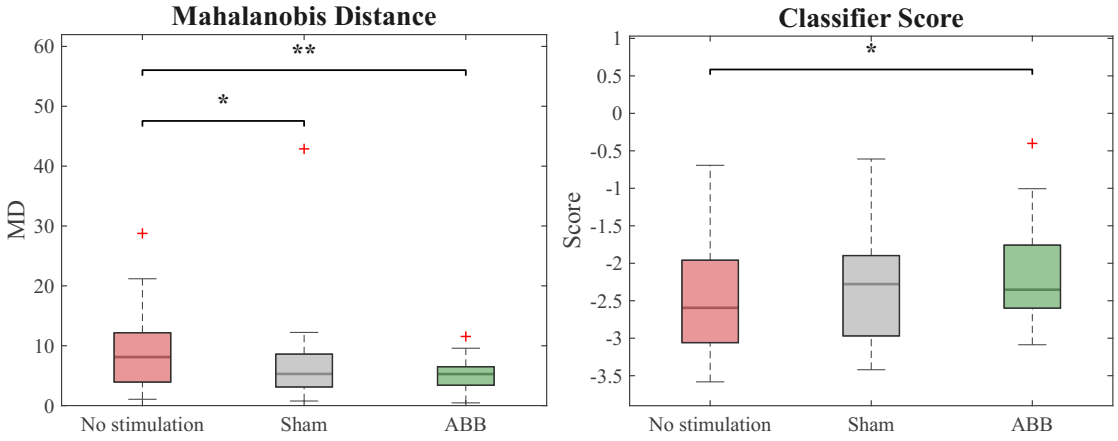

(a) MD (b) Score

Figure 3.4: Comparison among the three experimental task conditions for: (a) MD, representing the deviation of the overall physiological feature vector from the individual baseline profile; (b) Score, representing the output of the classifier. Global differences among conditions were evaluated using the Friedman test. Pairwise post-hoc comparisons were performed using the Wilcoxon signed-rank test for paired samples, with Holm–Bonferroni correction. Asterisks indicate significant comparisons after correction: * $p_{\text{Holm}} < 0.05$; ** $p_{\text{Holm}} < 0.01$; *** $p_{\text{Holm}} < 0.001$.

The comparison among the three task conditions suggests that the presence of an auditory stimulus was associated with a reduction in some physiological indices related to stress compared with the No stimulation condition [7]. In particular, HR decreased significantly ($p_{\text{Holm}} < 0.001$) in both the Sham and ABB conditions compared with No stimulation, while no significant differences emerged between Sham and ABB. This result suggests that part of the observed effect may be related to the presence of the acoustic stimulus itself, and not exclusively to the adaptive modulation of BB.

A similar interpretation can be applied to MD, which was significantly lower in both the Sham and ABB conditions compared with No stimulation. The reduction in MD in both stimulation conditions indicates that the physiological state of the subjects was closer to the baseline profile than in the condition without stimulation.

A relevant aspect is that the comparison between No stimulation and ABB was more significant ($p_{\text{Holm}} < 0.01$) than the comparison between No stimulation and Sham ($p_{\text{Holm}} < 0.05$), suggesting a possible greater effectiveness of adaptive BB stimulation. However, the absence of a significant difference between Sham and ABB indicates that, in this sample, the adaptive modulation of BB did not produce an effect clearly distinguishable from that of fixed acoustic stimulation.

The Score produced by the classifier showed a significant difference only between No stimulation and ABB after correction for multiple comparisons. This result suggests that adaptive BB stimulation may have produced a more marked reduction in the deviation estimated by the classifier compared with the No stimulation condition. However, since the comparison between Sham and ABB was not significant, this finding should be interpreted with caution. It indicates a possible specific contribution of adaptive stimulation, but does not provide sufficient evidence to clearly separate the adaptive effect from the general effect of acoustic stimulation or from a possible placebo effect.

For EDAder, the Friedman test indicated a significant global effect of condition, but none of the post-hoc comparisons remained significant after Holm–Bonferroni correction. This suggests that the dynamics of the EDA signal may vary across conditions, but also that the sample size was not sufficient to identify robust pairwise differences. CVRR and EDAstd, instead, did not show a significant global effect of condition, indicating that not all physiological features were equally sensitive to the different stimulation conditions.

Overall, compared with No stimulation, the ABB condition showed significant differences in a larger number of variables than the Sham condition. This result supports the possible benefit of adaptive binaural stimulation in reducing physiological indices associated with stress. However, none of the analyzed variables showed significant differences between Sham and ABB. Therefore, the specific effect of adaptive BB modulation is only partially supported and requires further investigation, consistently with the variability of BB effects reported in the literature [39].

3.4 Analysis of Stimulus Delivery Order

To investigate whether the order of stimulus administration could influence the physiological response of the participants, the sample of fourteen subjects was divided into two balanced subgroups of seven subjects each. In the first subgroup, referred to as S-ABB, the second mathematical task was accompanied by the sham stimulus and the third by the adaptive BB. In the second subgroup, referred to as ABB-S, the order was reversed, with the adaptive BB delivered during the second task and the sham stimulus during the third. The counterbalanced design was

introduced to help separate the effect of the stimulus from potential order effects, such as task learning or cumulative cognitive fatigue, which could otherwise affect the interpretation of the results.

As described in Section 3.3, the statistical analysis conducted on the full sample of fourteen subjects revealed a significant effect of the experimental condition for HR, EDAder, MD, and Score, identified through the Friedman test and subsequent Wilcoxon post-hoc comparisons with Holm–Bonferroni correction. It should be noted that this statistical framework was not applied to the direct comparison between the S-ABB and ABB-S subgroups. With only seven subjects per group, the sample size would be insufficient to ensure adequate statistical power, making formal statistical comparisons unreliable. Therefore, the subgroup analysis has an exploratory and descriptive purpose, aimed at visually assessing whether the differences observed in the full sample were maintained consistently regardless of the order in which the stimuli were delivered.

Figure 3.5 Figure 3.6, and Figure 3.7 report the boxplots of the three most relevant variables, HR, MD, and Score, respectively. For each variable, the three experimental conditions, No stimulation, Sham, and ABB, are shown separately for the full sample (All) and for the two subgroups defined by delivery order (S-ABB and ABB-S). This visual comparison allows an exploratory assessment of whether the effects observed in Section 3.3 were consistent across the two subgroups, or whether systematic differences related to the stimulus delivery sequence emerged.

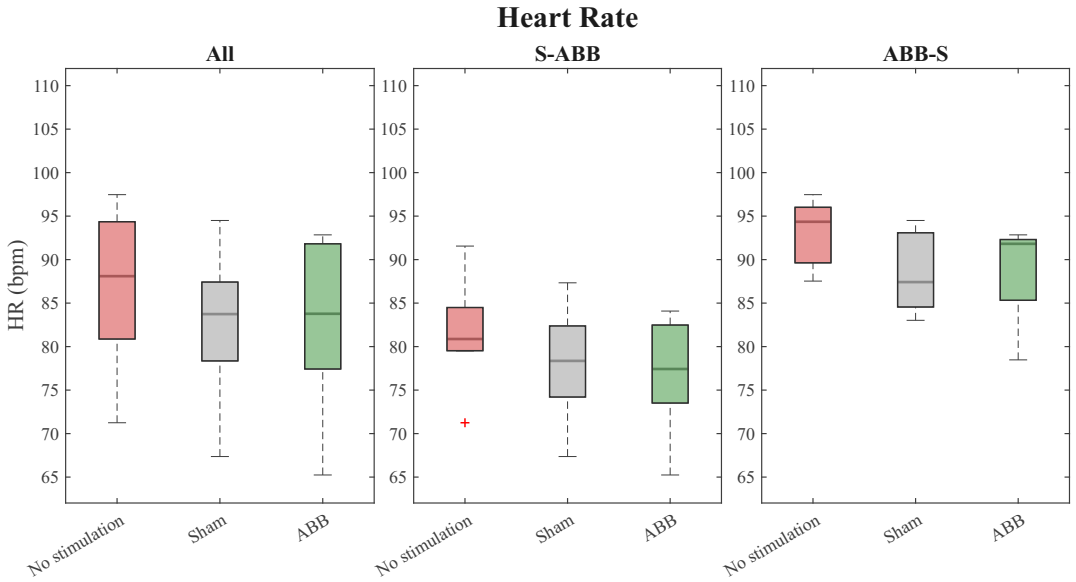

Figure 3.5: Comparison of HR, expressed in bpm across the three experimental conditions (No stimulation, Sham, ABB) for the full sample (All) and for the two subgroups defined by delivery order: S-ABB (Sham first, then ABB) and ABB-S (ABB first, then Sham).

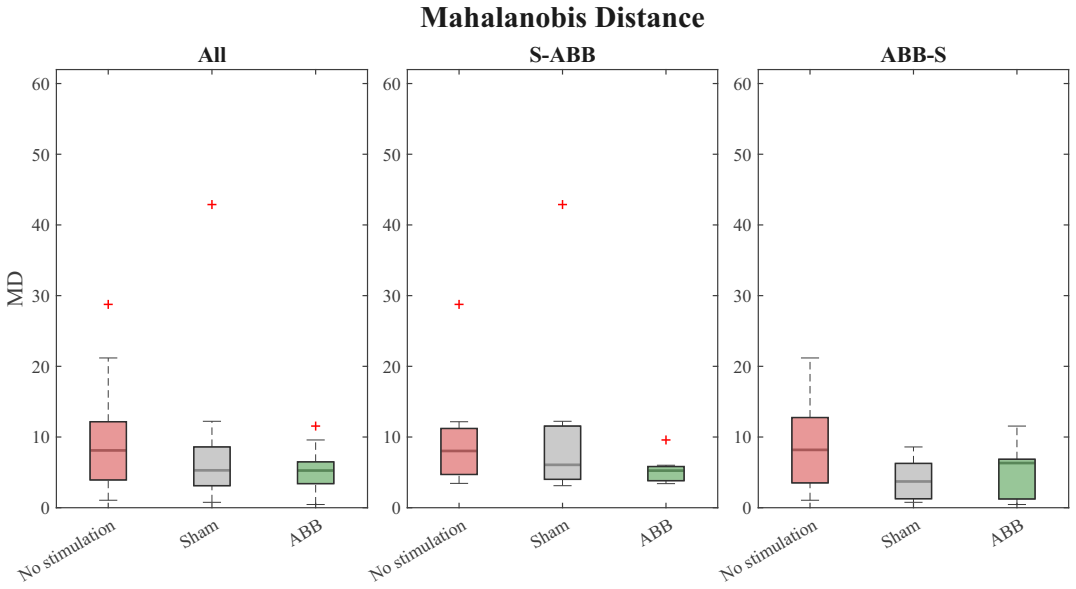

Figure 3.6: Comparison of MD, a composite index of the deviation of the physiological feature vector from the individual baseline profile, across the three experimental conditions for the full sample (All) and for the S-ABB and ABB-S subgroups.

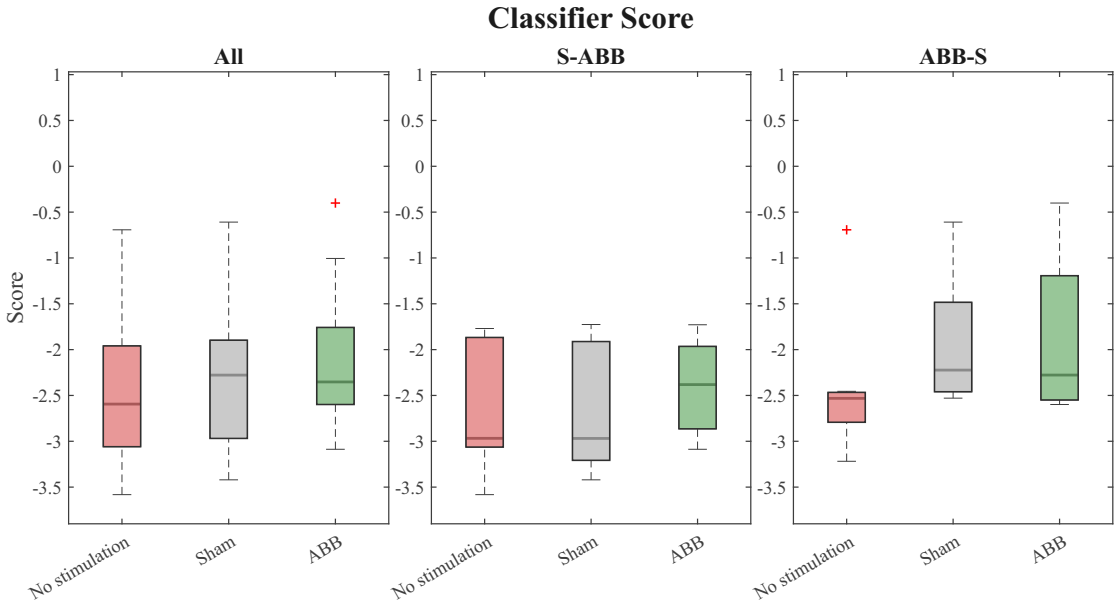

Figure 3.7: Comparison of the classifier Score across the three experimental conditions for the full sample (All) and for the S-ABB and ABB-S subgroups. The Score represents the raw output of the OC-SVM, with more negative values indicating a greater distance from the normality region defined during training.

Visually, the graph reported in Figure 3.5 does not show substantial differences as a function of delivery order. HR displays similar trends in the S-ABB and ABB-S subgroups, suggesting that this variable mainly reflects the presence of stimulation rather than its temporal position within the protocol.

The variable that appears most sensitive to the order of administration is MD, reported in Figure 3.6. In the ABB-S subgroup, in which the adaptive BB precedes the sham stimulus, MD values recorded during the Sham phase are comparable to, and in some cases lower than, those recorded during adaptive stimulation. This pattern is compatible with a possible carry-over effect: the adaptive stimulation administered during the second task may have induced a reduction in the stress-related physiological deviation that partially persisted into the subsequent session, lowering MD values even during the Sham condition. In the S-ABB subgroup, by contrast, prior exposure to the sham stimulus does not appear to produce an analogous effect. MD values decrease more clearly only when adaptive stimulation is introduced during the third task, suggesting that the observed improvement may be more closely related to the adaptive stimulus than to a generic order effect.

A qualitatively similar trend is observed for the classifier Score, reported in Figure 3.7. In the ABB-S subgroup, the values obtained during the Sham phase are comparable to those observed during adaptive stimulation, whereas in the S-

ABB subgroup the most pronounced improvement is observed only in the ABB condition. In addition, the values in the ABB-S subgroup appear globally better in both the Sham and ABB phases compared with the No stimulation condition, a pattern that is less evident in the S-ABB subgroup.

Among the variables analysed, MD appears to be the most informative for evaluating potential order effects. This is because it summarises, in a single scalar index, the deviation of the overall physiological feature vector from the individual baseline profile, while also accounting for the correlations among the ECG and EDA variables used by the classifier. Overall, this exploratory analysis suggests that the order of stimulus administration may have influenced the response observed in some variables, especially MD and Score. However, given the limited number of subjects in each subgroup, these observations should be interpreted with caution and considered as preliminary indications rather than conclusive evidence.

Chapter 4

Conclusions

This chapter reviews the most relevant methodological choices made throughout the thesis, highlighting their strengths and limitations. Possible future developments are also discussed, with the aim of improving and refining the proposed protocol.

4.1 Protocol Design

One of the most critical design choices in protocols of this type concerns the classifier. The adoption of a one-class approach, implemented through an OC-SVM classifier, proved particularly well suited to the application context considered in this work. This choice made it possible to avoid the need for a labelled two-class dataset and the associated difficulty of defining a priori thresholds to distinguish between stress and relaxation states. It is reasonable to assume, with a good degree of confidence, that the subjects were in a relaxed state during the baseline acquisition phase, since the experimental conditions were controlled and the subjects were at rest before any stress induction. Furthermore, once the mathematical tasks were administered, the recorded physiological parameters showed variations consistent with those expected during acute cognitive stress, as illustrated in Figure3.1. Therefore, under the assumption that the subjects were effectively relaxed during the initial seven-minute baseline, the classifier provided a reliable operational framework for detecting deviations from the individual resting state.

Another important choice concerned the selection of physiological signals. In many protocols similar to the present one, stress monitoring is mainly performed using ECG and EEG signals [45, 46]. Given the relative novelty of this field, the EDA signal, which is still less frequently used in this type of application, was selected to complement the ECG signal, with the aim of broadening the information available for classifying each subject's physiological state.

The ECG signal proved to be a reliable physiological signal for the real-time detection of stress. In particular, HR was the most informative feature under cognitive stress conditions. HR showed strong statistical significance both in the comparison between the Baseline and No stimulation phases and in the overall analysis across the Sham, No stimulation, and ABB conditions.

The EDA signal, on the other hand, presented some intrinsic limitations related to its spectral properties, which made it only partially compatible with the temporal constraints of the adopted protocol. The analysis windows, with a duration of 20 seconds and 50% overlap, were chosen to allow real-time adaptation of the ABB stimulus. However, this duration was not sufficient to reliably capture the phasic component of the EDA signal, which typically occurs in the 0.05–0.5 Hz range and requires longer observation windows to be adequately represented. This limitation negatively affected the statistical reliability of the features extracted from EDA. The features that provided the best results were EDAder and EDAstd, which showed good discrimination in the comparison between Baseline and No stimulation, but limited sensitivity during the second phase of the protocol, when the mathematical tasks were already underway.

The mathematical task proved effective in inducing stress in the participants. The physiological signals showed significant variations in response to its administration. In particular, HR and MD showed marked mean increases between the two phases under comparison. The magnitude of these variations in the most reliable features supports the conclusion that cognitive stress induction occurred in a consistent and reproducible manner.

4.2 Adaptive BB Effects

The results reported in Section 3.3 did not show a statistically significant advantage of ABB over the Sham condition. This may be due to a placebo effect associated with the acoustic stimulus itself, or to the progressive cognitive adaptation of the subject to the mathematical task. This adaptation may have reduced the level of induced stress over the course of the session, making any additional benefit of ABB less detectable. A further possible explanation may be related to the way in which BB were adapted during the mathematical task sessions. These aspects are discussed in Section 4.3, together with possible future improvements.

The descriptive analysis of stimulus delivery order, reported in Section 4.3, suggests that ABB stimulation may produce a partially persistent relaxation effect in the subsequent session. This pattern emerges visually for MD and classifier Score when comparing the ABB-S and S-ABB subgroups, whereas it is not observed for HR, which does not appear to vary according to the auditory stimulus delivered first. These observations remain preliminary and require confirmation on larger

samples.

4.3 Future Directions

As discussed in Section 2.7, the beat frequency difference, df, used for adaptive BB generation is mapped within the alpha band, between 8 and 12 Hz, and decreases in response to stress. This choice was based on the assumption that lower alpha frequencies may favour reduced arousal and greater relaxation.

However, the support for this specific modulation strategy in the BB literature remains limited. It is therefore important to emphasise that the direction of modulation adopted in this work represents an implementation choice that has not yet been fully validated experimentally. A further source of variability is the individual response to stimulation: not all subjects respond in the same way to the same entrainment frequency. In future work, it would therefore be useful to design an adaptive system capable of dynamically exploring the frequency space in both directions, in order to identify, for each subject, the stimulation frequency that produces the greatest relaxation response based on real-time physiological feedback.

A further consideration concerns the use of the EDA signal. As discussed above, the 20-second analysis window imposed by the real-time adaptation requirement of the ABB stimulus proved too short to adequately capture the phasic component of the EDA signal. Future protocols that include EDA should therefore adopt longer observation windows, allowing the phasic component to fully develop in response to external stimuli and remain detectable for a sufficient duration. This would make EDA-derived features more statistically informative for the study of induced stress and would allow the full potential of this signal as a complementary indicator of sympathetic nervous system activity to be exploited.

A further development could concern the mathematical task. Although it proved effective in inducing stress, its repeated administration over the course of the experimental session may progressively reduce its efficacy because of the cognitive adaptation of the subject. To mitigate this effect, future protocols involving multiple task administrations could vary the subtrahend value between tasks, rather than keeping it constant as in the present study. This could help preserve an adequate level of perceived difficulty and maintain the physiological stress response across conditions.

Finally, future recordings should be conducted on larger samples in order to increase the statistical power of the analyses and to more robustly assess the reliability of the extracted features, including in relation to the order of administration of the experimental conditions in the ABB-S and S-ABB subgroups.

Bibliography

[1] Ghasemi, F. et al. Stress and stress responses: A narrative literature review. *Journal of Health Psychology*, 2024.

[2] Charmandari, E., Tsigos, C., and Chrousos, G. Endocrinology of the stress response. *Annual Review of Physiology*, 67, 259–284, 2005.

[3] Selye, H. *The Stress of Life*. McGraw-Hill, 1956.

[4] Robinson, A. M. Let's talk about stress: History of stress research. *Review of General Psychology*, 22(3), 334–342, 2018.

[5] Carrasco, G. A., and Van de Kar, L. D. Neuroendocrine pharmacology of stress. *European Journal of Pharmacology*, 463(1), 235–272, 2003.

[6] Gu, H. F., Tang, C. K., and Yang, Y. Z. Psychological stress, immune response, and atherosclerosis. *Atherosclerosis*, 223(1), 69–77, 2012.

[7] Chu, B. et al. Physiology, Stress Reaction. *StatPearls / NCBI*, 2024.

[8] Stanfield, C. L. *Fisiologia*, IV edizione. Il sistema cardiovascolare: funzione cardiaca, pp. 360–378. EdiSES, 2017.

[9] AlGhatrif, M., and Lindsay, J. A brief review: History to understand fundamentals of electrocardiography. *Global Journal of Health Science*, 4(6), 51–56, 2012. Available from: `https://pmc.ncbi.nlm.nih.gov/articles/PMC3714093/`

[10] Hurst, J. W. Naming of the waves in the ECG, with a brief account of their genesis. *Circulation*, 98(18), 1937–1942, 1998. doi:10.1161/01.CIR.98.18.1937.

[11] Boucsein, W. *Electrodermal Activity*. Springer, 2012.

[12] Posada-Quintero, H. F., and Chon, K. H. Innovations in Electrodermal Activity Data Collection and Signal Processing: A Systematic Review. *Sensors*, 2020.

[13] Dawson, M. E., Schell, A. M., and Filion, D. L. The electrodermal system. In *Handbook of Psychophysiology*. Cambridge University Press, 2007.

[14] Greco, A. et al. Advances in Electrodermal Activity Processing. *IEEE Reviews in Biomedical Engineering*, 2016.

[15] Schmidt, P. et al. Introducing WESAD: A Multimodal Dataset for Wearable Stress and Affect Detection. In *Proceedings of ICMI*, 2018.

[16] Sun, F. T., and Morrell, M. J. Closed-loop neurostimulation: The clinical experience. *Neurotherapeutics*, 2014.

[17] Zanos, S. Closed-loop neuromodulation in physiological and translational research. *Cold Spring Harbor Perspectives in Medicine*, 2019.

[18] Hebb, A. O. et al. Creating the feedback loop: Closed-loop neurostimulation. *Neurosurgery Clinics of North America*, 2014.

[19] Mirza, K. B. et al. Closed-loop implantable therapeutic neuromodulation systems. *Frontiers in Neuroscience*, 2019.

[20] Wright, J. et al. A review of control strategies in closed-loop neuroprosthetic systems. *Frontiers in Neuroscience*, 2016.

[21] Guerrero Moreno, J. et al. Closed-loop neurostimulation for affective symptoms and disorders. *Biological Psychology*, 2021.

[22] Oster, G. Auditory beats in the brain. *Scientific American*, 1973.

[23] Schwarz, D. W. F., and Taylor, P. Human auditory steady state responses to binaural and monaural beats. *Clinical Neurophysiology*, 2005.

[24] Garcia-Argibay, M., Santed, M. A., and Reales, J. M. Efficacy of binaural auditory beats in cognition, anxiety, and pain perception: A meta-analysis. *Psychological Research*, 2019.

[25] Chaieb, L., Wilpert, E. C., Reber, T. P., and Fell, J. Auditory beat stimulation and its effects on cognition and mood states. *Frontiers in Psychiatry*, 2015.

[26] Sitaram, R., Ros, T., Stoeckel, L. et al. Closed-loop brain training: The science of neurofeedback. *Nature Reviews Neuroscience*, 2017.

[27] Can, Y. S., Arnrich, B., and Ersoy, C. Stress detection in daily life scenarios using smart phones and wearable sensors: A survey. *Journal of Biomedical Informatics*, 2019.

[28] Mitchell, T. M. *Machine Learning*. McGraw-Hill, 1997.

[29] Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, 2006.

[30] Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer, 2009.

[31] Cortes, C., and Vapnik, V. Support-vector networks. *Machine Learning*, 20(3), 1995.

[32] Vapnik, V. *Statistical Learning Theory*. Wiley, 1998.

[33] Schölkopf, B., and Smola, A. J. *Learning with Kernels*. MIT Press, 2002.

[34] Schölkopf, B., Williamson, R. C., Smola, A. J., Shawe-Taylor, J., and Platt, J. Support Vector Method for Novelty Detection. *Advances in Neural Information Processing Systems*, 12, 582–588, 1999.

[35] Manevitz, L. M., and Yousef, M. One-Class SVMs for Document Classification. *Journal of Machine Learning Research*, 2, 139–154, 2001.

[36] Kim, H. G., Cheon, E. J., Bai, D. S., Lee, Y. H., and Koo, B. H. Stress and heart rate variability: A meta-analysis and review of the literature. *Psychiatry Investigation*, 15(3), 235–245, 2018.

[37] De Witte, N. A. J., Buyck, I., and Van Daele, T. Combining Biofeedback with Stress Management Interventions: A Systematic Review of Physiological and Psychological Effects. *Applied Psychophysiology and Biofeedback*, 44, 71–82, 2019.

[38] González Ramírez, M. L., García Vázquez, J. P., Rodríguez, M. D., Padilla-López, L. A., Galindo-Aldana, G. M., and Cuevas-González, D. Wearables for Stress Management: A Scoping Review. *Healthcare*, 11(17), 2369, 2023.

[39] Ingendoh, R. M., Posny, E. S., and Heine, A. Binaural beats to entrain the brain? A systematic review of the effects of binaural beat stimulation on brain oscillatory activity, and the implications for psychological research and intervention. *PLOS ONE*, 18(5), e0286023, 2023.

[40] OT Bioelettronica S.r.l. *MotemaSens – Getting Started and Firmware Description*, Version 1.1. Internal technical document provided by the manufacturer. OT Bioelettronica S.r.l., Turin, Italy, 2024.

[41] OT Bioelettronica S.r.l. *MotemaSens – USB/Bluetooth Communication Protocol*, Version 1.0. Internal technical document provided by the manufacturer. OT Bioelettronica S.r.l., Turin, Italy, 2024.

[42] Seeed Studio. Grove GSR Sensor. Available from: https://wiki.seeedstudio.com/Grove-GSR_Sensor/

[43] Thoma, M. V., La Marca, R., Brönnimann, R., Finkel, L., Ehlert, U., and Nater, U. M. The effect of music on the human stress response. *PLOS ONE*, 8(8), e70156, 2013. doi:10.1371/journal.pone.0070156.

[44] Bella musica rilassante per dormire per alleviare lo stress • Calma la mente, medita (volando). YouTube. Available from: https://www.youtube.com/watch?v=3nM8TLXKjwc. Accessed during experimental data collection.

[45] Battù, G., and Lupo, L. "Short and Long term Investigation of Brainwave Entrainment induced by Binaural Beats: A Multivariate EEG Analysis and In-Depth Characterization of Alpha Rhythm Modulation." Tesi di Laurea Magistrale, Politecnico di Torino, A.A. 2024–2025.

[46] Capelli, R., and Donati, L. "Alpha Binaural Beats and Brain Entrainment: A Machine Learning-Based Assessment of Short-Term Effects." Tesi di Laurea Magistrale, Politecnico di Torino, A.A. 2024–2025.