



**Politecnico
di Torino**



Politecnico di Torino

Master's Degree in Biomedical Engineering

A.a. 2025/2026

Graduation Session July 2026

Pupil-com: a portable pupil-based BCI for autonomous communication in Completely Locked-In State

Supervisors:

Silvestro Roatta
Rosita Rabbito

Candidate:

Giada Bernabè

*Per chi non ha una voce,
e per chi ha cura di loro.*

Abstract

Amyotrophic lateral sclerosis (ALS) is a rare neurodegenerative disease of the motor neuron system, characterised by progressive impairment of upper and lower motor neurons while cognitive functions often remain intact. In advanced stages, patients may enter a Locked-In State (LIS), marked by near-total paralysis requiring assisted feeding, ventilation, and alternative communication systems based on residual movements, commonly eye motility. Still, as the patient transitions to the completely-locked-in state (CLIS), with total absence of skeletal muscle control, communication through conventional assistive devices, like eye-trackers, becomes ineffective. Research has shown that the autonomic nervous system is largely spared in ALS, including the pupillary accommodative response (PAR), a physiological reflex that adjusts the pupil size according to changes in gaze depth. Importantly, this autonomic response can be voluntarily elicited, and offers a promising communication channel. Although previous studies have demonstrated the feasibility of pupillometry-based communication systems, the prototypes that had been developed lacked the robustness, portability and usability required for real-time communication in clinical and home-care settings.

This work aimed to develop a fully functional, user-friendly brain-computer interface (*Pupil-com*) that enables autonomous use by ALS patients, also in CLIS, and their caregivers, without requiring technical assistance. The system combines a commercial eye-tracking device with a custom Python-based software platform capable of real-time pupillometry, pupillary state classification, data management and storage, and graphical user interface control. Running on a compact mini-PC and compatible with standard external displays, *Pupil-com* offers a flexible solution suitable for both clinical and domestic environments. Voluntary PARs elicited through fixation on a near visual target are translated into user commands. Through a scanning interface, users can select options, type on a keyboard with predictive suggestions, and play a simple game. Additional functionalities are assigned to longer intentional PAR events, such as pausing or switching modes. Manually adjustable parameters, including classification sensitivity and scanning speed, allow tailoring the applications to the user's physiological responsiveness.

Pupil-com was experimentally validated in a cohort of healthy subjects to assess system performance, usability and user learning capabilities. Participants reviewed the user experience and provided suggestions for improvement via a questionnaire. Quantitative results demonstrated a 95% accuracy of the system, while qualitative feedback reported a highly positive experience, confirming its ease of use. Preliminary testing was also conducted on a motor neuron disease patient in an almost Locked-In State to evaluate feasibility of using *Pupil-com* as a BCI. The validation protocol consisted of ten simple yes/no questions with known answers, allowing assessment of the system's reliability into conveying patient's intention.

Pupil-com now offers a robust and stable stand-alone brain-computer interface that can support autonomous communication with minimal caregiver assistance. The results highlight its potential for deployment in both research and real-world settings.

Table of Contents

Abstract	II
List of Tables	V
List of Figures	VI
Introduction	1
1 Background	2
1.1 Motor Neuron System: physiology and possible pathologies	2
1.2 Amyotrophic Lateral Sclerosis (ALS)	3
1.2.1 Clinical presentation and diagnosis	3
1.2.2 Aetiology	4
1.2.3 Epidemiology	5
1.2.4 Disease progression and prognosis	5
1.3 Locked-In State	6
1.4 Transition to Completely-Locked-In State	7
1.5 Assistive technologies for LIS and CLIS patients	8
1.5.1 Low-tech AAC and Eye-Tracking Systems	8
1.5.2 Brain-Machine and Brain-Computer Interfaces	9
1.5.3 Pupillometry as a novel approach to AAC for ALS	10
1.6 Pupillary Accommodative Response (PAR) as an HCI modality	10
1.6.1 Anatomy and physiology of the eye	11
1.6.2 PAR	11
1.6.3 PAR in HCIs	13
1.6.4 Evolution of acquisition setups and interfaces	14
1.6.5 Challenges: robustness and reliability for pupillometry	15
1.7 Aim of the study	16
2 Materials	17
2.1 System Architecture	17
2.2 Hardware	18
2.2.1 Acquisition module	19
2.2.2 Computational unit	22
2.3 Software architecture	22
2.3.1 Acquisition module software	23
2.3.2 Custom software overview	25
2.3.3 Data ingestion	26
2.3.4 Signal preprocessing	26
2.3.5 PAR detection algorithm	27

2.3.6	Graphical User Interface	30
3	Methods	40
3.1	Iterative developer testing	40
3.2	Validation protocol	41
3.3	First trial involving an ALS-LIS Patient	42
3.4	Data management and analysis	44
3.4.1	Manual dataset organisation	44
3.4.2	Feature extraction	44
3.4.3	Survey processing	45
3.4.4	Statistical analysis	45
4	Results and discussion	47
4.1	System evolution and robustness	47
4.2	Impact of user experience	49
4.3	Tuning and interface usability	50
4.4	User feedback	53
4.5	Testing on an ALS-LIS patient	55
5	Conclusions and future directions	58
	Appendix A: Eye-Trackers Comparison	61
	Bibliography	65

List of Tables

2.1	Comparative evaluation matrix of commercial eye-tracking systems considered for the <i>Pupil-com</i> platform configuration.	21
2.2	Technical specifications of the MaxMini B8A Pro host computing device.	22
2.3	Summary of configurable GUI parameters across different application tabs, detailing their default values and system impact.	38
3.1	Questions asked during the trial with one ALS-LIS patient	43
3.2	Summary of subject participation and identification codes across the pilot and validation studies.	44
3.3	Summary of experimental metrics evaluated across the <i>Pupil-com</i> applications. .	46
4.1	Summary of subjective usability survey results.	54
1	Comparison of Eye-Tracking Devices (Part 1)	61
2	Comparison of Eye-Tracking Devices (Part 2)	62
3	Comparison of Eye-Tracking Devices (Part 3)	63
4	Comparison of Eye-Tracking Devices (Part 4)	64

List of Figures

1.1	Motor Neuron System, UMN and LMN distinction. Adapted from [3]	3
1.2	Eye anatomy, cross section. Reprinted from [33]	12
1.3	Pupillary Accommodative Response, graphical representation. Adapted from [34][35]	12
1.4	Qualitative relationship between pupil size and target distance. Reprinted from [41].	14
2.1	Schematic representation of <i>Pupil-com</i> 's architecture during experimental validation. Image generated utilising Gemini AI based on the author's prompts	18
2.2	<i>Pupil-com</i> 's hardware components	18
2.3	Pupil Labs's Pupil Core eye-tracking device, detailing: (a) world camera, (b) and (c) IR eye cameras, (d) USB cable. Image adapted from [46]	19
2.4	Gazepoint's GP3 eye-tracking device, detailing: (a) IR LED-array, (b) embedded IR camera, (c) tripod mount. Image adapted from [47]	20
2.5	MaxMini B8A Pro.	22
2.6	Pupil Capture interface, detailing: (a) world camera, (b) eye 1 camera, (c) live pupil segmentation. Image acquired by the experimenter	24
2.7	Gazepoint Control interface, detailing: (a) full camera view, (b) and (c) close up of each eye, (d) distance check bar. Image acquired by the experimenter.	24
2.8	Internal data processing pipeline	25
2.9	Signal lost error dialog	27
2.10	State machine diagram of ConstrictionMonitor .	28
2.11	PAR detection and type update over time. Postprocessed from experimental acquisitions.	29
2.12	T_{exit} recovery after double crossing T_{SMA} . Postprocessed from experimental acquisitions.	29
2.13	Startup screen to set the device in use and the subject's name.	30
2.14	Main Menu initialisation screen.	30
2.15	Main Menu while scanning the options by highlighting the corresponding button.	31
2.16	Result screen after completing one "Training" session.	32
2.17	Scanning interface highlighting one button at a time	33
2.18	"Yes or No?" state machine loop.	33
2.19	"Keyboard" layout, row scanning active.	33
2.20	"Keyboard" state machine diagram	35
2.21	"Shuttle game" layout.	36
2.22	"Shuttle game" state machine	37
2.23	Settings dialog in the "General" tab	39
4.1	Example of missed PAR: the subject, due to teary eyes, had to close their eyelids for longer than one second, resulting in signal loss. When a new constriction was requested, they were not ready to respond. Close up of a <i>Training</i> acquisition	48
4.2	Inter-study comparison of true algorithm detection rates distribution.	48

4.3	Intra-group average reaction times comparison. Group A (returning) subjects exhibited significantly faster responses than Group B (naive).	49
4.4	<i>Shuttle Game</i> best score distribution for Group A (returning) than Group B (naive) subjects.	50
4.5	Importance of threshold multiplier setting: in the plot above, an example of too low T_{SMA} multiplier (0.85), that caused one Type 1 and one Type 2 misses; in the plot below, the optimal value for this subject (0.9).	51
4.6	Optimal threshold comparison: <i>Training</i> vs <i>Shuttle Game</i>	52
4.7	Accidental triggers during <i>Yes or No?</i> pause attempt: exit fail vs exit success distribution.	52
4.8	Patient <i>Training</i> plots, relative to first trial (above) and second trial (below) . .	56
4.9	Patient <i>Shuttle Game</i> trial, 15–45 s close-up. Each red marker represents a successful spaceship launch. The vertical red line represents a successful collision with a planet.	57

Introduction

This thesis addresses the clinical need for an accessible and portable device to restore communication with completely paralysed individuals affected by Amyotrophic Lateral Sclerosis or similar motor neurodegenerative diseases. Since traditional eye-tracking-based assistive devices fail due to loss of oculomotor control, this project proposes an alternative brain-computer interface (BCI), *Pupil-com* based on the Pupillary Accommodative Response (PAR), an autonomic nervous system reflex that can be voluntarily modulated. To provide a thorough understanding of the problem and introduce the proposed solution, the present manuscript is outlined as follows:

- “**Chapter 1: Background**” provides an accurate overview of the target patient population, the pathophysiology, and the disease’s complications, with a focus on communication issues as the pathology progresses. A breakdown of the solutions implemented in the past decades follows, highlighting strengths and weaknesses of both commercial and research-oriented devices. Lastly, attention is focused on the innovative pupil-based communication channel proposed by Professor S. Roatta’s research group: its working principle, advantages over existing assistive communication devices, and the technical challenges it aims to overcome.
- “**Chapter 2: Materials**” reports the core developmental phase of this work, detailing the hardware choice, the definitive setup organisation, and the Python software that the interface runs on. In particular, the chapter describes the integration of a commercial eye tracker into the custom software, real-time data handling and processing, and the algorithm for classifying voluntary pupillary constrictions. Then, the various interface applications are presented through detailed state machine diagrams and operational descriptions, offering an overview of all existing functionalities and interaction possibilities, highlighting their intuitiveness and simplicity of use, and demonstrating user-friendliness.
- “**Chapter 3: Methods**” firstly presents the methodology that guided the developer towards the final framework definition; it describes in detail the protocols conducted on the healthy subjects who took part in a preliminary Pilot Study and in the official Validation Study, with a focus on the evaluation metrics extracted from each trial; then, the chapter explains the methodologies followed during the first real-world experiment, conducted on a paralysed patient with Motor Neuron Disease. Finally, it describes the user experience questionnaire each participant completed, and the descriptive and inferential statistical analysis conducted on the quantitative and qualitative results obtained during the experimental phases.
- “**Chapter 4: Results and discussion**” displays and speculates on the results obtained during the validation phase. In particular, the system’s performance metrics across all experimental phases and the effect of training on returning subjects are evaluated, along with a discussion of post-acquisition feedback.
- “**Chapter 5: Conclusions and future directions**” acknowledges the achievements of this work, highlighting its strengths and proposing solutions for improvements. Future perspectives include new implementations to enhance user experience and automate parameter settings.

Chapter 1

Background

This chapter describes the clinical and technological foundations of this work. After an introduction to the pathologies affecting the target population, the communication issue is thoroughly addressed and explored. An analysis of currently available assistive devices identifies the bottlenecks encountered when ALS patients enter a completely locked-in state, outlining the problem and how this project aims to resolve it. The phenomenon of PAR is described from a physiological and operational perspective; the chapter finally defines the technical requirements for establishing communication through the pupil and lays the foundation for the device's design and development choices.

1.1 Motor Neuron System: physiology and possible pathologies

The motor neuron system (Figure 1.1) is composed of upper (or primary) and lower (or secondary) motor neurons (UMN and LMN, respectively): the former are located in the primary motor cortex of the brain, and their axons, comprising the corticobulbar tract, project to the brainstem and the spinal cord; the latter are present in the motor nuclei of the brainstem and the grey matter of the spinal cord, and their axons represent the connection with muscles of the head region and in the trunk and limbs, respectively. Any lesion to motor neurons can result in a loss of motor function and movement capacity, including control of respiration, mastication, and swallowing. Multiple neurodegenerative pathologies are known for affecting UMN and/or LMN and causing their failure: these are collectively classified as motor neuron diseases/disorders (MND) and are categorised depending on the system that is mainly affected. Four main types have been identified: the Primary Lateral Sclerosis (PLS), concerning the upper motor neurons only, showing signs of spasticity and hyperreflexia, is the least diffused (2% of MND cases); the Progressive Muscular Atrophy (PMA), on the contrary, only affects lower motor neurons, resulting in muscle fatigue, weakness and fasciculations – i.e. twitching — and involving around the 5% of patients; the Progressive Bulbar Palsy (PBP) affects the muscles of the bulbar region, is characterised by dysphagia, dysarthria, tongue atrophy and fasciculation; last, the Amyotrophic Lateral Sclerosis: among the MNDs, ALS is the most frequent to arise in adults, involving the 80–90% of the patients. This pathology involves both upper and lower motor neurons, leading progressively to their dysfunction and impairment, with consequent loss of skeletal muscle control, inevitably followed by paralysis and atrophy of the muscular tissue itself. This gradual deterioration of the motor system will initially mainly affect voluntary movements and, over time, progressively involve the respiratory and facial muscles responsible for speech, mastication, and swallowing.

These impairments lead to the patient's death unless they opt for life support treatments, such as mechanical ventilation and artificial feeding [1][2].




Figure 1.1: Motor Neuron System, UMN and LMN distinction. Adapted from [3]

1.2 Amyotrophic Lateral Sclerosis (ALS)

The first to identify and formally define the disease was the French physician Jean-Martin Charcot in the 1870s, but for the following decades, almost no research was conducted about it, until 1939, when ALS gained media attention and became commonly known as “Lou Gehrig’s disease”, after the famous American baseball player’s career terminated because of the onset of the illness. His condition managed to raise awareness in the public, and studies were conducted about possible medication and treatment, risk factors, hereditary forms and gene mutations involved in the onset of ALS. Lambert first proposed neurological criteria in 1969 to recognise the disease, which included assessment of sensory and motor nerve function via EMG. In the 1990s, the topic became of major interest in the field of neurodegenerative diseases, and many publications in those years laid a strong foundation for current knowledge and ongoing research.

1.2.1 Clinical presentation and diagnosis

The absence of specific biomarkers unique to the disease underscores the need for a clinical diagnosis, which is typically confirmed only 9–24 months after symptoms appear. While it is generally easier to recognise the signs of a bulbar onset, characterised by dysarthria and dysphagia, as already mentioned, it is more complex to diagnose ALS with a spinal onset, as its

main symptoms, limb weakness, are common to various other types of neuromuscular disease. Other onset symptoms associated with these include respiratory insufficiency (which generally arises only in later stages), axial weakness, limb fasciculation and cramps, emotional lability, and weight loss. Moreover, there's substantial evidence of cognitive decay related to ALS onset, with changes observed in 50% of patients and 10–15% diagnosed with frontotemporal dementia (FTD). The diagnostic criteria that were first proposed in 1969 were reviewed in 1990 and published as the El Escorial criteria, which operate by dividing the body into four susceptible regions (bulbar, cervical, thoracic and lumbo-sacral) and assessing both UMN and LMN conditions in each of them with EMG- and neuroimaging-based tests; the outcome defines the degree of certainty for a diagnosis of ALS – ranging between definite (UMN and LMN involved in 3+ body regions), probable (in 2+ body regions), possible (UMN or LMN in 2+ body region) and suspected (only one body region involved). These criteria, though, were too stringent and were revised multiple times: today, the Gold Coast criteria defined in 2019 are universally used as a reference and significantly simplify the disease identification, as dysfunction of both UMN and LMN, in the same body region, or LMN only, in at least 2 body regions, is considered enough to determine ALS presence. To detect anomalies, tests mainly rely on EMG techniques, aimed at recognising changes in conduction velocity, number of motor unit potentials and their neurogenic, and fasciculation potentials.

The following techniques are still not standardised, as they, alone, are not enough for diagnosis, but provide additional information that might be useful: Transcranial Magnetic Stimulation (TMS) is a non-invasive technique commonly used to assess muscle innervation and correct functioning; studies on electroencephalography (EEG) have recently shown that ALS onset is correlated to an increased cortical connectivity of frontal and motor regions of the brain; magnetic resonance imaging (MRI) is usually recommended for excluding other diseases, but anomalies in the corticobulbar region and in the diffusion tensor tractography related to ALS were recently found; ultrasound represents a good choice for assessing fasciculation potentials and for monitoring muscle thickness [4].

1.2.2 Aetiology

Although many factors, both genetic and environmental, are possible causes of the onset of ALS, this disease's aetiology is still not completely clear. Among the risk variables connected to lifestyle and environment that have been found, cigarette smoking is considered to have a key role; of major importance are also exposure to agricultural chemicals, heavy metals such as lead and mercury, neurotoxins and low-frequency (3–3000 Hz) electromagnetic fields, and even sports with a high incidence of concussion [5].

Research conducted in 2019 found shared polygenic risk and causal inferences in ALS. From a socio-economic perspective, while results are often conflicting, some research suggests a correlation between a higher education level and a lower chance of developing the disease: in this context, educational attainment probably reflects the subject's occupation, therefore their exposure to higher risk environments, rather than being an indicator of their innate cognitive abilities, the role of which in ALS aetiology remains hypothetical.

Light and leisure-time physical activity throughout life was found to be an important factor in reducing the risks as well, while moderate to high intensity sports showed a correlation with ALS onset: this may suggest that exercise is good for preventing muscle atrophy, but practising it at higher volumes probably releases excessive free radicals and/or glutamate excitotoxicity. Another important risk factor is abnormal lipid levels in the blood: it is well known that high cholesterol is common to many other neurological diseases, although the mechanisms are still not fully understood [6].

According to the most recent studies, around 120 genes are implicated in ALS development; all are involved in abnormal protein accumulation in neural cells. The gene C9ORF72, which is normally involved in membrane trafficking and autophagy, plays a crucial role in ALS, where it harbours anomalous repeats of a non-coding hexanucleotide. It was observed that this gene

is subject to mutation earlier in men than in women. The first to be discovered, though, was the gene SOD1, which converts superoxide radicals to hydrogen and oxygen peroxide, and for long had been thought to be the primary mutation to cause ALS; other important genes involved are the following: dynactin 1 (responsible for backwards transport in axons), PFN1 (regulating globular and filamentous actin production, that determines nerve extension) and Tubulin 4A (for microtubules organisation). Familial ALS, in addition, has been shown to harbour mutations in the TDP-43 and FUS genes, which cause the accumulation of RNA-binding proteins in the cytoplasm of neurons rather than the nucleus [5].

1.2.3 Epidemiology

It was observed that sporadic cases of ALS account for 90–95% of all cases, with male individuals being twice as common as females, and the remaining 5–10% consist of familial forms of the disease. According to the European ALS Epidemiology Consortium (EURALS), the yearly incidence of the disease is 2–3 cases per 100,000 individuals, with the highest scores in Scotland and Sweden, but it is significantly variable across the continents and is particularly low (down to 0.6–0.8) for some indigenous populations of America and Asia. ALS is therefore considered a rare incidence, with a lifetime risk of developing it of 1 : 350 for men and 1 : 400 for women, although heritability is more probable in mother-daughter pairs [7][8]. According to recent studies, the worldwide prevalence of ALS patients is 4.1–8.4 per 100,000 people, changing significantly depending on the ethnicity of the population: for instance, it was observed that the prevalence for the European-American population is 5.4, while for Afro-American it is 2.3. These statistics about incidence and prevalence are expected to rise in the following years, considering the better standards and guidelines for diagnosis and the inevitable ageing of the population. The disease is known to arise around the age of 60–75 years: more precisely, the median onset is even 10 years earlier for populations from Asia and South America, while in Europe and North America it is more likely to start later; moreover, it's been noticed that the familial cases tend to show up earlier than the sporadic ones, but this remark might also be influenced by a closer attention to the symptoms, that leads to an earlier diagnosis, and by genetic reasons. Lastly, it is important to consider data regarding the type of onset, an important focus for this work, as will be discussed in the next sections. It was proven by many studies that younger men are more prone to develop ALS with a spinal onset, while cases of bulbar onset include more women, elderly people and subjects with cognitive impairment [5][9].

1.2.4 Disease progression and prognosis

All ALS cases are commonly characterised by the gradual decline of the patient's motor skills, and it's been observed that the disease's spreading signs and their timing are highly correlated with the type of onset: prospective studies of large patient populations have recently mapped the dynamics of symptoms as the condition evolves. To objectively quantify this progression in clinical settings, the ALS Functional Rating Scale-Revised (ALSFRS-R) is universally employed; it assesses 12 functional domains - including speech, swallowing, and respiratory function - to provide a score that correlates with disease severity and survival probability [10]. Using such metrics, it has been observed that the disease spreads in a directional, contiguous pattern; the site of onset is an essential indicator of how ALS will progress. In particular, the spinal onset cases have shown to degenerate more slowly and in a caudal-rostral direction, affecting LMNs first and resulting in limb weakness and atrophy, and impairing UMNs in later stages only, with the consequent loss of autonomy in vital skills, such as eating and breathing; the bulbar onset, on the other hand, involves the UMN region, with all its complications, since the beginning of the disease, and its rostral-caudal spreading to LMNs is proven to be much faster than in the spinal cases [11]. Most ALS patients have a poor prognosis: only 10% of them are affected by a slow-evolving version of the disease, which allows them to live longer than 10 years after the diagnosis, while the remaining 90% survives other 24–50 months, but can live up to 5 more years

if they undergo a tracheostomy. It was observed that patients with longer survival probability are men, younger ones, with spinal onset and who manage to gain weight after the diagnosis. On the contrary, for women, people with respiratory and/or genitourinary comorbidities, cognitive impairment, depression, and subjects who face weight loss after the disease onset, the symptoms generally evolve faster, leading to a shorter life expectancy after the diagnosis [9].

1.3 Locked-In State

While symptoms in the early stages of ALS are highly dependent on the onset type, the progressive loss of muscle control eventually leads all patients to near-total paralysis. It is estimated that 80–96% of them lose the ability to communicate verbally: facial muscles are proven to be the last to succumb to neurodegeneration, hence this condition is typical of the advanced stages of the disease [9][11].

Locked-In-Syndrome (LIS) is clinically defined as a state in which a subject is not capable of moving or speaking, but preserves vertical eye movement and blinking. In this stage, cognitive capabilities remain typically intact. Consciousness is assessed through standardised protocols that involve blinking or gaze control when questions are asked, and further exchange with the patient is possible thanks to great advances in the Augmentative and Alternative Communication (AAC) systems, which will be discussed in the next section [12].

Because the eyes serve as the primary channel for maintaining social interactions, extensive research has focused on vision and oculomotor function in ALS. Findings indicate that extraocular muscles are the most resilient voluntary muscles: they are characterised by differential gene expression and a unique calcium-buffering capacity, and their decay is significantly slower compared to limb and bulbar muscles [13][14]. However, substantial evidence suggests that, even in relatively early stages, higher-level oculomotor tasks may be impaired and various types of dysfunction may occur. About 70% of ALS patients report at least one ocular symptom, including dryness or blurred vision [15]; furthermore, the 50% of them present oculomotor dysfunction such as abnormal “cogwheeling” smooth pursuit and central nystagmus triggered by head movement, while the 41% are affected by saccadic dysmetria, which may lead to inaccurately following a target with gaze. Less frequently, patients experience square-wave jerks - rapid and involuntary horizontal eye movements - that reduce eye-tracking efficacy due to noisier acquisitions [16].

A 2015 longitudinal study evaluating visual search performance in mid-to-late ALS stage demonstrated that while reflex-driven oculomotor skills, e.g., prosaccades, are fully preserved, higher-level tasks requiring attention and voluntary control show increased error rates. These patients often require more fixations and corrective saccades to track a target. Although the study showed improvement through training, it demonstrated that eye-tracking is an optimal communication solution only for patients who still retain precise control of their gaze, whereas for others, new types of AAC become a clinical necessity [13].

1.4 Transition to Completely-Locked-In State

Despite their relative resilience, extraocular muscles eventually succumb to degeneration as well, leading ALS patients into a full paralysis state, clinically defined as Complete Locked-In State (CLIS). As confirmed by electroencephalography (EEG) or electrocorticography (ECoG), a patient in such a condition is fully conscious and may survive for years, even with a good quality of life, if supported by a positive environment and provided they receive adequate clinical and nursing care [14][17]. Multiple studies show that the LIS-to-CLIS transition is gradual and highly patient-specific, with the presence or absence of eye movements alternating for extended periods. Permanent CLIS is therefore declared after negative results in prolonged observations through electrooculography (EOG) and electromyography (EMG) of face muscles. However, caregivers at times report voluntary movements in response to “yes/no” questions even after the signals failed to be detected: this may suggest that alternative channels of communication are necessary to make the interaction with patients easier [18].

Research conducted in 2008 investigated the effectiveness of EEG-based Brain-Computer-Interfaces (BCIs) as a communication tool for patients with different levels of motor impairment. Findings showed that LIS subjects performed at above-chance levels and could improve with training, whereas CLIS patients had a significantly lower success rate. This study is the first to propose the fundamental hypothesis that, having transitioned to a completely locked-in state, the patient experiences deprivation of contingent external feedback because their effort produces no outcome; hence, they succumb to cognitive decay and “extinction of goal-directed thinking”, and eventually lose the capability to respond. They address the need to train LIS subjects before they transition to CLIS, so they can continue using BCIs even in later stages of the disease [19].

While it is a single-patient study, a 2010 publication by Ramos-Murguialday et al. reinforces this theory due to its longitudinal nature: the LIS-to-CLIS transition was mapped continuously using ECoG and non-invasive EEG, and the presence of somatosensory evoked potentials was investigated. It was observed that, while skin sensitivity diminished and became more selective, vision could be preserved through eye care, and proprioception and auditory pathways remained mostly intact [14].

Even up-to-date, more extensive research still focuses on proving the extinction hypothesis and preventing it, yet states that CLIS is a heterogeneous condition and that each case is unique; it is not possible to generalise. A 2021 review investigated the effectiveness of various non-invasive (EEG or fNIRS-based) and invasive (microelectrode array or ECoG-based) BCIs in patients with different stages of paralysis, observing that, while all LIS subjects successfully communicated regardless of the device, only a few single-case CLIS patients were able to answer “yes/no” questions with an above-chance success rate. This evidence, along with neurophysiological and neuropsychological measurements and observations, suggested that prolonged inactivity of muscles and respective cortical areas dedicated to movement and language articulation may be the cause of semantic learning and thinking extinction: CLIS subjects often exhibited a lack of attention, signs of boredom and habituation, and frequently fell asleep, proving that they were not actively engaged in BCI tasks. Furthermore, the study addresses other factors as possible causes for poor BCI performance in these patients: all the interfaces rely primarily on vision, which may be limited or even compromised in advanced CLIS, making it impossible to establish communication; regarding non-invasive devices, because of their low SNR, there is an actual technological barrier to the possibility of implementing an effective spelling or multiple choice BCI [18].

Consequently, the primary objective of this thesis work is to provide ALS patients during the critical transition from LIS to CLIS an effective communication tool, as they need to maintain their contact with the outside world after the loss of voluntary oculomotor control. By implementing a human-computer interface that requires minimal training and is complemented by auditory feedback, this research aims to overcome the limitations of traditional AAC systems and offer a robust, low-cognitive-load solution for communication in complete paralysis.

1.5 Assistive technologies for LIS and CLIS patients

Alternative and Augmentative Communication (AAC) systems represent a broad category of solutions designed to assist individuals with disabilities in performing daily tasks and enhance their quality of life by providing diverse means of expression. As the progression of ALS eventually deprives patients of their ability to speak, over the past decades, research has focused on systems that restore communicative autonomy for paralysed subjects, relying on their residual functional capacity. Thanks to the ongoing development of new technologies and the sustained support for research from caregivers and family members, the applications of AAC for ALS patients are continually improving in terms of accessibility and usability.

However, the primary challenge in managing neurodegenerative conditions lies in the progressive functional decline, which demands iterative fine-tuning and adaptation of the features to the patient's needs. Furthermore, in the case of complete paralysis, a significant challenge remains: achieving an optimal balance between effective communication and practicality in daily care environments [20][21].

1.5.1 Low-tech AAC and Eye-Tracking Systems

When dysarthria first manifests, the most commonly used aids are communication boards that display sentences, words, or single letters; the target is selected by the patient's gaze, a laser pointer, or a specific movement that can be interpreted by the caregiver. These low-tech methods, though, are reported to be limited and exhausting for both sides because of the high effort required, dependency on the patient's specific residual skills, and a significantly low communication rate [22].

Since the 1990s, eye-tracking devices have been developed primarily to overcome the ALS language barrier, but they were adopted in many other research fields - neuropsychology, computer gaming, driver monitoring, the advertising industry, and military weaponry, due to their typical low cost and high effectiveness. The principle of eye-tracking is based on live measurement of the eye's point of gaze or its motion in relation to the head. Such information can be collected through different techniques, divided into four main types.

The first one is based on an invasive acquisition: a magnetic search coil in a contact lens detects the eye's motion by measuring changes in the electromagnetic field, which are then recorded by an external mechanical or optical device. This method, though, is not suitable outside a research environment because it requires costly instrumentation and topical ocular anaesthesia.

Electrooculography (EOG) leverages the eye's dipole nature: the corneoretinal standing potential is measured by multiple surface electrodes placed around the orbit, and its changes indicate where the subject is gazing. EOG, though, is not the preferred technique due to several drawbacks, including reduced mobility from electrode wiring, usability limitations, and an uncomfortable setup.

Video-based oculography (VOG) represents the optimal eye-tracking solution in terms of versatility, reliability, and convenience. The system relies on live images of the subject's eyes or the entire facial area, acquired with visible- or infrared-spectrum cameras, and on segmentation algorithms that highlight univocal markers of gaze. The passive method uses the reflection of ambient light from the limbus, the eye's region between the iris and the sclera; however, this marker does not provide a stable, continuous signal due to intra-subject variability in eyelid shape, which may cover part of the contour. For this reason, infrared lighting and imaging are preferred, and the pupil is the standard reference to most of the commercial eye-tracking devices: depending on the illumination angle, the pupil either reflects the source (bright-pupil effect, obtained with on-axis rays) or is the darkest element in the picture (dark-pupil effect, due to off-axis radiation). The non-invasiveness of these eye-tracking devices, which may be wearable (e.g., headsets and glasses) or remote (e.g., webcam-based), makes them the most suitable choice for monitoring gaze in

ALS-LIS patients in a clinical or home environment. Although representing the gold standard for communication rates, ensuring a typing speed of almost 20 words per minute (wpm), they become unreliable when the patient is affected by nystagmus or other gaze disturbances, and cannot overcome the limitations introduced by eye motility loss in CLIS [23].

1.5.2 Brain-Machine and Brain-Computer Interfaces

As patients with LIS/CLIS present with motor impairment while their central nervous system functions are intact, the purpose of BMI/BCI technology in this case is to bypass the speech channel and establish a direct connection between the brain and the communication device. A standard BMI workflow comprises a signal acquisition module that collects live electrical or hemodynamic brain activity; a signal processing unit for communication-feature extraction; and a feedback actuator that triggers the desired command or interface output [21].

Non-invasive BMIs

Surface EEG-based technologies record electrical signals from the subject’s scalp and monitor localised cortical activity, its patterns, and responses to stimuli. Electroencephalography remains the most clinically accessible option for BCI, thanks to its relevant temporal resolution, but is limited in spatial resolution and has an inherently low signal-to-noise ratio. Furthermore, many of the possible applications require strong cognitive capabilities or long training, as they rely on motor imagery and other mentally challenging tasks to trigger the interface response; however, for Event-Related Potentials (ERPs) or Steady-State Visual Evoked potentials (SSVEPs) driven interfaces, a good performance can be achieved for multiple choice and spelling applications even with little training, as the brain rhythm is temporarily influenced by the external stimulus [24]. Meta-analyses indicate that these systems, in ALS patients, have had approximately a 70% classification accuracy for the last 15 years; performance even drops to chance levels for CLIS patients, likely due to the already mentioned “goal-directed thinking extinction hypothesis” and the high cognitive load required to maintain attention during tasks [25]. Hemodynamic monitoring, on the other hand, offers a valid alternative to EEG, as the activation of task-specific cortical areas is reflected in the neurovascular coupling effect: changes in oxygenated and deoxygenated haemoglobin in the surrounding blood vessels are detected through fNIRS and fMRI acquisitions, providing live cortical imaging via activation heatmaps. Despite being based on the same physiological phenomenon, these two techniques differ in terms of signal acquisition, spatial and temporal resolution, and sensitivity to motion artefacts and noise. Functional near-infrared spectroscopy (fNIRS) relies on tissue absorbance variations, recorded through emitter/receiver scalp electrodes, and it is characterised by a better spatial resolution and higher SNR than surface EEG; fMRI is a live magnetic resonance scan featuring HbO/HbR as contrast agents and, while providing higher quality imaging, it is less suitable for BCI because of too long time-domain responsiveness and is, overall, not as easy to set up and use [26].

Magnetoencephalography (MEG), even if purely for research purposes, has been applied to speech synthesis with significant success: the system measures the magnetic fields produced by the brain’s electrical activity, with a temporal resolution comparable to EEG, a much higher spatial resolution, and the advantage of no signal distortion due to interlayering tissues. A 2020 study selected the best sensor placement to achieve high accuracy in distinguishing imagined phrases, as the neural speech-decoding method did not rely on single-word spelling but on the intended mechanical articulation of the phrase or the imagined speech. However, a MEG machine can’t be used in healthcare or home environments, as it is large and expensive and requires a magnetically shielded room to prevent interference from external magnetic fields [27].

More recently, hybrid systems that combine EEG with fNIRS or eye-tracking have been developed to enhance system robustness and decoding accuracy through integrating multimodal neural dynamics [25].

Invasive BMIs

To overcome the limitations of scalp-level recordings, invasive approaches place electrodes directly on or within the cortex to acquire high-resolution signals and achieve the best SNR [25].

Microelectrode Arrays (MEAs) are implanted into the motor or premotor cortex to record individual neuronal firing. Recent and ongoing research has developed successful MEA-based speech synthesising and cursor control devices, achieving the highest communication rate among AAC options and over 90% accuracy in transforming neural patterns into acoustic features using deep learning algorithms. Electrocorticography (ECoG) is currently being investigated as a valid alternative: the electrode grid is placed on the brain surface (dural or subdural) and causes less tissue damage, while still providing a robust signal with minimal interference from EMG and higher stability throughout the years, promising long-term reliability even in home-use settings [28][25].

However, despite their higher performance and the potential for more natural communication, invasive solutions face significant challenges regarding long-term stability and biocompatibility. Electrode implants require surgery, raising the risks of complications for individuals with already compromised health, and sensors may detach or be compromised by the brain's biological environment, resulting in signal degradation over time [21][28].

1.5.3 Pupillometry as a novel approach to AAC for ALS

All the AAC solutions and devices presented in the previous sections encounter significant impediments at transition to CLIS: eye-trackers become unusable due to the absence of ocular motility, while BCIs impose a very high cognitive load and require access to expensive technologies. The most recent advances in BCI, which include machine learning (ML) algorithms to predict keyboard choices, showed they could speed up communication rates, making spelling more natural and less frustrating, but the core problems of using such devices with CLIS patients remain, as their learning rate and communication efficacy are significantly lower than control subjects and ALS patients in earlier stages due to the goal-directed-thinking extinction they may suffer [23]. This work proposes a different approach: by relying on an autonomic nervous system response unaffected by disease spread, it ensures an uncompromised channel of communication for CLIS patients with preserved visual function. Live pupillometry can be easily and reliably performed by any eye-tracking device and provides a novel foundation for a Human-Computer Interface (HCI) that is not only highly accessible in terms of training and cognitive load, but also convenient in terms of non-invasiveness and reduced costs, as the required instrumentation is relatively inexpensive and highly portable.

1.6 Pupillary Accommodative Response (PAR) as an HCI modality

The use of an autonomic nervous system (ANS) response associated with a voluntary action as a communication channel has been investigated by very few studies. Still, it represents a valid approach for ALS AAC, since the disease spreading does not involve parasympathetic and sympathetic responses. Research and clinical trials have focused on pupil autonomic control as the main communication tool, as it is completely spared in CLIS patients: pupil size in correlation with emotional activity and the possibility of controlling its diameter by engaging in cognitive tasks, such as mental calculation, were investigated. Due to the high variability of the pupillary response and the fatiguing nature of the activity, these principles were not optimal for an HCI, and alternative strategies for pupil-area variations related to vision are currently being studied. This work is based on voluntary modulation of one's pupil size in relation to a shift in the depth of gaze: previous studies have proven that it is possible to establish an effective communication

channel with a CLIS patient by monitoring their pupil diameter and leveraging on its autonomous response to a voluntary act. This section aims to introduce PAR and break down the findings to thoroughly justify the choices that led to the final device design and development.

1.6.1 Anatomy and physiology of the eye

From a purely engineering perspective, the human eye (Figure 1.2) can be modelled as a visible-spectrum camera because it has multiple anatomic components that play fundamental roles in converting light into an image. The eyeball’s outer layer comprises the sclera, a white vascularised tissue that surrounds and protects the organ, and the cornea, a round transparent window in the eye’s anterior cavity that refracts light towards the inner elements. Behind the cornea, which is typically 11.5mm in diameter, are the iris and the pupil, the two concentric elements that characterise the centre of the eye. The latter is the aperture through which light is gathered and conveyed to the crystalline lens, responsible for focusing and impressing an image on the retina, the sensitive film in the back of the eye that ultimately encodes visual information into electrical stimuli and sends them to the visual cortex, in the occipital area of the brain, through the optic nerve [29].

The eyeball is located inside the orbit, and six extraocular muscles are inserted into it from different angles to permit motility of gaze. These muscles, together with the levator palpebrae superioris, are innervated by three cranial nerves that originate in the brain stem. Voluntary and involuntary fixation movements are both mediated by the activation of extraocular muscles, but influenced by stimuli from different brain areas; the former, referred to when the subject actively searches and focuses on a target, rely on the primary motor cortex area, while the latter are performed non-willingly to adjust the focus and pursue an object, such as saccadic and optokinetic movements, and are controlled by the secondary vision cortex.

Similarly to the aperture of an objective, the amount of light that enters the eye is controlled by the pupil size, which changes according to the activation of the sphincter muscle present inside the iris. Like in any optical system with a convex lens, the curvature and thickness of the crystalline lens determine the focal length, i.e., the distance for all the light rays to converge in the back of the eye. To allow focus on multiple depth planes, the crystalline lens must reshape accordingly, and the ciliary muscle is deployed for this purpose: it is attached to the lens through a set of suspensory ligaments and, while contracted, it pulls them, making the refractive element thicker and therefore increasing its optical power [29][30].

1.6.2 PAR

Pupillary accommodative response, along with convergence of the visual axis and lens modulation to ensure focus, forms what is known as the near vision complex [31] - also frequently referred to in clinical literature as the near triad -, which is regulated exclusively by the ANS. The system is based on closed-loop feedback, starting with a blurred image on the retina that travels along the afferent pathway - the optic nerve - to the visual cortex, where the reflex is triggered. The efferent pathway involves the oculomotor nerve, responsible for convergence, and the Edinger-Westphal nuclei of cranial nerve III, which control lens accommodation and pupillary constriction by activating the ciliary muscle and the iris sphincter muscle [32]. The near triad elements, although working synchronously to ensure acuity of vision at any gaze depth and sharing neural pathways, are independent from one another, meaning that the impairment of convergence, typical of late ALS stages, does not prevent PAR from happening, provided the far and near targets are aligned and do not require eye movements to shift focus [30, 31].

Miosis, the pupil’s constriction, occurs stereotypically when the eye’s focus is shifted from a far to a closer target: the drop in the pupil area is rapid (100–300 ms), as it is governed by the parasympathetic nervous system, and lasts as long as the fixation is maintained on the same point. Analogously, when gaze shifts from a near to a more distant object, pupils dilate, albeit with a longer latency (500 ms), as the ciliary and iris sphincter muscles relax, driven by the

parasympathetic system. This reflex (Figure 1.3) is observed in any subject who does not present with complete vision impairment, is equal in both eyes, and shows no significant difference in individuals with uncorrected refractive errors, such as myopia or presbyopia [31].




Figure 1.2: Eye anatomy, cross section. Reprinted from [33]




Figure 1.3: Pupillary Accommodative Response, graphical representation. Adapted from [34][35]

1.6.3 PAR in HCIs

Due to its replicability, voluntary trigger, and autonomic nature, PAR is a competitive candidate for establishing communication between ALS-CLIS patients and their caregivers. Studies on ANS involvement in ALS suggested that pupil modulation is not affected [36], and subsequent research has presented methodologies for systematically developing a human-computer interface.

In 2019, Ponzio et al. conducted multiple experimental protocols to assess the environmental and setup conditions in which PAR is maximised. As expected, increased eye illuminance showed reduced pupil basal size, but the characteristic drop of PAR was observed with no significant difference regardless of brightness. Monocular vision, achieved by positioning a patch over one subject's eye, tended to elicit a larger response than binocular vision, but it was noted that it may be uncomfortable and more fatiguing. Target specifics were also defined: PAR is increased significantly by showing detailed targets rather than plain ones, as focusing is easier when the gaze can anchor to a precise point; furthermore, a near target was found to be brighter than a far target, making it the optimal setup. It was observed that the short duration of a far-to-near fixation does not reach the maximal constriction, which can be achieved after 3–4 s; a long duration, lasting 8 s or longer, revealed the phenomenon of pupil escape, i.e., the pupil's dilation after a PAR constriction, explained by the difficulty of maintaining long-lasting close focus. A first attempt was made to test the possibility of a communication channel: the healthy subjects were asked to answer obvious "yes/no" questions by performing a voluntary PAR when a specific audio cue was presented. The promising results (96% accuracy, 15 bits/min information transfer rate, and low response latency) set the stage for further validation in individuals with neurological impairment [37].

In 2021, a single case study conducted by Lorenzo et al. involved a patient with brain stem stroke and strong nystagmus, and featured a "yes/no" interface that the user could control through PAR: despite the eye drifting due to the neurological condition, voluntary constriction of the pupil could be clearly detected by the system, with an 80% accuracy [38]. In 2022, Cardinali et al. proposed training individuals to voluntarily modulate pupil size during gameplay: a 2D side-scroller game with a spaceship moving across the screen at a constant horizontal speed and a height determined by the user's pupil diameter. The game's goal is to drive the aircraft through as many objects as possible by performing consecutive constrictions and dilations, which result in up and down displacements, respectively. The interface was tested in a healthy control group and in 3 ALS-LIS patients, after a brief training session to understand and familiarise with the mechanism of pupillary modulation. Healthy individuals reported not feeling fully in control of the space shuttle; some found the game boring, and its difficulty showed high intra-subject variability, but skill improved progressively through repetition. Among the three patients, only one could successfully interact with the interface and even play with low effort, improving through replays, and found it entertaining [39]. The same hardware (which will be described and discussed in the following sections) was used by Chiarion et al. to develop e-Pupil, a stand-alone portable device that featured a "yes/no" interface, the possibility of calling the caregiver with a 3–6 s long miosis, and to ask for help with a pupillary constriction lasting more than 6 s. This study opened the possibility of multiple rather than binary classification, increasing information transfer rate, and demonstrated that a successful HCI can be implemented with minimal training, as reported by users [40].

In 2025, a new PAR application, the Music Machine, was proposed by Colucci et al., aiming to go beyond binary modality, not in a matter of duration, but of pupil size modulation instead: the setup included multi-level targets, placed at five different distances from the subject, thus PAR could be deliberately performed with variable amplitude by shifting gaze on multiple depth planes. Each target represented a different note the user could sound on a virtual saxophone to harmonise with the jazz track playing in the background. This research produced a significant breakthrough, demonstrating that PAR offers a broader range of possible variants than simple "yes/no" answers, even if it was conducted on a preliminary level only, and still requires extensive validation of the relation between target distance and pupil size, which is not linear, and repeatability across subjects [41]. The chart in Figure 1.4 shows the non-linear relationship between pupil diameter

and target distance that was qualitatively observed by Colucci et al., but finds strong theoretical foundation in both optics and physiology: accommodative demand measured in diopters is inversely proportional to the object’s distance, thus becoming non-linear on a metric scale [42]; the pupil has a physiologic minimum constriction, below which a saturation phenomenon occurs regardless the increase of proximity of the target [43]; intra-subject variability of baseline pupil size and dynamic range occurs across all age groups and refractive errors, which ultimately justifies a qualitative profile over absolute quantitative values [44].




Figure 1.4: Qualitative relationship between pupil size and target distance. Reprinted from [41].

1.6.4 Evolution of acquisition setups and interfaces

To translate the physiology of PAR into a functional assistive device, three basic technical requirements must be satisfied: video-based image acquisition of the subject’s ocular area, with correct illuminance, real-time pupil segmentation and metrics, and a classification algorithm that ultimately detects voluntary miosis, discerning it from background noise. Numerous prototypes were developed over the years, featuring changes in hardware and software architectures to meet the requirements and evolve, from complex to more portable set-ups, but encountered challenges of robustness and affordability, which demanded compromise.

The initial prototype, developed by Lorenzo Villalobos et al. [38], consisted of a commercial, low-cost eye-tracking device (EyeTribe, Denmark) connected via USB cable to a host laptop running a custom MATLAB program. This program performed live processing for artefact removal, low-pass filtering, and PAR detection. The interface that the patient interacted with was a simple “yes/no” desktop screen, with a vertical bar scanning back and forth between the two possible selections. Once a PAR was detected, the system would vocalise the translated patient’s choice.

To improve clinical practicality and reduce costs, the following systems’ design moved to custom-made, standalone options: a 2 MP infrared (IR) camera, coupled with IR LED illumination, was mounted on a wearable glass frame and acquired a close image of the eye, subsequently processed by a custom program on a Raspberry Pi 4 Single Board Computer. The interface was displayed on a touch screen, and, along with the “yes/no” feature, gameplay and the “music machine” described in the previous section were introduced in different versions of the same hardware architecture. A parallel work evaluated the possible usage of a visible-light camera with higher

resolution (8–12 MP), either remote or mounted on an analogous headset; however, the SNR of this type of acquisition was inherently low due to light reflections in the subject’s eye, thus lacking robustness and adaptability to ambient factors [38][39][41].

To achieve high portability and wearability at very low cost, a final legacy prototype was developed as part of a master’s thesis project [45]. A visible-light camera and a Raspberry Pi 0 were mounted on a glasses frame: the device was fully wireless and communicated via Wi-Fi to a laptop interface, featuring a spelling keyboard with word completion; the system overcame binary communication by implementing PAR recognition on two different depth targets and relying on PAR duration for multiple functionalities.

1.6.5 Challenges: robustness and reliability for pupillometry

The knowledge derived from prior research and art set the direction and highlighted the core design constraints for developing a fully functional PAR-based human-computer interface. The hardware choice is a primary challenge: previous studies have shown that live pupillometry can be performed successfully with both custom and commercial setups. The former would be preferred, ideally, because it is more cost-effective; however, many technological and structural bottlenecks were encountered and could not be easily overcome. The latter, while more expensive, relies on more advanced technology and ensures performance and stability.

A wearable acquisition architecture must be stable and ensure comfort and non-invasiveness. The lightweight spectacle frame meets initial ergonomics criteria; however, the unilateral integration of a camera introduces a considerable mass asymmetry across the frame axis, and counterbalancing the weight increases the overall load. Standard glasses frames lack rigid anchor points; hence, they are more likely to slip over extended use, and even a small shift in the camera direction or angle may result in signal loss or false positives.

The signal SNR represents another drawback of a custom headset mount: IR-light cameras enhance the contrast of the pupil against the surrounding area and reduce the effects of reflections and scattering in the visible-light range, thus representing the optimal choice for pupil monitoring; unfortunately, no low-cost IR-light cameras on the market are sufficiently small and satisfy the working distance requirement of a glass-frame setup (< 5 cm). The attempts with visible-light cameras, as anticipated, were too noisy and highly dependent on environmental conditions and the colour of the subject’s eye; the uncertainty and difficulty of pupil recognition made this setup unsuitable for further development.

Based on all the listed findings, an effective PAR-based communication tool must meet three basic requirements:

1. **Infrared illumination and imaging:** it maximises the pupil-to-iris contrast, making the system independent of room lighting and eye colour.
2. **Mechanical stability:** the hardware must remain motionless during the entire acquisition to prevent signal drift and ensure robust pupil recognition.
3. **Data streaming:** the device must provide reliable, real-time pupillometric data to ensure a smooth interface and exceed the minimum sampling rate of 30 Hz required to detect miosis events.

To overcome the physical instabilities and signal degradation inherent in custom-built wearable prototypes, a systematic trade-off analysis of commercial eye-trackers was conducted, evaluating real-time data transparency, software openness, non-invasiveness, mechanical stability, and cost-efficiency. A comprehensive technical comparison of these commercial platforms, including wearable headsets and remote-tracking systems, is presented in the comparative evaluation table in Appendix 5.

High-performance wearable eye-trackers, such as the already-owned Pupil Core headset, employ built-in IR illumination LEDs (860 nm) and high-speed eye cameras (up to 200 Hz) to perform pupil segmentation using the dark-pupil technique. The integration of proprietary eye-tracking models enables these wearable systems to produce high-fidelity metrics. Consequently, the prototyping and preliminary software development phases of the *Pupil-com* platform were successfully executed, utilising a Pupil Core headset as the data-acquisition source.

However, to fully overcome the limitations of wearable hardware and positioning constraints, the definitive configuration of *Pupil-com* was switched to a remote eye-tracking setup using the Gazepoint GP3 device. While other wearable and remote options offer high-resolution imaging, they are bound to proprietary analysis suites (e.g., D-LAB software for Dikablis Glasses 3) that prevent real-time integration in custom interfaces. Conversely, the Gazepoint GP3 offers an open-access real-time API (Open Gaze API) that streams raw gaze coordinates and binocular pupil diameter metrics over a local TCP/IP network at 60 Hz. The remote IR camera mounted on a static tripod guarantees mechanical stability and removes the physical and psychological annoyance of wearable instrumentation. The detailed hardware integration and software architecture supporting this remote setup are described in Chapter 2.

1.7 Aim of the study

In light of the limitations characterising current eye-tracking and BCI technologies for ALS patients in completely locked-in states, there is a clear clinical requirement for robust, accessible, and low-fatigue communication alternatives. As highlighted by the evolution of acquisition setups, overcoming the physical and signal-related constraints of custom prototypes is essential to translate PAR physiology into a functional clinical device. Therefore, the aim of the project was to develop and test *Pupil-com*, an AAC system operable through the user's voluntary modulation of pupil size in complete autonomy. The main objectives that were pursued are as follows:

- Integration of a low-cost, commercial eye-tracking device into a portable hardware setup suitable for both laboratory and home-care environments.
- Development of a custom software architecture featuring an adaptive signal-processing algorithm for the real-time detection of voluntary pupillary constrictions.
- Design of a user-friendly graphical interface enabling binary responses, keyboard typing, and entertainment. Particular focus on minimalism and intuitiveness to facilitate autonomous at-home use, without the intervention of an expert.
- Validation of the system's robustness, usability, and responsiveness through experimental trials on both a cohort of healthy volunteers and target ALS patients.

The detailed hardware integration and software architecture fulfilling these technical objectives are described in Chapter 2.

Chapter 2

Materials

This chapter provides a thorough description of *Pupil-com*'s framework. It details hardware and software specifics, along with the features introduced to support research-oriented data processing and improve user experience.

2.1 System Architecture

The communication system design meets standard BCI requirements and integrates features specifically required for PAR performance and monitoring. The setup comprised the following elements:

- **Acquisition module:** the eye-tracking device, wearable or remote, performing live pupillometry and providing real-time data;
- **Computational unit:** the PC, receiving a live data stream from the acquisition module, and running the program for the graphical user interface (GUI);
- **GUI display:** a monitor, showing a scanning-mode interface that the subject interacts with, and working as the far target;
- **Audio output:** external speakers, to vocalise cues and selections within the GUI;
- **External input devices:** mouse and keyboard, for the initialisation of the GUI and fine-tuning of parameters by the researcher or caregiver;
- **Universal laboratory stand:** holding the close target, a transparent plexiglass rectangle with a blue circular dot that the subject fixates their gaze on;
- **Seat:** for the subject to comfortably sit and maintain stillness while shifting the depth of focus.

Figure 2.1 illustrates the general system architecture during an experimental protocol.




Figure 2.1: Schematic representation of *Pupil-com*'s architecture during experimental validation. Image generated utilising Gemini AI based on the author's prompts

2.2 Hardware

The diagram in Figure 2.1 illustrates the hardware components and highlights the data flow. The computational unit is the core component of the structure, receiving raw data streams from the acquisition module and external inputs (key presses or clicks) and processing them to produce real-time audio and GUI outputs. This section describes each system block at the technical and functional levels.




Figure 2.2: *Pupil-com*'s hardware components

2.2.1 Acquisition module

A commercial eye-tracking device was used to acquire pupil data. As anticipated in Chapter 1, the choice of a commercial eye tracker was dictated by the need for stable and high-quality IR imaging of the eye, along with a robust, built-in pupil segmentation and metrics algorithm. The work involved two different eye-trackers: one wearable, in a lightweight, pervasive glass frame, and the other remote, as most traditional devices. Table 2.1 compares the two devices' specifications and PC requirements relevant to this study.

Wearable eye-tracker

Pupil Core (Figure 2.3) is the first headset developed by Pupil Labs, a German company operating in the field of wearable eye-tracking since 2014 and now a market leader. The ultra-light, ergonomic frame design ensures comfort during extended use. It integrates three cameras: one visible-light-range camera (a) placed on the front provides an image from the user's point of view, commonly known as the "world" camera; two IR cameras (b) and (c) and the respective IR LEDs are mounted on two symmetrical arms with length and pointing direction adjustable to the subject's facial features; a USB cable (d) connects the device to the computational unit, providing power and exchanging data. The presence of two independent cameras enables both binocular and monocular eye tracking. However, due to a malfunction in the right-side camera ("eye 0"), this work relied solely on monocular acquisition from the subjects' left eye ("eye 1"). Moreover, the "world" camera was unnecessary, since no gaze-pointing information was required. IR cameras provide a close-up image with a 192x192 px resolution at a maximum nominal frame rate of 200 fps, which may vary significantly depending on the receiving unit CPU and is inspectable live in a dedicated window. Acquisitions performed during this work ran at 120fps. The full datasheet is available on the Pupil Labs website [46]. Pupil Core runs on an open-source platform, Pupil Capture, provided by the house company. The segmentation algorithm is based on the dark pupil technique (Chapter 1). The upcoming Software section analyses the specifics of the data stream and its management.




Figure 2.3: Pupil Labs's Pupil Core eye-tracking device, detailing: (a) world camera, (b) and (c) IR eye cameras, (d) USB cable. Image adapted from [46]

Remote eye-tracker

Gazepoint 3 (Figure 2.4) is a remote eye-tracking module developed and commercialised by the homonymous company from Canada. This device is an affordable solution for this research, as it features a 60 Hz IR camera that meets the minimum fps requirement and enables high-resolution remote pupillometry. This configuration is preferable because it does not require the subject to wear a headset and allows them to use their own glasses, maximising comfort and supporting prolonged use of the communication system. An IR-LED array (a) illuminates the subject's facial area uniformly, and the integrated IR camera (b) collects the image at 0.3 MP. Two USB-C cables connect the device to the host PC: one for LED power and the other for data streaming. The acquisition module is mounted on a tripod at a height of 9 cm and can be angled as needed, via a spherical joint (c). The full datasheet is available on the company website [47]. GP3 runs on the manufacturer-provided private software, Gazepoint Control, installed after purchase, which performs pupil detection using the bright-pupil technique (Chapter 1) and streams pupil centre coordinates and diameter (in pixels or, after calibration, in mm) via an open-standard API, as described in more detail in the Software section of the present chapter.




Figure 2.4: Gazepoint's GP3 eye-tracking device, detailing: (a) IR LED-array, (b) embedded IR camera, (c) tripod mount. Image adapted from [47]

Table 2.1: Comparative evaluation matrix of commercial eye-tracking systems considered for the *Pupil-com* platform configuration.

Technical Feature	Pupil Labs - Pupil Core	Gazepoint - GP3
Setup Type	Wearable headset	Remote, desktop tripod mount
Eye-Tracking Tech.	Dark pupil technique	Bright pupil technique
Tracking Mode	Binocular and monocular	Binocular and monocular
Camera Specifications	2× IR custom CMOS Resolution: 192×192 px Frame rate: Up to 200 Hz	1× IR proprietary CMOS Resolution: 0.3 MP Frame rate: 60 Hz
Illumination	2× IR LEDs (860 nm) per eye	Multi-IR (860 nm) LED array
Pupil Measurements	2D: Position and diameter in eye-camera pixels 3D: Eye model parameters and diameter in mm via 3D mode (requires binocular tracking)	Position and pupil diameter in camera pixels (or estimated in mm after calibration procedure)
Cabling Infrastructure	1× USB-C cable	2× USB-C cables
Host Computing Specs	Intel i7 CPU or greater ≥ 8 GB RAM Cross-platform compatibility	Intel i5 CPU or greater ≥ 8 GB RAM OS: Windows 10 or greater
Software Ecosystem	Native suite: <i>Pupil Capture</i> Data Streaming: Real-time API for gaze, pupillometry, and video feeds over local network	Native suite: <i>Gazepoint Control / Analysis</i> Data Streaming: Open Gaze API for real-time gaze and pupillometry streaming over TCP/IP
Physical Dimensions	Weight: 22.75 g $160 \times 175 \times 51$ mm	Weight: 118 g $234 \times 47 \times 43$ mm
Eyewear Compatibility	No (requires custom corrective lenses)	Yes (compatible with glasses and contact lenses)
Working Distance	Close range (< 5 cm from eye axis)	Remote range (50 – 80 cm tracking depth)

2.2.2 Computational unit

The system architecture utilised a mini PC (MaxMini B8A Pro) as the core computational unit. This product meets all the project's requirements while remaining cost-effective. The mini PC offers high portability without compromising the performance levels of a standard laptop or desktop. Equipped with an AMD Ryzen™ 7 8745HS processor (8 cores, 16 threads), it not only fulfils the CPU demands for eye-tracking but also supports the development of multi-threaded, high-throughput user interfaces. These interfaces can run simultaneously with the acquisition software, providing low latency and avoiding graphical rendering delays. Additionally, its low thermal design power ensures stable operation under continuous loads and minimal fan noise compared to similarly performing processors, such as the Intel i9. Table 2.2 reports the unit's technical specifications, and Figure 2.5 displays it.

Table 2.2: Technical specifications of the MaxMini B8A Pro host computing device.

Component	Specification
Model	MaxMini B8A Pro
Processor	AMD Ryzen 7 8745HS (8 cores, 16 threads)
Clock Frequency	Base: 3.8 GHz Burst: 4.9 GHz
Graphics	AMD Radeon 780M
Memory	16 GB DDR5
Storage	512 GB NVMe SSD
OS	Windows 11 Pro
Connectivity	LAN: 1× 2.5 Gbps RJ45 Port Wi-Fi: 802.11a/b/g/n/ac/ax Bluetooth: 5.2
I/O Ports	2× USB 3.0, 2× USB 2.0, 1× Type-C 1× HDMI 2.1, 1× DisplayPort 2.1 1× 3.5 mm Headset jack, 1× DC 19 V
Dimensions	Weight: 430 g 126 × 112 × 52 mm



Figure 2.5: MaxMini B8A Pro.

2.3 Software architecture

The custom software for *Pupil-com* developed throughout the project is the core of this human-computer interface. Based on Python object-oriented language, it runs synchronously with the current eye-tracking device acquisition software. It is responsible for the data processing pipeline, the graphical user interface (GUI), and its real-time updates and responses. This section explores all aspects of custom software, including how it integrates and manages external inputs.

2.3.1 Acquisition module software

The commercial eye trackers employed in this work allowed the data stream to third-party applications, such as a custom-built interface, via a standard API, but only when the manufacturer’s acquisition software was running in the background. A brief description of both the programs follows.

Pupil Capture

Pupil Capture is an open-source software application provided by Pupil Labs, available for download on the company’s website [48]. It enables inspection of all three cameras’ views, helping adjust them to the user’s eyes and the surrounding environment, and choose between monocular and binocular eye tracking. Recording and calibration options are present, but this work did not make use of them because the former was unnecessary, and the latter requires eye movement, impaired in the target population of the system; however, calibration is not strictly necessary for full functioning pupillometry, since it only enhances the gaze pointer by integrating a 3D model, and provides the pupil diameter in millimetres instead of pixels. Real-time charts are available to monitor CPU, frame rate and other parameters that may vary during the acquisition. The main window (Figure 2.6, a) displays sidebars with multiple settings and plugins that enhance personalisation and user experience; however, they were all left at their default settings in this study, since only pupil recognition and metrics were relevant. The eye camera (“eye 0” or “eye 1”) is activated manually via a dedicated button in the main window, and its view is displayed in another window (b). The camera must be adjusted to frame the user’s eye and make sure the image is in focus (by manually rotating the objective, if needed): when the pupil is detected, a red overlay line delimits its border and a dot marks its centre (c); if the acquisition is not stable, due to unfavourable light conditions or out-of-focus image, the line is orange or does not appear at all. This video stream serves as the primary reference for initial configuration, and once that is complete, the external support can start the main GUI. The automatic pupil recognition algorithm is based on a dark pupil technique: its pipeline defines a region of interest (ROI) including iris and pupil, converts the image into greyscale, then uses Canny edge detection and global thresholding to identify the pupil region; connected components and curvature continuity criteria are applied to the edge points and fitted to an ellipse that represents the final pupil segmentation.

Gazepoint Control

Gazepoint Control is the manufacturer’s acquisition software for the GP3 eye tracker. The acquisition starts immediately after the program execution, and a one-window UI displays the full remote camera field of view (Figure 2.7, a) and, when pupils are detected, a close-up of both eyes below (b and c). The software is designed to provide an accurate gaze pointer, which can be enhanced with 5-point or 9-point calibration, and it can potentially integrate other biometric data, such as heart rate, for further analysis. Nevertheless, as already mentioned in the previous paragraph, the only data streams necessary for this project are the pupil centre coordinates and its size, thus making many of the eye tracker’s built-in features, including calibration, redundant. Placed between the settings bar and the live camera image, a horizontal line with a green circle signals the distance of the user’s head in respect with the GP3 (d); this feature is helpful during initial configuration: to achieve optimal positioning, the marker must be at the bar’s centre; if the distance is too small or too large, the marker is colored in orange and is placed at one of the bar’s ends - with the messages “Close” and “Far”, respectively - and the device may be repositioned accordingly. Unlike Pupil Capture, which is open source, Gazepoint Control is proprietary software. The manufacturer does not disclose the specific image processing pipeline or the algorithm that drives its bright pupil technique. The program essentially functions as a black box, processing the video stream internally and outputting the metrics via a standard server/client architecture.




Figure 2.6: Pupil Capture interface, detailing: (a) world camera, (b) eye 1 camera, (c) live pupil segmentation. Image acquired by the experimenter




Figure 2.7: Gazepoint Control interface, detailing: (a) full camera view, (b) and (c) close up of each eye, (d) distance check bar. Image acquired by the experimenter.

2.3.2 Custom software overview

The custom software driving *Pupil-com* was developed in Python within the Visual Studio Code environment. The full code is available in the dedicated GitHub repository [49]. To manage the distinct requirements of real-time data acquisition, signal processing, and user interaction, the application is built with a modular architecture that utilises specialised open-source libraries, categorised according to their functional role as follows:

- **User interface and real-time updates:** *Pupil-com* design uses the PyQt5 library, which provides a robust event-driven framework for a scanning interface.
- **Data acquisition:** Pupil Core eye-tracker required the `pymq` library, while GP3 relied on standard network sockets.
- **Signal processing and data management:** the libraries NumPy and SciPy performed numerical computation; Pandas and the native `csv` module served for data saving, while `pyqtgraph` ensured lighter and faster charting capabilities, compared to other plotting libraries.
- **Audio feedback and cues:** speech synthesis employed the library `pyttsx3`, and auditory triggers used `sounddevice` and `pygame`.
- **System logging and time management:** the native `time` and `datetime` modules were core for accurate timestamping, thread synchronisation, and internal clock management.

Figure 2.8 is a block diagram illustrating the data flow, *Pupil-com*'s internal processing, and its event-driven outputs. The following sections offer a detailed breakdown of the framework structure.




Figure 2.8: Internal data processing pipeline

2.3.3 Data ingestion

Pupil-com receives raw data packages via a standard API: a dedicated background execution thread acts as the client, pulling pupil centre coordinates and metrics from the server, which is managed by the eye tracker’s acquisition software. While the background thread unifies the data stream into the same format for the rest of the processing pipeline, the underlying communication protocols differ for the two recording modules: Pupil Capture uses a distributed messaging framework based on ZeroMQ (ZMQ) sockets, specifically a Publish-Subscribe (Pub/Sub) architecture. This allows the server to send compressed, thus lighter, data in binary MessagePack format, which the client unpacks into Python dictionaries with dedicated `msgpack` library tools. Gazeport Control communication is established over a standard TCP/IP (Transmission Control Protocol/Internet Protocol) socket. The GP3 server streams data continuously, in plain-text XML format, from which the Python client extracts the required gaze coordinates and pupil dimensions. The pupil diameter is initially measured in pixel units and converted by the receiving block into area metrics. Using quadratic rather than linear measurements, as area is monotonically related to diameter ($A = \frac{\pi}{4}d^2$), does not alter the underlying classification logic; it may also improve performance, since a percentage reduction in pupil size is mathematically enhanced when tracking area. This quadratic scaling expands the dynamic range of the signal drop during a PAR trigger, which can be advantageous for threshold-based detection algorithms. Furthermore, it aligns with the analysis performed by legacy systems, ensuring continuity in data processing and decision-making.

2.3.4 Signal preprocessing

The raw metrics stream, as unpacked by the data ingestion module, may be inaccurate and highly variable due to involuntary movements, eye blinks, and sudden changes in illuminance. Real-time filtering is crucial for stable, high-quality pupillometry to prevent miosis recognition errors and enhance the user experience. The `AreaFilter` execution block performs frame-by-frame sequential operations: as the server delivers a new raw area value, a fixed-length double-ended queue (deque) is updated, and the preprocessing pipeline, based on and adapted to legacy lab frameworks, is applied in series, in this specific order:

1. **Minimum-maximum area value range (in px²):** fixed to avoid errors in miosis detection caused by signal loss or misrecognition of objects as the pupil. The latest in-range value temporarily substitutes discarded frames; a timeout loop checks the area validity and triggers an error after 5 consecutive seconds of signal loss, freezing the GUI and waiting for the support user to fix and restart the acquisition manually (Figure 2.9).
2. **Event-based filter (EBF):** a derivative filter introduced to remove high-frequency spikes that may occur due to motion artefacts, eyelid lowering and blinking. It evaluates the absolute difference between two consecutive frames to monitor the signal’s rate of change. It rejects new frames that exceed a specific threshold ($> 100 \text{ px}^2$ in GP3 and $> 500 \text{ px}^2$ in Pupil Core) from a dynamic baseline, calculated as the average of the previous 0.25 s deltas. An internal counter prevents system lockout: after 0.8 s during which anomalies have persisted, the EBF stops holding the pre-artefact values and adapts to the new signal level.
3. **Moving average filter (MAF):** a real-time low-pass filter applied to remove high-frequency noise and improve signal stability. The MA window length is 0.5 s, a parameter that represents the optimal trade-off between good smoothing and the detection of pupil size variations without introducing noticeable latency.

It is important to note that the code is hardware-independent. By explicitly setting the frame rate and area boundaries, the filtering block applies the same mathematical operations to any commercial eye-tracking data stream. The filtered area values are then used in parallel to provide

real-time updates on the pupil area in the UI and, along with the pupil centre coordinates, for digital 2D eye rendering. Furthermore, the PAR detection algorithm performs its analyses on the output of the preprocessing block.




Figure 2.9: Signal lost error dialog

2.3.5 PAR detection algorithm

A dedicated code block, `ConstrictionMonitor`, handles the filtered data stream and performs a frame-by-frame pupil classification, out of four possible states: relaxed (Type 0), when no PAR is detected; short-duration constriction (Type 1), for PAR events lasting less than 3 s; long-duration constriction (Type 2), when PAR has been held for longer than 3 s, but less than 5 s; in an extra-long constriction (Type 3), for PAR lasting longer than 5s. Inspired by the implementation described by Chiarion et al. [40], this multiple-duration miosis detection strategy enables communication beyond binary and allows the user to access various GUI functionalities. The classification algorithm utilises dynamic thresholding and internal timers to identify and update the pupil's current state. The operational pipeline consists of the following steps:

1. **Baseline collection/update:** A deque buffer stores the last 2 s of the filtered area stream. The monitor starts checking for a constriction only after completing baseline collection. During the first 2 seconds of the acquisition, the user is invited to look at the screen (far target) to ensure the pupil is in a relaxed state and to provide a reliable area value. Once the buffer is full, the latest value is appended, while the oldest is dropped.
2. **Adaptive threshold calculation:** To ensure both adaptability to slow pupil area drifts due to environmental changes and sensitivity to fast, intentional PAR events, the system calculates a dynamic Simple Moving Average (SMA) threshold, T_{SMA} . This is computed frame-by-frame by multiplying the current baseline mean ($\mu_{baseline}$) by a manually calibrated sensitivity factor k (in the 0.7–0.9 range):

$$T_{SMA} = \mu_{baseline} \cdot k$$

3. **Trigger detection (Pupil status check):** The monitor continuously compares the incoming area to the adaptive threshold. If the area falls below this value, the algorithm enters a listening state loop: it registers the current T_{SMA} and locks it as a fixed exit reference (T_{exit}). Simultaneously, an internal timer begins tracking the duration of the miosis event.
4. **Duration-based classification (Listening state):** After the counter starts, the monitor waits for the minimum short-constriction duration (0.2 s) before validating a Type 1 event,

to avoid misclassifying unintended area drops while maintaining a fast response when miosis is intentional. If the signal remains below T_{exit} over the long and extra-long timer limits, the monitor dynamically upgrades the classification to Type 2 and Type 3 events, respectively.

5. **Fatigue compensation and dynamic reset:** In standard conditions, the monitor resets (Type 0 event) when the user relaxes their eyes by looking at the far target and the pupil area rises back above the fixed T_{exit} . However, empirical observations indicate that prolonged miosis, or successive short constrictions, is typically followed by a smaller delta in mydriasis, most likely due to ciliary muscle fatigue; this often results in a pupil dilation that is insufficient to return above the fixed threshold, potentially trapping the monitor in an infinite listening loop. To prevent lockouts and enable the user to perform multiple pupillary constrictions within a short period, the algorithm implements a double-condition reset. The T_{SMA} is continuously updated and checked in the background: it is mathematically certain that, while being computed with constricted data, its value is dragged below the pupillary area signal. By tracking the intersections between the signal and T_{SMA} , the algorithm can identify a new constriction attempt even if the pupil fails to dilate fully. If the signal crosses above this lowered T_{SMA} and subsequently drops back below it, indicating a new intentional constriction, the monitor forces a reset, releases the fixed T_{exit} , and immediately begins tracking a new event.

Figure 2.10 provides a state machine diagram of the PAR detection algorithm.

In Figure 2.11, the graph demonstrates how the system recognises and updates PAR events, while Figure 2.12 shows how the threshold is recovered after a double-crossing.




Figure 2.10: State machine diagram of ConstrictionMonitor.




Figure 2.11: PAR detection and type update over time. Postprocessed from experimental acquisitions.




Figure 2.12: T_{exit} recovery after double crossing T_{SMA} . Postprocessed from experimental acquisitions.

2.3.6 Graphical User Interface

The front-end architecture of the *Pupil-com* design offers the user full navigation autonomy once the supporting person (researcher/caregiver) has completed the session initialisation and setup. All the messages, labels, and text-to-speech features are in Italian.

Session initialisation

Upon launch, the software opens an initialisation dialog (Figure 2.13) that requires standard manual input from an external operator. The screen prompts the selection of the eye-tracking device in use (Pupil Core or GP3) to configure the correct network sockets, and requests the subject’s name for data logging purposes. To accommodate privacy requirements, one can initiate an anonymous session via a dedicated “Skip” button.

After completing initialisation, the view transitions to the application Main Menu in a standby state (Figure 2.14). The caregiver must manually start data acquisition by pressing the dedicated button (“Avvia il dispositivo”). Once the device is correctly set up, the live area label and 2D eye rendering are updated in sync, and the green start button can be pressed (“Dispositivo pronto”). The internal UI clock for interval scanning starts, and total control is handed to the user.




Figure 2.13: Startup screen to set the device in use and the subject’s name.




Figure 2.14: Main Menu initialisation screen.

Main Menu

After baseline collection, the scanning interface activates: the buttons on the GUI Main Menu screen are highlighted in sequence (Figure 2.15), while a speech synthesiser vocalises the corresponding label. The scanning interval is 3s by default and, along with other parameters, can be manually modified by the supporting person in a dedicated Settings dialogue, which will be described further at the end of this section. The *ConstrictionMonitor*, running in the background, constantly checks for a short-duration pupil constriction, which indicates the voluntary selection of the currently highlighted button. The GUI maintains the scanning mode active until a PAR event is recognised and interpreted as a click command, switching from the Main Menu screen to the Application corresponding to the user’s choice. If necessary, the supporting person can press all interface buttons by mouse click.




Figure 2.15: Main Menu while scanning the options by highlighting the corresponding button.

Application 1: "Training"

This application, as its name suggests, was introduced to provide the user with an initial training tool and help them familiarise with the system, but it has also been included in the system's validation protocol, since its results define the PAR recognition algorithm's performance relative to the manually selected threshold sensitivity and the subject's responsiveness. The 2D eye rendering is displayed to help the user visualise, in a mirror, the dynamic behaviour of their pupil. The session can be initiated either by the user, when the scanning highlights the "Start" button, or by a click. In each training session, the user is guided through 5 consecutive PAR-relaxation events by vocal instructions that invite them to focus on the near target, then shift to the far target, and repeat. The session comprises 3 short-type and 2 long-type constrictions, requested in a randomised order. The system monitors pupillary activity, records all detected events, and provides a score at the end of the session indicating how many constrictions were correctly recognised. Moreover, a chart appears on the screen, displaying the filtered area stream over time, the active threshold, the instant a constriction was requested, and, if so, the instant the constriction was detected. An example of the results screen presentation is provided in Figure 2.16.

Application 2: "Yes or No?"

The simple front-end design of this application is a recurring interface in AAC devices: two large buttons, labelled "SI" ("Yes") and "NO" ("No") respectively, occupy the majority of the screen and are highlighted in turn (Figure 2.17). The module enables vocalisation of a binary answer to an external question with a short PAR, and features a "Pause/Return" menu accessible via a long PAR. The scanning logic is more refined than the simple highlight-and-select of the Main Menu, since the system must handle distinct event types and introduce a cooldown interval after a miosis detection to avoid accidental selections while the pupil is relaxing. To achieve these requirements, the closed-loop state machine (Figure 2.18) on which this widget operates was designed as follows:

- **Initialisation/Question formulation:** When the application is first started, the interface is in a standby state, and the `ConstrictionMonitor` module restarts baseline collection to prepare for the next PAR recognition. The scanning highlight is disabled, and the subject is instructed to focus on the screen and relax their gaze to achieve full mydriasis. During this




Figure 2.16: Result screen after completing one "Training" session.

standby state, the supporting person asks a "yes/no" question, then triggers the system's Active Scanning state by clicking the mouse.

- **Active Scanning:** The widget's idle state, in which it automatically toggles the highlight between "SI" and "NO" buttons at a fixed interval of 3.5 s. The system continuously polls the monitor for PAR triggers.
- **Selection and confirmation (Type 1 Event):** When the algorithm reports a Type 1 constriction, the scanning loop is instantly interrupted, and the current highlighted answer is frozen, providing visual feedback that the system is listening. To avoid vocal feedback activation before the PAR type has been defined, the trigger occurs only after the area signal returns above the fixed T_{exit} . Afterwards, the state machine switches to the Cooldown state.
- **Cooldown:** A mandatory 1.5 s timeout phase immediately following the answer. This prevents the system from misinterpreting post-constriction pupil settling as a new trigger. After the timer expires, the loop returns to the Initialisation/Question formulation state.
- **Pause and Exit Menu (Type 2 Event):** When the pupil remains constricted for longer than 3 s, the "SI"/"NO" buttons change to "RIPRENDI" (Resume) /"ESCI" (Exit) options, and the system switches back to Active Scanning. In this case, the user can choose to re-initialise the "Yes or No" widget or to navigate back to the Main Menu via a Type 1 selection.

Application 3: "Keyboard"

This is *Pupil-com*'s core application, enabling users to write full words by simply spelling them or selecting one of the three suggestions, as on a smartphone keyboard. The widget has the most complex interface, since it features an alternative row- and column-scanning approach to minimise latency and the number of PAR needed to transcribe a word, thereby raising the information rate.



Figure 2.17: Scanning interface highlighting one button at a time




Figure 2.18: "Yes or No?" state machine loop.

All three PAR types are employed in its operation. Figure 2.19 presents the interface: the text is displayed in a dialog box at the top, just below the title; the main keyboard layout is a 5×6 grid, in which all the Italian alphabet letters and main punctuation symbols are displayed in an arbitrary order, chosen according to the respective frequency in the language. There are two additional rows: one above the grid, containing three Suggestion buttons, dynamically updated by a dedicated algorithm to match user-spelt words to the most likely in the Italian language; the one below provides two keys for "Space" and "Backspace" functions.



Figure 2.19: "Keyboard" layout, row scanning active.

Predictive suggestions algorithm

The self-updating suggestions mechanism is powered by an internal routine that combines tools from two open-source Python libraries:

- **autocorrect**: through its **Speller** object, provides the best typographical matches to the current written word;
- **wordfreq**: provides a reference list of the 10,000 most common words in Italian.

Each time a new character is selected, the algorithm compares the current string to items in the **wordfreq** list starting with the same sequence: if it finds at least three candidates, those are displayed in the Suggestion buttons; otherwise, it uses the **Speller** to find words that most likely match the keyboard input, by changing one or more characters; and if no matches are found, the three initial default vocabulary words are displayed.

State machine loop

The internal state-machine logic was developed using the Application 2 architecture as a reference and expanded to support matrix navigation rather than a binary interface. Figure 2.20 represents the schematics, and its operational blocks function as follows:

- **Initialisation**: Upon widget launch, all keys are disabled while the user's pupil returns to a relaxed state and the monitor recalibrates to the baseline. When the initialisation timer expires (2 s), the system transitions to the Row Scanning state.
- **Row Scanning**: All the Keyboard rows, excluding the Suggestions row, are scanned on alternate turns at a fixed scanning interval (1.5 s, editable parameter). The user must perform a short PAR when the row containing the target letter is highlighted. The scanning loop continues until the **ConstrictionMonitor** reports a miosis event, and moves to the Cooldown state after confirmation of the event type.
- **Column Scanning**: Once a row is selected, it highlights the keys contained in the selected row, one at a time, at the scanning interval. Via a short PAR, the user selects the target letter displayed in the display box, which is then vocalised by the text-to-speech synthesiser. Moreover, the Suggestion buttons are updated accordingly. An additional feature enables the user to specifically trigger a simple backspace when performing a short-type PAR, or to activate a full cancel option via a long-type PAR on the "CANCELLA" button: this erases the last word completely.
- **Suggestions Scanning**: Entered and aborted via a Type 2 event, this state highlights the Suggestion boxes in turns during Rows or Column Scanning. The user must perform a short PAR to select the current option, or a long PAR to switch back to Rows Scanning.
- **Cooldown**: The intermediate state between an event detection and the subsequent Scanning state, introduced to prevent the system from misinterpreting commands during the pupil's relaxation phase. To make the Keyboard user experience smoother, this latency was set very low (0.5 s).
- **Pause and Exit Menu**: Triggered by a Type 3 PAR, regardless of the state the machine was previously in. As in Application 2, the screen displays two buttons, one to resume the Keyboard widget and the other to navigate back to the Main Menu, which are selectable via a simple short PAR.

The application is fully functional, though both the graphical design and the suggestions algorithm can be largely improved: these possibilities are further analysed in the Discussion chapter.




Figure 2.20: "Keyboard" state machine diagram

Application 4: "Shuttle Game"

The last application featured in this project is a single-player game, inherited from the BCIs that were previously developed by prof. Roatta's research group. The pilot version of the game is described in Chapter 1; the one featured in this work, however, is based on a later, unpublished version with more refined graphics and multiple options.

Gameplay and graphics

The primary aim of this 2D-scrolling game is to entertain the user, who can launch a spaceship by simply performing a short PAR to hit the planets and score 1 point, passing horizontally across the screen at a fixed speed and a random height. At the beginning of the game, the player is at Level 1 and is given 5 Lives, meaning they can miss a planet 5 times before Game Over. Every 5 successful hits, a Level-Up is granted, an extra Life is provided as a reward, and the planet's speed is raised, thus requiring faster response and increasing the overall gameplay difficulty. For Level 5 and above, an additional feature is introduced: asteroids may appear on the screen and hitting one of them costs 1 lost Life. The game's scope is therefore to reach the highest score possible by hitting the planets and avoiding the asteroids. After selecting the "Game" button in the Main Menu, the GUI opens a game lobby where the user is welcomed, and the game is briefly explained via a text message. Via a simple 2-options-scanning, the user can start the game or navigate back to the Main Menu. After "Game Over", a message screen reports the user's current game results, along with their three all-time top scores. After 5 seconds, the system automatically goes back to the lobby, where another game session can be started after a brief cooldown. The game aspect (Figure 2.21) consists of a cartoon space shuttle on a full-screen starry sky background; an upper bar shows a variable number of heart-shaped Lives indicators, and a label with the current Level and Score. An "Exit" button is accessible via mouse click in the upper-right corner. The planets' aesthetics are randomised from a collection of cartoon-like versions of the solar system's celestial bodies.

Game development and internal loop

Unlike the other Applications, SpaceShuttle's graphics, sounds, and game rendering logic were not developed during this thesis project but were inherited from previous research work, which was adapted to Pupil-com's internal architecture and miosis detection algorithm. A description




Figure 2.21: "Shuttle game" layout.

of the state machine loop that drives the widget follows:

- **Initialisation:** Upon game launch from the lobby, two messages ("Ready?"/"Go!") are displayed by default to briefly prepare the user. During this 3-second interval, the `ConstrictionMonitor` collects a fresh baseline, then the machine moves to the Active Gameplay state.
- **Gameplay Update:** The targets (planets and asteroids from Level 5 on) spawn one at a time on the right side of the interface, and disappear when they reach the left side, or when the spaceship collides with them after launch. The system continuously polls the background `ConstrictionMonitor` for a miosis event and updates the objects' positions on the screen accordingly.
- **Shuttle Launch:** As soon as a Type 1 constriction is detected (0.2 s below the active threshold), launch is triggered. Unlike the "Keyboard" and "Yes or No?" applications, which wait for the pupil's relaxation to confirm the event type and, therefore, activate the respective command, in this case, the absence of latency is essential to ensure an instant response to the player's command and maximise user experience. While the shuttle travels vertically across the screen, three checks are performed: collision with a target, target miss or Type 2 event confirmation.
- **Score (and Level) Update:** When a collision with a planet is confirmed, Score and Level are upgraded accordingly, and the GUI plays a positive audio feedback; then the state returns to Gameplay Update.
- **Life Lost:** A target miss, or an asteroid collision, causes the loss of one Life: one of the heart-shaped indicators is deleted and a negative audio feedback plays. If the Lives count is still above zero, standard Gameplay continues; if it's null, the next machine state is Game Over.
- **Pause:** When the `ConstrictionMonitor` recognises a long PAR, the game is paused, and a dark screen with two standard "Resume"/"Exit" buttons is highlighted alternatively. A Type 1 selection can restart the game where it was left off or go to the Game Over screen.

- **Game Over:** Final screen described in the previous paragraph.

A diagram of the state machine is found in Figure 2.22.




Figure 2.22: "Shuttle game" state machine

GUI personalisation

To ensure an optimal user experience, an external support person can manually adjust many of the GUI's parameters. A dedicated Settings dialog (Figure 2.23), accessible from the Main Menu, presents all the parameters divided into tabs, each for a different program's application. Table 2.3 lists all editable settings and their impact on usability.

Data logging and storage

Pupil-com implements a background thread dedicated to session logging and data saving, with dual storage: local on the host PC and cloud synchronisation via a Google Drive folder. The former option is available unconditionally, since it's offline and relies on native Python libraries (`os`, `csv`); the latter requires internet access and the pre-configured API permissions required by the cloud storage system, via the `google-auth` library for Python. This cloud integration was introduced to support remote data access; by enabling real-time telemetry, the researcher can monitor the patients' at-home usage performance during prolonged, unsupervised periods without requiring physical presence in longitudinal trials.

Upon session initialisation, a new directory is created in the "Experimental Results" folder. It is named after the subject, if such information is provided, or with the current timestamp, if the anonymous option was chosen. Three custom Python objects are dedicated to data handling in the `HelperClasses` modules:

- **SessionLogger:** Constantly updates the `session_log` text file, a global timestamped report of the current session activity, the user's actions, and warning/error messages when they occur. It offers essential support for troubleshooting in the software development phase and provides a solid reference for assessing anomalies during post-processing.

Table 2.3: Summary of configurable GUI parameters across different application tabs, detailing their default values and system impact.

Dialog Label (Italian)	Parameter	Default	Impact
General Tab			
<i>Dispositivo attivo</i>	Device in use	Session def.	Selects acquisition software and internal fps
<i>Soglia</i>	Threshold sensitivity	0.8	Multiplier to calculate T_{sma}
<i>Durata costrizione breve</i>	Type 1 PAR minimum duration	0.2 s	Type 1 PAR dwell time
<i>Durata costrizione lunga</i>	Type 2 PAR minimum duration	3.0 s	Type 2 PAR dwell time
<i>Tempo inizializzazione</i>	Main Menu initialisation timer	2.5 s	Latency prior to scanning sequence
<i>Intervallo scansione</i>	Scanning interval	3.0 s	Button highlight speed
Yes or No Tab			
<i>Tempo inizializzazione</i>	Widget initialisation timer	2.5 s	Latency prior to scanning sequence
<i>Intervallo scansione</i>	Scanning interval	3.5 s	Button highlight dwell time
<i>Tempo cooldown</i>	Cooldown interval	1.5 s	Latency after answer
Keyboard Tab			
<i>Tempo inizializzazione</i>	Widget initialisation timer	2.5 s	Latency prior to scanning sequence
<i>Cicli massimi</i>	Looping time	1	Forces reset to Row Scanning to prevent lock-outs
<i>Intervallo scansione</i>	Scanning interval	1.5 s	Button highlight dwell time
<i>Tempo cooldown</i>	Cooldown interval	0.5 s	Latency after selection
<i>Durata cost. extra-lunga</i>	Extra-long constriction min. duration	5.0 s	Type 3 PAR dwell time
Training Tab			
<i>Tempo inizializzazione</i>	Widget initialisation timer	5.0 s	Latency prior to guided session
<i>Durata richiesta breve</i>	Short task duration	2.0 s	Type 1 request timer
<i>Durata richiesta lunga</i>	Long task duration	5.0 s	Type 2 request timer
<i>Durata richiesta lontano</i>	Far task duration	3.0 s	Pupil relaxation timer
Game Tab			
<i>Tempo inizializzazione</i>	Widget initialisation timer	3.0 s	Latency prior to game
<i>Tempo cooldown</i>	Cooldown interval	2.0 s	Latency after pause trigger
<i>Velocità astronave</i>	Shuttle speed	1 (a.u.)	User avatar velocity
<i>Velocità base pianeta</i>	Planet transit time (Level 1)	7 s	Time required for a target to cross the screen
<i>Numero vite</i>	Initial Lives count	5	Maximum target miss before Game Over




Figure 2.23: Settings dialog in the "General" tab

- **DataSaver:** Each time an application is started, this object stores frame-by-frame data in temporary buffers. When exiting the application and returning to the Main Menu, these buffers are saved in a dedicated `.csv` file. The exported dataset structure includes a timestamp, raw and filtered pupillary area, MA-based and exit thresholds, PAR event code, and, when necessary, additional application-specific information (e.g., selected characters or state machine condition). This file serves as the foundation for advanced offline data analysis.
- **DataPlotter:** It runs in parallel with the **DataSaver** and employs `matplotlib` to generate and save a `.png` chart with the essential data from each application use: filtered area, the two thresholds, and constriction detection marks, against time. This additional feature enables immediate visual assessment of the user's physiological response and the system's accuracy, without processing the CSV file, and is especially useful for manual calibration of threshold sensitivity.

Chapter 3

Methods

This chapter details the methodology that guided the research workflow throughout its operational phases: the system’s iterative software development, the subsequent experimental validation and performance testing of the final prototype on a pool of healthy subjects, and the ultimate protocol designed for presentation to target users. Finally, it outlines the data analysis framework conducted to assess system efficacy and quantify user-experience-related metrics for future improvement.

3.1 Iterative developer testing

The initial phase focused on developing a reliable custom software architecture and enhancing the PAR detection algorithm, accomplished through iterative feature development and thorough testing by the developer.

During early development, the acquisition device in use was the Pupil Core wearable headset. The device’s accurate, high-speed metrics were essential for evaluating the raw pupillary area signal, mapping temporal signal trends, and designing an operational workflow to correctly detect miosis and mydriasis events.

The implementation of the current double-threshold detection algorithm (detailed in Chapter 2) is the direct result of these iterative trials. Technical prototype testing revealed that relying on an SMA thresholding method alone precluded the detection of secondary commands (Type 2 and Type 3 PARs), due to its characteristic delayed drift below the signal; conversely, a simple fixed exit threshold introduced a high risk of user lock-out if the pupillary area failed to recover to baseline fully.

Before the official validation phase, a pilot study (PS) was conducted on a sample of 11 healthy volunteers (4 females, 7 males); the sample included a diverse age distribution (14–50 years old) to capture a broad spectrum of physiological baselines, and the subjects were recruited from a university peer group or immediate community; these tests were performed on the alpha prototype of the system, which still relied on Pupil Core for data acquisition and ran on a personal laptop. Up to this point, system benchmarking had been restricted to developer testing, meaning the algorithm parameters were calibrated exclusively to a single subject’s physiological response and training level. This pre-validation step was therefore critical for assessing inter-subject physiological variability, optimising lighting constraints, and defining ideal target distance and thresholding requirements.

This pilot version of the Pupil-com software featured three applications: “Training”, “Yes or No?” and “Shuttle Game”. In this stage, the users were asked to perform multiple “Training” trials, until full familiarisation with the voluntary PAR act, and then to freely navigate the interface and interact with the other modules.

The data acquired during the pilot sessions provided essential information to finalise the double-thresholding method and enhance personalisation by introducing manually editable parameters in the dedicated Settings dialog. Furthermore, the subjects' qualitative verbal feedback on ergonomics, visual comfort, and user experience established the design requirements for the subsequent development phase; this directly influenced the definitive system architecture, including the transition to the remote eye-tracking device GP3, which maximised user comfort.

3.2 Validation protocol

The final system configuration, based on the Gazepoint GP3 eye-tracker and detailed comprehensively in Chapter 2, underwent a formal validation study (VS) involving a pool of healthy subjects. The primary objectives were to assess the efficiency and stability of the PAR recognition algorithm, evaluate user-specific operational performance, and quantify overall BCI accessibility.

The 15 healthy participants (8 females, 7 males) were recruited from an academic cohort aged 21–30 years. To evaluate the impact of familiarisation with the system, the validation cohort was divided into two distinct subsets: Group A, consisting of 7 individuals (2 females, 5 males) who had previously participated in the pilot investigation, while Group B included the remaining 8 participants who were approaching the Pupil-com platform for the first time. This study took place exclusively in the Laboratory of Human Physiology, Department of Neuroscience “Rita Levi Montalcini”, University of Turin.

The protocol design was strictly standardised and comprised the following steps for each participant:

- **Introductory Briefing:** The investigator delivered a standardised presentation detailing the physiological mechanism of the PAR, instructing participants on how to voluntarily trigger the reflex by simply shifting the gaze from a far fixation point to a near target, and how to use this as a command to navigate the Pupil-com interface.
- **Demonstration:** For Group B subjects, the introductory briefing was followed by a live operational demonstration. The investigator interacted with the system in real time, allowing participants to observe the GUI controls and behaviour before initiating their own trials.
- **Experimental Setup:** The volunteer was seated in an ergonomic chair positioned at a fixed distance of 1.60 m from the GUI screen. The chair height was adjusted according to the subject's profile to align their line of sight with the screen centre, maximising visual comfort and postural stability. Subsequently, the investigator positioned the near target 10 cm from the participant's eyes, precisely aligning it with the centre of the display. This spatial configuration was designed to isolate the pupillary accommodation reflex by eliminating horizontal or vertical eye movements during fixation shifts. The subject was instructed to remain still during acquisitions.
- **System initialisation:** After entering the subject's identifier on the dedicated GUI welcome screen, the investigator initialised the acquisition device. By utilising the Gazepoint Control software interface, it was possible to place the remote camera at the optimal operational distance and angular orientation relative to the individual's face (target range: 40–60 cm) and verify robust binocular pupil segmentation before proceeding. Before starting the formal acquisition session, the subject was asked to shift their focus from the screen to the near target to provide an initial qualitative confirmation of voluntary PAR elicitation. Furthermore, the investigator encouraged the user to perform multiple successive miosis-mydriasis cycles until they felt fully confident with the gesture. The objective of this phase was to calibrate the sensitivity of the T_{SMA} to the individual's specific physiological response using an iterative, trial-and-error approach. By utilising the Training widget (Application 1),

the participant completed a guided session that prompted three Type 1 PAR and two Type 2 PAR events in a randomised sequence. The application generated a final performance score metric: $S = \text{successfullydetected}/\text{prompted}$. Based on these initial metrics and a visual analysis of the signal telemetry, the investigator adjusted the sensitivity multiplier parameter (*Soglia*) as needed before executing a new session. This optimisation loop was repeated until an optimal T_{SMA} sensitivity configuration was identified that maximised the performance score.

- **“Training” module:** The participant completed three consecutive, standardised Training sessions. This step was necessary to establish a homogeneous dataset for subsequent statistical analysis and, as the activity’s name suggests, to increase the user’s confidence in the PAR mechanism.
- **“Yes or No?” module:** Functional validation of Application 2’s binary selection interface was assessed by presenting the subject with two personalised questions with unambiguous answers (one expecting a “Yes” response, the other a “No”). Additionally, the participant was instructed to access the pause feature via a Type 2 PAR execution and, if successful, autonomously navigated back to the Main Menu.
- **“Keyboard” module:** The user initialised Application 3 and was permitted to type either a self-selected string or the standard freely suggested word (“TEST”). They were instructed to interact with all the extra functionalities, including the predictive text suggestions layout, the multiple character deletion (“Full-canc”), and the pause-to-exit trigger. Because this interface layout utilised the matrix-scanning algorithm and required three distinct types of PAR, the task introduced significantly higher cognitive load than the previous modules in this protocol. Consequently, quantitative performance metrics were not recorded for this application, and evaluation was restricted to qualitative accessibility training.
- **“Shuttle Game” module:** As the final operational block of the experimental acquisition, the participant interacted with the arcade mode featured in Application 4. To accommodate the fast-paced game mechanics and maximise success rates, the investigator typically increased the sensitivity threshold by 0.05–0.1. Participants could play multiple rounds, and were encouraged to try again to improve their score.
- **Psychometric Survey:** After concluding the acquisition protocol, each participant completed a dedicated digital questionnaire. The primary objectives were to document general demographic parameters (age, gender) and to collect quantitative data on perceived user experience (UX), cognitive workload, and physical ergonomics.

3.3 First trial involving an ALS-LIS Patient

Following the overall successful validation phase, *Pupil-com* was tested by one subject (male, 67 years old) affected by MND since 1994. At the moment of the experiment, the patient had been tetraplegic and artificially ventilated and fed for 16 years, but presented with intact eye movement and some facial expressions control. He communicated fluently with a standard eye-tracking device. The trial was conducted in a home care environment, in the presence of the caregiver, one speech therapist and two researchers. The study followed a protocol specifically adapted from the previous validation phase for system performance assessment with a target user. Particular focus was put on the “Yes or No?” application due to its ease of use and adaptability for basic communication. The experiment proceeded as follows:

- **Setup Preparation in the Homecare Environment:** The portable device setup was adapted to the patient’s home, utilising a standard external display as the interface screen.

All components listed in Chapter 2 were arranged according to the available space and resources.

- **Introductory Briefing:** The patient and caregiver received a detailed explanation of PAR, including methods for elicitation and its use as an interface command. Prior experience with other AAC devices facilitated rapid comprehension of the system.
- **Acquisition Launch:** The experimenter initialised *Pupil-com* with an anonymous session to ensure the subject's complete privacy, according to the Ethical Committee rules.
- **PAR Qualitative Verification:** After ensuring that the user's eyes were correctly detected by the tracker, the investigator invited the patient to shift his focus from the screen to the near target. This verified the correct voluntary elicitation of PAR before the official acquisition session began.
- **Iterative Threshold Calibration and "Training" module:** To determine the optimal T_{SMA} value, consecutive "Training" exercises were performed, each prompting three Type 1 PARs and two Type 2 PARs in a randomised sequence, and the sensitivity value was adjusted as needed.
- **"Yes or No?" module:** The investigator prepared 10 questions (adapted from [38], displayed in Table 3.1) with obvious "YES/NO" answers that were asked to the user one at a time. For each question, response time and answer correctness were evaluated.
- **Free use of the interface:** The patient was allowed to navigate the interface and try the *Keyboard* and *Shuttle Game* Applications as he preferred.
- **Psychometric Survey:** After concluding the acquisition, an adapted version of the previous validation phase questionnaire was proposed to the patient. The questions focused on the interface's usability and flow, ease of PAR elicitation, and the overall subjective experience of this novel communication approach.

Table 3.1: Questions asked during the trial with one ALS-LIS patient

Question	Expected answer
Is your name <i>subject's name</i> ?	Yes
Is your name <i>wrong subject's name</i>	No
Are we in Australia?	No
Are we in Italy?	No
Are you a man?	Yes
Are you a woman?	No
Is there anyone else in this room?	Yes
Are you alone in this room?	No
Are you over 18?	Yes
Are you under 18?	No

3.4 Data management and analysis

3.4.1 Manual dataset organisation

The first stage of data analysis, regardless of the study’s nature (pilot, validation, patient testing), always aimed to remove redundant information to obtain a clean dataset including only relevant files. Each acquisition session folder, which stores all auto-saved data per experimental session, underwent a manual cleanup: empty folders, duplicates, and irrelevant CSV files or PNG charts from aborted or incomplete sessions were deleted. All remaining data (log file text, CSVs, and the respective PNGs of the various modules utilised during the acquisition protocol) were collected in the subject directory, with the number of sessions required to gather all the information specified (due to external interruptions or unexpected software crashes). Each subject was assigned a unique ID, which was reused if they took part in multiple studies, with a prefix indicating the study. Table 3.2 summarises the subject organisation.

Table 3.2: Summary of subject participation and identification codes across the pilot and validation studies.

Subject	Pilot Study ID	Validation Study ID	Group
Subjects A–H	PS_01, ..., PS_08	VS_01, ..., VS_08	Group A
Subjects I, J, K	PS_09, PS_10, PS_11	<i>Did not return</i>	N/A
Subjects L–P	<i>Not in pilot</i>	VS_09, ..., VS_15	Group B

3.4.2 Feature extraction

Raw telemetry from CSV files and operational TXT logs were parsed and synchronised to extract subject-specific performance metrics, algorithm classification parameters, and scores representative for each different *Pupil-com* application.

Because the preliminary pilot study used a simpler architecture and a classification algorithm that was still evolving, feature extraction for this phase was restricted to core metrics compatible with the subsequent validation study. Conversely, the increased structural and algorithmic complexity of the final version of *Pupil-com* provided more insights: it allowed the extraction of many additional features to evaluate all the newly integrated applications.

To ensure data integrity, the extraction pipeline incorporated rigorous timestamp-based synchronisation between the CSV files and the TXT session logs. While the former were essential to extract time series of instantaneous pupillary area, threshold values T_{SMA} and T_{exit} , and PAR classification, the latter were just as fundamental for a complete understanding of the system’s behaviour and an evaluation of the quality of human-computer interaction, since they recorded all the state-machine transitions, user prompts and subsequent command triggered. This double-source data mining enabled accurate isolation of the physiological reaction times (metrics available only after one patch release, mid-validation study), responses in the *Yes or No?* widget and the entries evolution while using the *Keyboard*, as well as the successful use of extra functionalities, such as the predictive Suggestions and the "Pause and Exit" option. It was important to track not only successful interactions, but also the accidental secondary command triggers (Type 2 and Type 3) to evaluate the impact on usability of the command overlap – i.e. misrecognition of the PAR type because they only differ in dwell-time, and it may be difficult to have full timing control, especially during first or second use.

Table 3.3 outlines the specific metrics extracted across both pilot (PS) and validation (VS) phases, along with their operational definitions.

3.4.3 Survey processing

The subjective and qualitative data gathered through a post-acquisition survey were analysed to complement the quantitative telemetry features. The questionnaire combined standard Likert-scale ratings (1–5) and open-ended responses, allowing participants to share their personal impressions. The questions focused on the following topics:

- **Ergonomics and comfort:** assessments of physical constraints, posture during acquisitions, and potential ocular fatigue.
- **System usability:** evaluation of the intuitiveness of the PAR mechanism, ease of voluntary actuation, and the overall responsiveness of the application to user commands.
- **Application-specific user experience:** analysis of participants’ perceptions of the interface features—such as training usefulness to familiarise with the PAR gesture, options-scanning speed in the Main Menu and Yes or No?, difficulty of Shuttle Game gameplay, and the impact of keyboard typing flow and predictive suggestions. Participants expressed their preferences regarding usability, as well as opinions and suggestions for future improvements.

All Likert-scale ratings were extracted and processed from an .xlsx file, while open-ended feedback was manually inspected to enrich the discussion phase.

3.4.4 Statistical analysis

The final stage of the data pipeline consisted of statistical processing of metrics extracted from the acquisitions and the post-experiment survey. This stage evaluated overall system performance, reliability, and usability. All statistical analyses were conducted using MATLAB [50] (MathWorks, UK). A Shapiro–Wilk test was applied to each variable to assess normality. Descriptive statistics were reported based on distribution characteristics. Normally distributed metrics were summarised as mean $\pm$ standard deviation, while non-normally distributed data were reported as median $[IQR]$ values. Subsequent analyses examined performance improvements and the potential presence of a learning curve across consecutive trials. The investigation was organized as follows:

- **Inter-study and intra-group comparison:** Core performance metrics common to both the Pilot and Validation Studies were compared quantitatively to assess improvements and robustness introduced by the finalised *Pupil-com* framework. Since the upgrade encompassed both hardware and software changes, and eight participants had prior experience with the interface during the Validation Study, observed changes were attributed to a combination of system and participant-specific factors. Consequently, no formal inferential tests were conducted between these two studies or between the returning participants’ first and second acquisitions.
- **Inter-group comparison:** Performance metrics for Group A (returning participants) and Group B (naive participants) in the Validation Study were compared to determine the statistical significance of prior experience in PAR actuation and gameplay. Independent-sample statistical tests (significance threshold $\alpha = 0.05$) were applied to Training and Shuttle Game scores from the Validation Study. Due to differing sample sizes ($N_{GroupA} = 8$ and $N_{GroupB} = 7$), the Mann–Whitney U test was used for non-normally distributed data, and the Welch t-test was used for normally distributed data.

Table 3.3: Summary of experimental metrics evaluated across the *Pupil-com* applications.

Metric	Operational Definition	Study
Training Application		
Absolute score (out of 5)	Total displayed at session termination	PS & VS
Subject's score (out of 5)	Number of physiologically valid PARs executed by the user	PS & VS
Algorithm score	Absolute algorithm accuracy computed relative to the subject's score	PS & VS
Average subject's reaction time	Average of latency (seconds) between prompt and PAR actuation	VS
Optimal threshold	Manually adjusted optimal T_{sma} sensitivity	VS
Game Application		
Subject's best score	Maximum game score achieved across all sessions	PS & VS
Optimal threshold	Average of manually adjusted optimal T_{sma} sensitivity	VS
Yes or No? Application		
Answer success	Intended Type 1 PAR to answer	VS
Pause and exit success	Intended Type 2 PAR to pause and exit	VS
Accidental answer	Unintended Type 1 PAR	VS
Accidental pause	Unintended Type 2 PAR or unsuccessful pause and exit	VS
Keyboard Application		
Total keyboard interactions	Character/suggestion selection or "CANC" actions	VS
Correct keyboard interactions	Intended character/suggestion selection or "CANC" actions	VS
Secondary command trigger	Type 2 and Type 3 PAR to trigger extra commands	VS
Accidental secondary command trigger	Unintended Type 2 and Type 3 PAR	VS
Predictive suggestions correctly utilised	Number of correct suggestions prompted	VS
Pause and exit success	Task 3 PAR execution success to pause	VS

Chapter 4

Results and discussion

The quantitative and qualitative metrics extracted from all experimental phases (Pilot Study and Validation Study on healthy subjects, alongside the clinical case study involving an ALS-LIS patient) were essential for assessing and confirming the overall usability and compatibility of Pupil-com for autonomous communication. These metrics also highlighted the challenges users may encounter. Specifically, the data underscore the learning curve required to achieve full control of the BCI. They also identify architectural improvements to enhance the user experience. This chapter aims to present the results, discuss their impact on the system’s characterisation, and propose solutions to the bottlenecks encountered.

4.1 System evolution and robustness

The features extracted from both the PS and VS experimental phases were processed and evaluated using descriptive and inferential statistics, as detailed in Chapter 3. An initial comparison of metrics common to PS and VS studies evaluated the state-of-the-art framework’s accuracy in PAR recognition.

The prototype architecture underwent multiple changes between Pilot and Validation studies, including the acquisition module (from wearable to remote) and the PAR detection algorithm (from simple moving average to adaptive double-thresholding); for this reason, it is not possible to perform rigorous statistic comparison between the two studies’ results, rather the PS findings provide a baseline reference, validating any improvement reported by the VS.

Classification accuracy evaluation was the primary objective of the comparative analysis: the *Training* application guided the subject through five consecutive PAR events; it provided a final score, on a 0–1 scale, (PS 0.8 [0.31] vs VS 0.9 [0.2]) reflecting the algorithm’s recognition capability, though this was highly dependent on the subject’s responsiveness. Figure 4.1 shows the example of a PAR missed by the subject, who was not ready to elicit one among the five requested: this resulted in an overall system score of 0.8, while the algorithm’s score for this trial should have been equal to 1, with 4 PAR events detected out of 4 correctly performed by the subject.




Figure 4.1: Example of missed PAR: the subject, due to teary eyes, had to close their eyelids for longer than one second, resulting in signal loss. When a new constriction was requested, they were not ready to respond. Close up of a *Training* acquisition

To overcome these discrepancies, evaluating the subject-specific performance with post-processing tools enabled the isolation of a univocal metric for the algorithm's true detection rate. With this metric, the Validation Study exhibited a 15% increase in median accuracy (0.95 [0.19] vs 0.8 [0.26], Figure 4.2). This finding suggests that the current system architecture introduced a substantial overall improvement in PAR recognition, compared to the Pilot Study framework.




Figure 4.2: Inter-study comparison of true algorithm detection rates distribution.

4.2 Impact of user experience

The definitive *Pupil-com* framework, as documented in the previous section, demonstrated high accuracy and efficiency for users of all experience levels. However, inter-study and intra-group analyses suggested that returning subjects, who are already familiar with PAR elicitation, perform faster and more stable intentional pupillary constrictions. In fact, the intra-group statistical test performed on average reaction times showed that returning subjects responded significantly faster than naive subjects (Figure 4.3 Group A, 1.11 [0.077] s, vs Group B, 1.33 [0.607] s, $p=0.0037$). This finding solidly confirms that practice increases the user’s familiarity with the PAR gesture. In particular, it enables considerably faster response already from the second use, promising increased interaction speed and communication rates with very little training.




Figure 4.3: Intra-group average reaction times comparison. Group A (returning) subjects exhibited significantly faster responses than Group B (naive).

Furthermore, quantitative comparisons and inferential tests performed on the *Shuttle Game* best scores (the number of planets the user managed to hit before Game Over) supported the hypothesis above. A first PS-VS comparison revealed a significantly higher median best score during the Validation Study (VS: 8.00 [6.50]; PS: 3.5 [4.00]). The large IQRs, indicating high variance among subjects, suggest that game results are highly dependent on user skill and timing. Subsequent longitudinal comparison of Group A’s best scores across the first and second experimental phases showed a significant increase in gameplay outcomes (PS 4.00 [3.25] vs VS 9.00 [4.00]) and a reduced IQR in VS compared to the entire sample. These findings may imply that experience enhances the subject’s responsiveness and synchronisation to the game dynamics. Nevertheless, they are also highly dependent on the different setup architectures and software, and significant improvements can only be investigated via intra-group analysis within the validation framework. Although the increased performance of Group A compared with Group

B was marginally significant ($p=0.087<0.1$), the median and IQR values (Group A: 9.00 [4.5]¹ vs Group B: 4.00 [5.5], Figure 4.4) suggest a clear trend toward greater confidence among returning subjects and promising scores among first-time users.




Figure 4.4: *Shuttle Game* best score distribution for Group A (returning) than Group B (naive) subjects.

4.3 Tuning and interface usability

Manual parameter tuning is one of the core strengths of *Pupil-com*, enabling experimenters and caregivers to fully adapt the system to the user’s pupillary physiological response and specific needs.

The essential adjustable parameter that ensures flexibility of use among all subjects is the multiplier associated with the T_{SMA} threshold. The optimal threshold value is considered optimal when the system recognises PAR events with high accuracy and near-zero false positives; during the experimental protocol, it was determined via iterative manual calibration. Figure 4.5 show how the correct threshold value impacts on miosis recognition accuracy during *Training*. Its value was almost unanimously required at 0.85, with a few exceptions. The result strongly suggests that this value should be set as the default at *Pupil-com* activation and adjusted only if necessary. Interestingly, the optimal threshold for successful gameplay showed higher median values compared to the other GUI applications (0.90 [0.05]): the subjects experienced difficulty

¹The discrepancy in the IQR reported for Group A’s VS scores (4.00 in the inter-study vs 4.50 in the inter-group analysis) is not an error. Due to a data-saving failure, subject PS_07’s game score is missing. To maintain data alignment for paired statistical testing, this subject was excluded from the intra-group comparison but included in the independent inter-group analysis, as their Validation Study score was correctly recorded.




Figure 4.5: Importance of threshold multiplier setting: in the plot above, an example of too low T_{SMA} multiplier (0.85), that caused one Type 1 and one Type 2 misses; in the plot below, the optimal value for this subject (0.9).

in Type 1 PAR elicitation while controlling the game interface, and an increment of 0.05–0.1 (Figure 4.6 for a distribution comparison) had beneficial effects in false negatives reduction and final score increase.

The applications newly integrated in the definitive *Pupil-com* structure provided additional insights into usability and features that may need improvement. The *Yes or No?* widget, based on a binary choice, and the *Keyboard*, which relies on multi-symbol communication, explored system usability and evaluated the integration of Type 2 and Type 3 PAR as commands.

From the *Yes or No?* validation protocol, it emerged that all the 15 subjects could answer at least one of the two questions correctly, while 12 of them responded to both. The Type 2 command for this application was successfully used by 13 subjects, but 11 of them navigated back to the Main Menu. A specific investigation into accidental triggers (unintended "YES"/"NO" responses or "RESUME" selections instead of "EXIT") while attempting to complete this final task revealed that the majority of subjects had at least one false positive before succeeding. Figure 4.7 boxplots show that those who managed to pause and exit the application had lower median accidental triggers compared to those who failed this task; this generally high error

rate, though, indicates clearly higher difficulty in performing Type 2 PAR, compared to Type 1. A retrospective code analysis revealed a bug, isolated within the *Yes or No?* widget and not affecting any other functionalities of *Pupil-com*, which interfered with optimal PAR recognition: the **ConstrictionMonitor** was inadvertently processing the incoming GP3 raw data, bypassing the regular smoothing and artifact removal pipeline of the **AreaFilter** module. With no control over signal loss from natural blinking or sudden movements, the algorithm could not ensure stability, especially during prolonged PARs: a single frame lost during a pupillary constriction would cause misclassification of the entire event.




Figure 4.6: Optimal threshold comparison: *Training* vs *Shuttle Game*




Figure 4.7: Accidental triggers during *Yes or No?* pause attempt: exit fail vs exit success distribution.

An analogous observation was applied to the *Keyboard* application: for each participant, an Interaction Success score was calculated as $IS = \frac{\text{correct interactions}}{\text{total interactions}}$, and a Character Success score as $CS = \frac{\text{correct characters typed}}{\text{all characters typed}}$. These two performance values reflect the high complexity of this interface ($IS = 0.69 \pm 0.27$, $CS = 0.85 [0.36]$) that, compared to all the others, has a higher cognitive load and, due to its matrix-scanning sequence, makes it easier to unintentionally select the wrong character. Additionally, the evaluation of accidental vs intended commands while

attempting to use the predictive suggestions (Type 2 PARs) or to pause and exit the application (Type 3 PARs) found that 63.5% of the commands were not intentional, due to false positives or PAR Type misrecognition.

Taken together, all these results suggest that *Pupil-com*'s critical issues are mostly related to secondary commands that must be triggered with long and extra-long PARs. As evidenced by the developer's ability to operate all command types with minimal errors, only further investigation will clarify whether this command-overlap effect can be reduced solely by user expertise, or whether algorithmic optimisation would be necessary to achieve near-optimal performance.

4.4 User feedback

The post-acquisition questionnaire distributed to the validation cohort ($N = 15$) provided critical insights into *Pupil-com*'s usability and UX. The subjective evaluations, presented extensively in Table 4.1, indicate an overall positive experience, validating the system's ergonomic design (5 [1]) and responsiveness (4 [1]).

A key finding emerged regarding the intuitiveness and ease of use of the BCI: while subjects reported unanimous theoretical understanding of the PAR elicitation mechanism (5 [0]), the practical execution of a voluntary constriction showed a wider inter-subject variability (4 [1.5]). This discrepancy confirms the necessity of the dedicated *Training* application, which was evaluated as essential (5 [1]). It is interesting to highlight a specific trend observed when comparing this score between naive (Group B) and returning subjects (Group A): the former reported a higher median rating (5 [0]) than the latter (4 [1.25]), suggesting that first-time users heavily rely on the initial guided familiarisation to bridge the gap between understanding and successfully actuating a PAR gesture.

Furthermore, the data highlight that prolonged use may cause visual fatigue: although overall visual fatigue was rated moderately low (4 [1.5]), isolated reports of epiphora (tearing eyes) in two subjects suggest that session durations must be limited if the user experiences ocular strain.

Regarding the application-specific modules, the survey responses effectively mirrored the distinct cognitive demands of each interface. The *Yes or No?* application scanning speed was perceived as optimal with a tightly clustered consensus (3 [0]). Conversely, the *Keyboard* app was reported to have significantly lower fluidity and control (3 [1]), a result that can be attributed to the complex layout and the bottlenecks encountered when accidentally selecting the wrong row or column. In this context, the high cognitive load was only partially mitigated by the predictive suggestions, as the subjects reported that the execution of advanced features—accessible via Type 2 and Type 3 PAR events—was not easy (3 [2]).

Finally, while the quantitative results for the *Shuttle Game* application suggest an optimally balanced initial speed (3 [0.5]) and progressive level difficulty (3 [1]), direct observation and open-ended feedback revealed moments of user frustration while playing the game. This was primarily triggered by perceived false negatives from the PAR recognition algorithm: under the pressure of a dynamic screen and strict timing requirements, subjects struggled to produce voluntary pupillary constrictions of sufficient amplitude. Although this issue was temporarily alleviated during the sessions by manually raising the threshold (T_{SMA}), it highlights the distinct physiological and cognitive challenges introduced by real-time, dynamic tasks compared to static scanning interfaces.

Table 4.1: Summary of subjective usability survey results.

Survey Item	Evaluation Scale (1 – 5)	Median [IQR]
General & Training		
System setup comfort	1=Very uncomfortable; 5=Very comfortable	5 [1]
Visual fatigue during use	1=High; 5=Low	4 [1.5]
Understanding of PAR mechanism	1=Difficult; 5=Easy	5 [0]
Training phase utility	1=Useless; 5=Essential	5 [1]
PAR execution difficulty	1=Difficult; 5=Easy	4 [1.5]
System responsiveness	1=Poor; 5=Excellent	4 [1]
Game Application		
Perceived game speed	1=Slow; 5=Fast	3 [0.5]
Level progression difficulty	1=Difficult; 5=Easy	3 [1]
Pause execution difficulty	1=Difficult; 5=Easy	3 [1.5]
Yes or No? Application		
Perceived scanning speed	1=Slow; 5=Fast	3 [0]
Keyboard Application		
Perceived scanning speed	1=Slow; 5=Fast	3 [0.5]
Predictive suggestions and exit difficulty	1=Difficult; 5=Easy	3 [2]
Typing fluidity and control	1=Poor; 5=Excellent	3 [1]

4.5 Testing on an ALS-LIS patient

The final stage of the experimental phase of this work consisted of a one-patient trial, performed in their home-care environment. The protocol for patient testing outlined in Chapter 3 was followed almost entirely, providing essential feedback on *Pupil-com*’s real-world usability for a target patient interacting with it for the first time. The system’s portable nature allowed the researchers to adapt the setup to the bedridden patient’s room, utilising their own television screen as the far target, the interface support, and audio output. The most challenging part of the setup preparation was the near-target positioning: the same laboratory stand used during the validation study on healthy subjects was placed on an overbed tray table and adjusted to the patient’s gaze; however, perfect alignment could not be achieved, as positioning the target too precisely would have caused it to enter the eye-tracking field of view and interfere with accurate pupillary detection. In this case, the patient still had full control of the extraocular muscles and presented no nystagmus or fixation problems; despite the slight axis mismatch, the user successfully elicited PARs, switching focus from the far to the near target with additional eye movements. However, this solution would not be suitable for patients in a CLIS state, and alternative target holders must be engineered to ensure the device’s universal usability.

Since the first *Training* session (score: 4/5), the subject showed complete understanding of the PAR mechanism: the system correctly detected the three Type 1 events, while one of the Type 2 constrictions was classified as Type 1 because of the subject’s delayed response, and the other one was successfully detected. A repositioning of the targets and the eye-tracking device was needed because the system experienced multiple signal losses when the subject elicited a PAR (as indicated by the flat lines observed from $t=28$ s in Figure 4.8, plot above). The second trial resulted in 5/5 constrictions recognised, and it was decided not to administer additional *Training* sessions to avoid excessive eye fatigue. Figure 4.8, plot below, represents the results of this second acquisition: the greater pupillary area span between relaxed and constricted states proves that the optimal target distances were found. By observing these two plots, one may argue that, in addition to the true positives, the system produced a discrete number of false positives, indicated by the black dots (Type 1) and squares (Type 2), which would compromise interface responsiveness during actual use. However, each Type 1 request was set to last longer than 3 seconds, allowing users sufficient time to respond while still enabling reliable detection of pupillary constriction. In this context, the occurrence of a Type 2 response during a Type 1 request does not constitute a false negative; rather, it may reflect a rapid Type 1 response that remains below the detection threshold throughout the request. This scenario can lead to the system erroneously registering a “phantom” Type 2 event, even though the underlying physiological response aligns with a fast Type 1 constriction. Moreover, it is worth noting that the observed Type 1 false positives in the first trial are likely attributable to the suboptimal positioning of the near-target, which made fixation unstable and challenging for the user. In contrast, the Type 1 false positives detected during the second trial can be explained by the subject’s anticipatory behaviour; these events both occurred approximately one second before the request was vocalised, suggesting that the subject attempted to predict and respond to the upcoming prompt before its actual onset.

The core activity of the patient trial was the *Yes or No?* application testing. Out of 10 questions with known answers, the patient provided 8 correct responses. Wrong answers were given to questions 6 and 10, both of which expected “NO” as the answer and the response time was 10.12 s and 2.33 s, respectively. Moreover, the average response time was 5.56 ± 3.21 s across 9 answers, excluding the outlier question 5 that took 31.12s. It is important to note that the bug previously discussed for the Validation Study was still present during the patient trial. Both real-time and post-processing observations showed that the subject attempted to answer correctly multiple times before succeeding, but, since the **ConstrictionMonitor** was operating on the raw signal, with no blink or motion artifact removal, not all the PARs were detected. This issue significantly increased the perceived difficulty of use and response time of the interface to commands, leading to false negatives and subject’s frustration. However, the trial results are even more encouraging: the achievement of an 80% success rate with raw data demonstrates the




Figure 4.8: Patient *Training* plots, relative to first trial (above) and second trial (below)

ConstrictionMonitor pipeline is solid and the performance is expected to improve substantially with smoothed data.

Subsequently, the patient interacted with the *Shuttle Game* module. After a short familiarisation with the game layout and dynamics, he successfully completed a game: he launched the spaceship multiple times (Figure 4.9), hitting 3 moving planets in total, but just slightly missing many others, simply because he was not perfectly synchronised with the game. With continued practice, he could easily surpass this and achieve much higher scores.

After gameplay, because the patient exhibited signs of fatigue, the *Keyboard* widget was presented for informational purposes only, but no typing attempts were made.

After completing the trial, the patient and his caregiver provided immediate verbal feedback and completed the user satisfaction questionnaire. They assigned an overall subjective rating of 4/5 to the communication system, indicating the perceived potential of this new solution, and declared they would try it again in the future. Regarding general usability metrics, the patient reported a comfort level of 3/5 (not too comfortable, nor too uncomfortable), and visual strain was rated 4/5 (almost no ocular fatigue). However, the interface control and the perceived importance of the initial training session were rated as poor (1/5). In retrospect, the patient succeeded in the training session almost immediately, fully understood the PAR mechanism (rated 4/5), and likely found the guided exercise session redundant. The low rating for interface control may reflect the time required and the trial-and-error process involved in finding the optimal setup, as well as the novelty of the communication system. Such initial challenges are to be expected when adapting to a new interface that is substantially different from the traditional eye-tracking device he is accustomed to for everyday communication, and only prolonged use may improve familiarity and proficiency.




Figure 4.9: Patient *Shuttle Game* trial, 15–45 s close-up. Each red marker represents a successful spaceship launch. The vertical red line represents a successful collision with a planet.

Chapter 5

Conclusions and future directions

This Master’s thesis objective of restoring communication with completely paralysed individuals affected by ALS was achieved by developing and testing the AAC device *Pupil-com*. To overcome the inherent limitations eye-tracking-based systems encounter after the patient’s loss of oculomotor control, the system was based solely on live pupillometry. This innovative BCI enabled rapid interaction with a multiple-choice scanning interface, by exploiting the Pupillary Accommodative Response (PAR), an autonomic nervous system reflex spared by the disease’s progression and elicitable voluntarily.

Pupil-com’s definitive design combined legacy device features with new technologies, primarily to achieve robust signal performance, autonomous operation with minimal external intervention, and tunability to the user’s physiological response. This was achieved by integrating a low-cost commercial eye tracker into the data acquisition pipeline; compared with the previous custom-made acquisition setups, it provided significantly more stable pupillometry, thanks to the embedded infrared illumination and camera, as well as to the underlying powerful image processing algorithm for live pupil segmentation.

Portability and adaptability to any home or clinical environment were ensured by selecting a fast mini-PC with standard audio and video outputs, a cost-effective solution that maintained high performance. *Pupil-com*’s interface runs on a custom-built Python platform, and it comprehends four main applications: *Training*, for initial familiarisation with the PAR via an exercise guiding gaze depth changes between a far and a close target; *Yes or No?*, for binary responses to simple questions; *Keyboard*, for more complex communication, with the possibility to write full sentences and with the help of predictive suggestions; *Shuttle Game*, for light-hearted yet challenging entertainment. The interface appearance was developed to be as minimal as possible, thus allowing for immediate understanding and ease of use. In addition, to enhance the user experience, it was complemented with auditory feedback and vocalisation of selected items, as well as a set of manually adjustable parameters, such as the threshold settings for optimal detection of PAR events.

Experimental trials conducted on healthy volunteers validated the general framework’s usability; in particular, they demonstrated the robustness and effectiveness of the double-thresholding PAR detection algorithm. They underscored its high accessibility as well, since subjects could be operate it easily even with little or no training. The encouraging results and positive feedback from participants opened the possibility of presenting the device to target users. A single-case study involving a target patient in a locked-in state allowed an initial investigation of the device’s performance in a real-world setting. Despite difficulties with near-target positioning and general

setup comfort, the participant demonstrated complete understanding and control of the PAR voluntary mechanism, successfully responded to binary-choice questions, and briefly played the video game. Feedback from the patient after the trial session revealed strong potential for this novel communication system and expressed interest in participating in upcoming experimental phases.

This work represents a significant advance in the field of assistive communication technologies, marking the first successful coupling of robustness and user comfort with an architecturally complex interface that offers multiple functional applications. Prior prototypes, developed by this research group and previous Master’s thesis projects, demonstrated the feasibility of communication but remained limited in terms of readiness, usability, and functionality. This system, instead, achieved reliable performance and practical usability for both healthy volunteers and an ALS patient in a real-life context.

One relevant feature is the integration of prolonged PAR activation, unlocking advanced options within the interface, such as pausing an application and accepting predictive suggestions. While the algorithm proved highly effective in distinguishing PARs of different duration during guided sessions, the experimental phase revealed that maintaining a prolonged constriction during autonomous use was generally difficult. It is important to determine whether the observed limitations are mainly due to insufficient user training or whether the algorithm itself requires further refinements to better detect longer PAR events. However, the investigation on healthy subjects highlighted the importance of even a short training phase. The current sample size, though, is too limited to support definitive conclusions. Therefore, further longitudinal studies with larger groups would be necessary to state clearly the effects of user familiarisation. In this respect, integrating other video games into the user interface may be an effective strategy to increase patient engagement and facilitate effortless training.

The promising results and overall positive feedback reinforce the potential of this novel communication system. Nevertheless, *Pupil-com* is far from perfection, and further improvements are necessary to increase usability and meet the target user’s real needs.

From the system organisation point of view, the still non-rigorous setup for a home-care environment is still a major challenge: a practical solution for target positioning that recreates the controlled laboratory environment must be found to ensure signal stability and reliable responses.

From a software point of view, the enhancement of the *Keyboard* application is a priority, since it represents the fundamental application to restore communication with patients in a completely locked-in state; its layout must consider that CLIS patients do not have control over extraocular muscles, thus may not be able to identify all the keys that currently span the whole screen size. A possible solution would be to display a few characters at a time, in the centre of the screen, where the patients direct their gaze – presumably corresponding to the foveal area –, or to have the GUI vocalise the currently highlighted options. The integration of machine learning into the current predictive suggestions algorithm would represent a leap forward in usability and communication speed, as the system could be continuously trained to predict and suggest entire sentences or expressions commonly used by the patient, thus bypassing long, even frustrating typing delays. The speech synthesiser itself could be updated to vocalise keyboard entries only when the user wants it to, avoiding repetition of every new key press.

One more useful improvement would be to replace the currently manual threshold setting for PAR detection with an automatic calibration process. By monitoring a set of consecutive guided pupillary constrictions, the system could automatically identify the optimal multiplier for the adaptive thresholding algorithm. Moreover, background monitoring during prolonged use could account for slow drifts or changes in the patient’s responsiveness, such as fatigue or variations in ambient luminance, although the current T_{SMA} implementation already adapts to slow pupil size fluctuations.

Finally, from the testing perspective, the implementation of longitudinal studies on a larger group of ALS patients is strictly necessary. This could enable research to understand the intra-patient physiological variability of PAR and its perceived complexity as an interface command, and ultimately refine *Pupil-com*’s setup for the target users.

The current framework paves the way for these crucial future investigations: the device is intuitive, and acquisition is easy to initialise, enabling extensive at-home use without expert intervention. *Pupil-com* could be mounted at the patient's home and entrusted to the caregiver for multiple days. The patient could perform daily trials at their own pace, with no external pressure, while the recorded sessions could be accessed remotely by the researchers. By remotely monitoring the log files and the pupillary area recordings, it would be possible to track the patient's commitment and actual use of the device, as well as the changes in the PAR quality and responsiveness in relation to the progression of the disease.

This Master's thesis work finds its conclusion in a functioning prototype that makes a meaningful contribution to the field of Alternative and Augmentative Communication devices for individuals with severe motor impairment. *Pupil-com* demonstrates that, through a multidisciplinary engineering framework integrating physiology, optics, electronics, and programming, restoring a voice to those who have lost it is a goal well within reach.

Appendix A: Eye-Trackers Comparison

Feature	Gazepoint GP3	Pupil Core
Setup type	Remote, tripod mount	Wearable
Eye-tracking technology	Bright pupil technique	Dark pupil technique
Binocular eye tracking	Binocular and monocular	Binocular and monocular
Cameras	0.3 MP, 60 fps	192x192 px, 200 fps
Illumination	Multi-IR (860nm) LED array	2x IR LEDs (860nm) per eye
Pupil measurements	Position and pupil diameter in camera pixels (or in mm, after calibration)	2D: position and pupil diameter in eye camera pixels 3D: eye model parameters and pupil diameter in mm using 3D eye mode (unavailable for monocular use)
Cables	2x USB-C	1x USB-C
Computing device	Laptop or desktop running Gazepoint Control. CPU required: Intel i5 or greater RAM required: 8GB or more OS: Windows 10 or greater	Laptop or desktop running Pupil Core software. CPU required: Intel i7 or greater RAM required: 8GB or more OS: Windows 10 or greater
Pupil data	Gazepoint Analysis software. Real-time API: Gaze data and pupillometry data over local network.	Pupil Capture software. Real-time API: Gaze data, pupillometry data, and videos over local network.
Dimensions	Weight 118 g L 234 x D 47 x H 43 mm	Weight 22.75 g L 160x D 175x H 51 mm
Eyewear compatible	Yes	No
Working distance	50-80 cm	<5 cm
Price	855€	2995€

Table 1: Comparison of Eye-Tracking Devices (Part 1)

Feature	Smart Eye AI-X	Neon
Setup type	Remote, mounting plate	Module mounted on glasses frame
Eye-tracking technology	Dark pupil technique, corneal reflection	NeonNet (real-time neural network)
Binocular eye tracking	Binocular	Binocular and monocular
Cameras	60 fps	2 x IR eye cameras 192 x 192 @200 Hz
Illumination	Not specified	-
Pupil measurements	Position and pupil diameter in camera pixels (with associated quality index)	3D position and pupil diameter in mm
Cables	1xUSB 3.0	1x USB-C
Computing device	OS: Windows 10 or greater	Mobile computing device: Samsung S25 OS: Windows 10/macOS 13/Android 13
Pupil data	Supports iMotion (metrics analysis third-party software) or custom software with special API needed	Pupil Companion for real-time Pupil Cloud for post-hoc
Dimensions	Weight 147 g L 268 x D 32 x H 13 mm	Weight 7,3g L 40 x D 10 x H 35 mm
Eyewear compatible	Not specified	Yes
Working distance	45-85 cm	<5 cm
Price	1200€	5515€

Table 2: Comparison of Eye-Tracking Devices (Part 2)

Feature	AdHawk MindLink	Tobii Pro Glasses 3
Setup type	Wearable	Wearable
Eye-tracking technology	Laser reflection geometry	Corneal reflection, dark pupil technique, stereo geometry
Binocular eye tracking	Binocular and monocular	Binocular and monocular
Cameras	Camera-free, MEMS-based measurement @500 fps	2xIR sensors per eye 50 - 100 fps
Illumination	-	-
Pupil measurements	Position and pupil diameter in mm	3D position and pupil diameter in mm
Cables	1xUSB 3.0/1xUSB-C	1xUSB 2.0
Computing device	PC or mobile running the owner control app. CPU required: Intel i5 or greater RAM required: 8GB or more OS: Windows 10/macOS 13/Android 13	PC or mobile running D-LAB software. CPU required: Intel i7 or greater RAM required: 8GB or more OS: Windows 10
Pupil data	Custom SDK protocol	Controller app for real-time acquisitions. Tobii Pro Lab for post-hoc
Dimensions	Weight 27 g L 166 x D 146 x H 42 mm	76.5 g L 153 x D 168 x H 51 mm
Eyewear compatible	Yes	No
Working distance	<5 cm	<5 cm
Price	10000€	15000€

Table 3: Comparison of Eye-Tracking Devices (Part 3)

Feature	Dikablis Glasses 3	EyeLink Portable Duo
Setup type	Wearable	Remote, multiple mounts
Eye-tracking technology	Corneal reflection, dark pupil technique	Corneal reflection
Binocular eye tracking	Binocular and monocular	Binocular and monocular
Cameras	648x488 px 60 fps	Remote IR sensor up to 2000fps
Illumination	IR LEDs	-
Pupil measurements	Position and pupil diameter in mm	Position and pupil diameter in mm
Cables	1xUSB (cable acquisition) Battery powered router (wireless acquisition)	1xUSB 2.0
Computing device	Custom host PC included in the purchase	Not provided
Pupil data	D-LAB measurement and analysis software. No possibility for custom socket-level programming.	Custom PC interface
Dimensions	Weight 52 g L 164.5 x D 198 x H 58 mm	Not provided
Eyewear compatible	Yes	Yes
Working distance	<5 cm	42-62 cm
Price	25000€	38000€

Table 4: Comparison of Eye-Tracking Devices (Part 4)

Bibliography

- [1] Colin Quinn and Lauren Elman. «Amyotrophic Lateral Sclerosis and Other Motor Neuron Diseases». en. In: *Continuum* 26.5 (Oct. 2020), pp. 1323–1347. ISSN: 1080-2371, 1538-6899. DOI: 10.1212/CON.0000000000000911. URL: <https://continuum.aan.com/doi/10.1212/CON.0000000000000911> (visited on 05/14/2026) (cit. on p. 3).
- [2] Acary Souza Bulle Oliveira and Roberto Dias Batista Pereira. «Amyotrophic lateral sclerosis (ALS): three letters that change the people’s life. For ever». In: *Arquivos de Neuro-Psiquiatria* 67.3a (Sept. 2009), pp. 750–782. ISSN: 0004-282X. DOI: 10.1590/S0004-282X2009000400040. URL: http://www.scielo.br/scielo.php?script=sci_arttext&pid=S0004-282X2009000400040&lng=en&tlng=en (visited on 05/14/2026) (cit. on p. 3).
- [3] Jack N. Alpert. «The Six Major Anatomic Decussations with Clinical Correlation». In: *The Neurologic Diagnosis: A Practical Bedside Approach*. New York, NY: Springer New York, 2012, pp. 127–139. ISBN: 978-1-4419-6724-4. DOI: 10.1007/978-1-4419-6724-4_6. URL: https://doi.org/10.1007/978-1-4419-6724-4_6 (cit. on p. 3).
- [4] Mamede De Carvalho and Michael Swash. «Diagnosis and differential diagnosis of MND/ALS: IFCN handbook chapter». en. In: *Clinical Neurophysiology Practice* 9 (2024), pp. 27–38. ISSN: 2467981X. DOI: 10.1016/j.cnp.2023.12.003. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2467981X23000355> (visited on 05/14/2026) (cit. on p. 4).
- [5] Ryan G. Brotman, Maria C. Moreno-Escobar, Joe Joseph, Sunil Munakomi, and Gauri Pawar. «Amyotrophic Lateral Sclerosis». eng. In: *StatPearls*. Treasure Island (FL): StatPearls Publishing, 2026. URL: <http://www.ncbi.nlm.nih.gov/books/NBK556151/> (visited on 05/14/2026) (cit. on pp. 4, 5).
- [6] Sara Bandres-Ciga et al. «Shared polygenic risk and causal inferences in amyotrophic lateral sclerosis». en. In: *Annals of Neurology* 85.4 (Apr. 2019), pp. 470–481. ISSN: 0364-5134, 1531-8249. DOI: 10.1002/ana.25431. URL: <https://onlinelibrary.wiley.com/doi/10.1002/ana.25431> (visited on 05/14/2026) (cit. on p. 4).

- [7] Radu Eugen Rizea, Antonio-Daniel Corlatescu, Horia Petre Costin, Adrian Dumitru, and Alexandru Vlad Ciurea. «Understanding Amyotrophic Lateral Sclerosis: Pathophysiology, Diagnosis, and Therapeutic Advances». en. In: *International Journal of Molecular Sciences* 25.18 (Sept. 2024), p. 9966. ISSN: 1422-0067. DOI: 10.3390/ijms25189966. URL: <https://www.mdpi.com/1422-0067/25/18/9966> (visited on 05/14/2026) (cit. on p. 5).
- [8] Elisa Longinetti and Fang Fang. «Epidemiology of amyotrophic lateral sclerosis: an update of recent literature». In: *Current Opinion in Neurology* 32.5 (Oct. 2019), pp. 771–776. ISSN: 1350-7540. DOI: 10.1097/WCO.0000000000000730. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6735526/> (visited on 05/14/2026) (cit. on p. 5).
- [9] Orla Hardiman, Ammar Al-Chalabi, Adriano Chio, Emma M. Corr, Giancarlo Logroscino, Wim Robberecht, Pamela J. Shaw, Zachary Simmons, and Leonard H. Van Den Berg. «Amyotrophic lateral sclerosis». en. In: *Nature Reviews Disease Primers* 3.1 (Oct. 2017), p. 17071. ISSN: 2056-676X. DOI: 10.1038/nrdp.2017.71. URL: <https://www.nature.com/articles/nrdp201771> (visited on 06/17/2026) (cit. on pp. 5, 6).
- [10] André Maier et al. «ALSFRS-R-SE: an adapted, annotated, and self-explanatory version of the revised amyotrophic lateral sclerosis functional rating scale». en. In: *Neurological Research and Practice* 4.1 (Dec. 2022), p. 60. ISSN: 2524-3489. DOI: 10.1186/s42466-022-00224-6. URL: <https://neurolrespract.biomedcentral.com/articles/10.1186/s42466-022-00224-6> (visited on 05/14/2026) (cit. on p. 5).
- [11] Marta Gromicho et al. «Spreading in ALS: The relative impact of upper and lower motor neuron involvement». en. In: *Annals of Clinical and Translational Neurology* 7.7 (July 2020), pp. 1181–1192. ISSN: 2328-9503, 2328-9503. DOI: 10.1002/acn3.51098. URL: <https://onlinelibrary.wiley.com/doi/10.1002/acn3.51098> (visited on 05/14/2026) (cit. on pp. 5, 6).
- [12] Joe M. Das, Kingsley Anosike, and Ria Monica D. Asuncion. «Locked-in Syndrome». eng. In: *StatPearls*. Treasure Island (FL): StatPearls Publishing, 2026. URL: <http://www.ncbi.nlm.nih.gov/books/NBK559026/> (visited on 05/14/2026) (cit. on p. 6).
- [13] Malcolm Proudfoot, Ricarda A.L. Menke, Rakesh Sharma, Claire M. Berna, Stephen L. Hicks, Christopher Kennard, Kevin Talbot, and Martin R. Turner. «Eye-tracking in amyotrophic lateral sclerosis: A longitudinal study of saccadic and cognitive tasks». en. In: *Amyotrophic Lateral Sclerosis and Frontotemporal Degeneration* 17.1-2 (Feb. 2016), pp. 101–111. ISSN: 2167-8421, 2167-9223. DOI: 10.3109/21678421.2015.1054292. URL: <http://www.tandfonline>.

- com/doi/full/10.3109/21678421.2015.1054292 (visited on 05/14/2026) (cit. on p. 6).
- [14] A. Ramos Murguialday, J. Hill, M. Bensch, S. Martens, S. Halder, F. Nijboer, B. Schoelkopf, N. Birbaumer, and A. Gharabaghi. «Transition from the locked in to the completely locked-in state: A physiological analysis». en. In: *Clinical Neurophysiology* 122.5 (May 2011), pp. 925–933. ISSN: 13882457. DOI: 10.1016/j.clinph.2010.08.019. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1388245710006619> (visited on 05/14/2026) (cit. on pp. 6, 7).
- [15] Federica Cozza et al. «Ocular Involvement Occurs Frequently at All Stages of Amyotrophic Lateral Sclerosis: Preliminary Experience in a Large Italian Cohort». en. In: *Journal of Clinical Neurology* 17.1 (2021), p. 96. ISSN: 1738-6586, 2005-5013. DOI: 10.3988/jcn.2021.17.1.96. URL: <https://thejcn.com/DOIx.php?id=10.3988/jcn.2021.17.1.96> (visited on 05/14/2026) (cit. on p. 6).
- [16] Bong-Hui Kang, Jae-Il Kim, Young-Min Lim, and Kwang-Kuk Kim. «Abnormal Oculomotor Functions in Amyotrophic Lateral Sclerosis». en. In: *Journal of Clinical Neurology* 14.4 (2018), p. 464. ISSN: 1738-6586, 2005-5013. DOI: 10.3988/jcn.2018.14.4.464. URL: <https://www.thejcn.com/DOIx.php?id=10.3988/jcn.2018.14.4.464> (visited on 05/14/2026) (cit. on p. 6).
- [17] Majid Khalili-Ardali, Shizhe Wu, Alessandro Tonin, Niels Birbaumer, and Ujwal Chaudhary. «Neurophysiological aspects of the completely locked-in syndrome in patients with advanced amyotrophic lateral sclerosis». en. In: *Clinical Neurophysiology* 132.5 (May 2021), pp. 1064–1076. ISSN: 13882457. DOI: 10.1016/j.clinph.2021.01.013. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1388245721000341> (visited on 05/14/2026) (cit. on p. 7).
- [18] Ujwal Chaudhary, Natalie Mrachacz-Kersting, and Niels Birbaumer. «Neuropsychological and neurophysiological aspects of brain-computer-interface (BCI) control in paralysis». en. In: *The Journal of Physiology* 599.9 (May 2021), pp. 2351–2359. ISSN: 0022-3751, 1469-7793. DOI: 10.1113/JP278775. URL: <https://physoc.onlinelibrary.wiley.com/doi/10.1113/JP278775> (visited on 05/14/2026) (cit. on p. 7).
- [19] A. Kübler and N. Birbaumer. «Brain-computer interfaces and communication in paralysis: Extinction of goal directed thinking in completely paralysed patients?» en. In: *Clinical Neurophysiology* 119.11 (Nov. 2008), pp. 2658–2666. ISSN: 13882457. DOI: 10.1016/j.clinph.2008.06.019. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1388245708009115> (visited on 05/14/2026) (cit. on p. 7).

- [20] Janice Light, David McNaughton, David Beukelman, Susan Koch Fager, Melanie Fried-Oken, Thomas Jakobs, and Erik Jakobs. «Challenges and opportunities in augmentative and alternative communication: Research and technology development to enhance communication and participation for individuals with complex communication needs». en. In: *Augmentative and Alternative Communication* 35.1 (Jan. 2019), pp. 1–12. ISSN: 0743-4618, 1477-3848. DOI: 10.1080/07434618.2018.1556732. URL: <https://www.tandfonline.com/doi/full/10.1080/07434618.2018.1556732> (visited on 06/17/2026) (cit. on p. 8).
- [21] Kaishan Wang, Penghu Wei, Yongzhi Shan, and Guoguang Zhao. «Clinical practice of Brain–Machine interfaces in neurological disorders». en. In: *EngMedicine* 2.3 (Sept. 2025), p. 100090. ISSN: 29504899. DOI: 10.1016/j.engmed.2025.100090. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2950489925000363> (visited on 06/17/2026) (cit. on pp. 8–10).
- [22] David Beukelman, Susan Fager, and Amy Nordness. «Communication Support for People with ALS». In: *Neurology Research International* 2011 (2011), p. 714693. ISSN: 2090-1852. DOI: 10.1155/2011/714693. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3096454/> (visited on 06/17/2026) (cit. on p. 8).
- [23] Hilary O. Edughele, Yinghui Zhang, Firdaus Muhammad-Sukki, Quoc-Tuan Vien, Haley Morris-Cafiero, and Michael Opoku Agyeman. «Eye-Tracking Assistive Technologies for Individuals With Amyotrophic Lateral Sclerosis». In: *IEEE Access* 10 (2022), pp. 41952–41972. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2022.3164075. URL: <https://ieeexplore.ieee.org/document/9745906/> (visited on 06/17/2026) (cit. on pp. 9, 10).
- [24] Hitesh Yadav and Surita Maini. «Decoding brain signals: A comprehensive review of EEG-Based BCI paradigms, signal processing and applications». en. In: *Computers in Biology and Medicine* 196 (Sept. 2025), p. 110937. ISSN: 00104825. DOI: 10.1016/j.combiomed.2025.110937. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010482525012892> (visited on 06/17/2026) (cit. on p. 9).
- [25] Yuchun Wang et al. «Advances in brain computer interface for amyotrophic lateral sclerosis communication». en. In: *Brain-X* 3.1 (Mar. 2025), e70023. ISSN: 2835-3153, 2835-3153. DOI: 10.1002/brx2.70023. URL: <https://onlinelibrary.wiley.com/doi/10.1002/brx2.70023> (visited on 06/17/2026) (cit. on pp. 9, 10).

- [26] Yuxin Qin, Baojiang Li, Wenlong Wang, Xingbin Shi, Cheng Peng, and Yifan Lu. «Classification Algorithm for fNIRS-based Brain Signals Using Convolutional Neural Network with Spatiotemporal Feature Extraction Mechanism». en. In: *Neuroscience* 542 (Mar. 2024), pp. 59–68. ISSN: 03064522. DOI: 10.1016/j.neuroscience.2024.02.011. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306452224000617> (visited on 06/17/2026) (cit. on p. 9).
- [27] Debadatta Dash, Alan Wisler, Paul Ferrari, Elizabeth Moody Davenport, Joseph Maldjian, and Jun Wang. «MEG Sensor Selection for Neural Speech Decoding». In: *IEEE Access* 8 (2020), pp. 182320–182337. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.3028831. URL: <https://ieeexplore.ieee.org/document/9214489/> (visited on 06/17/2026) (cit. on p. 9).
- [28] Kimberley Wyse-Sookoo et al. «Stability of ECoG high gamma signals during speech and implications for a speech BCI system in an individual with ALS: a year-long longitudinal study». In: *Journal of Neural Engineering* 21.4 (Aug. 2024), p. 046016. ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2552/ad5c02. URL: <https://iopscience.iop.org/article/10.1088/1741-2552/ad5c02> (visited on 06/17/2026) (cit. on p. 10).
- [29] Maria Carles. *Human Anatomy and Physiology Laboratory Manual: Understanding How Structure Determines Function*. en. ROTEL (Remixing Open Textbooks with an Equity Lens) Project, Dec. 2024. ISBN: 9781964276403. URL: <https://rotel.pressbooks.pub/anatomyphysiology/> (visited on 06/17/2026) (cit. on p. 11).
- [30] John E. Hall and Arthur C. Guyton. *Guyton and Hall textbook of medical physiology*. eng. 13th edition. Student consult. Philadelphia, Pa: Elsevier, 2016. ISBN: 9781455770052 9781455770168 (cit. on p. 11).
- [31] Gunter K. Von Noorden and Emilio C. Campos. *Binocular vision and ocular motility: theory and management of strabismus*. eng. 6th ed. St. Louis, Mo.: Mosby, 2002. ISBN: 9780323011297 (cit. on pp. 11, 12).
- [32] Mahsaw Motlagh and Ragi Geetha. «Physiology, Accommodation». eng. In: *StatPearls*. Treasure Island (FL): StatPearls Publishing, 2026. URL: <http://www.ncbi.nlm.nih.gov/books/NBK542189/> (visited on 06/17/2026) (cit. on p. 11).
- [33] Niranjana Sudhakar Deshmukh, A. S. Todkar, and Hemant Katakhar. «Measurement of Eviscerated Eye for Retinal Prosthesis—A New Approach». In: *Recent Advances in Mechanical Engineering*. Ed. by Anoop Kumar Shukla, Bhupendra Prakash Sharma, Ahmad Arabkoohsar, and Pradeep Kumar. Singapore: Springer Nature Singapore, 2023, pp. 541–549. ISBN: 978-981-99-1894-2 (cit. on p. 12).

- [34] O. Pratt and J. W. S. R. Anaesthetics. «The Autonomic Nervous System: Basic Anatomy and Physiology». In: *Continuing Education in Anaesthesia Critical Care & Pain* 8.3 (2008). Reprinted online via e-SAFE Anaesthesia Platform. Available at: <https://www.e-safe-anaesthesia.org/>, pp. 86–91. DOI: 10.1093/bjaceaccp/mkn012 (cit. on p. 12).
- [35] Pearson Scott Foresman. *Visual accommodation process illustration*. Wikimedia Commons. Public domain. Available online at: [https://commons.wikimedia.org/wiki/File:Accommodation_\(PSF\).png](https://commons.wikimedia.org/wiki/File:Accommodation_(PSF).png). 2007 (cit. on p. 12).
- [36] Rositsa Baltadzhieva, Tanya Gurevich, and Amos D Korczyn. «Autonomic impairment in amyotrophic lateral sclerosis». en. In: *Current Opinion in Neurology* 18.5 (Oct. 2005), pp. 487–493. ISSN: 1350-7540. DOI: 10.1097/01.wco.0000183114.76056.0e. URL: <https://journals.lww.com/00019052-200510000-00002> (visited on 06/17/2026) (cit. on p. 13).
- [37] Francesco Ponzio, Andres Eduardo Lorenzo Villalobos, Luca Mesin, Claudio de’Sperati, and Silvestro Roatta. «A human-computer interface based on the “voluntary” pupil accommodative response». en. In: *International Journal of Human-Computer Studies* 126 (June 2019), pp. 53–63. ISSN: 10715819. DOI: 10.1016/j.ijhcs.2019.02.002. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1071581918301101> (visited on 06/17/2026) (cit. on p. 13).
- [38] Andres Eduardo Lorenzo Villalobos, Silvia Giusiano, Luca Musso, Claudio de’Sperati, Alessandra Riberi, Piotr Spalek, Andrea Calvo, Cristina Moglia, and Silvestro Roatta. «When assistive eye tracking fails: Communicating with a brainstem-stroke patient through the pupillary accommodative response – A case study». en. In: *Biomedical Signal Processing and Control* 67 (May 2021), p. 102515. ISSN: 17468094. DOI: 10.1016/j.bspc.2021.102515. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1746809421001129> (visited on 06/17/2026) (cit. on pp. 13–15, 43).
- [39] Leonardo Cardinali, Silvestro Roatta, Raffaele Pertusio, Marcella Testa, and Cristina Moglia. «Gaming for Training Voluntary Control of Pupil Size». en. In: *Electronics* 11.22 (Nov. 2022), p. 3713. ISSN: 2079-9292. DOI: 10.3390/electronics11223713. URL: <https://www.mdpi.com/2079-9292/11/22/3713> (visited on 06/17/2026) (cit. on pp. 13, 15).
- [40] Giovanni Chiarion, Francesco Ponzio, Stefano Terna, Cristina Moglia, Edoardo Patti, Santa Di Cataldo, and Silvestro Roatta. «e-Pupil: IoT-Based Augmentative and Alternative Communication Device Exploiting the Pupillary Near-Reflex». In: *IEEE Access* 10 (2022), pp. 130078–130088. ISSN: 2169-3536.

- DOI: 10.1109/ACCESS.2022.3229143. URL: <https://ieeexplore.ieee.org/document/9984643/> (visited on 06/17/2026) (cit. on pp. 13, 27).
- [41] Livia Colucci, Leonardo Cardinali, and Silvestro Roatta. «Playing Jazz with the Pupil Accommodative Response: A Novel Unexplored Pupil-Based Interaction Mode». en. In: *Computer-Human Interaction Research and Applications*. Ed. by Hugo Plácido Da Silva and Pietro Cipresso. Vol. 2371. Cham: Springer Nature Switzerland, 2025, pp. 186–197. ISBN: 9783031838446 9783031838453. DOI: 10.1007/978-3-031-83845-3_11. URL: https://link.springer.com/10.1007/978-3-031-83845-3_11 (visited on 06/17/2026) (cit. on pp. 13–15).
- [42] F Schaeffel, H Wilhelm, and E Zrenner. «Inter-individual variability in the dynamics of natural accommodation in humans: relation to age and refractive errors.» en. In: *The Journal of Physiology* 461.1 (Feb. 1993), pp. 301–320. ISSN: 0022-3751, 1469-7793. DOI: 10.1113/jphysiol.1993.sp019515. URL: <https://physoc.onlinelibrary.wiley.com/doi/10.1113/jphysiol.1993.sp019515> (visited on 06/17/2026) (cit. on p. 14).
- [43] Sanjeev Kasthurirangan and Adrian Glasser. «Characteristics of pupil responses during far-to-near and near-to-far accommodation». eng. In: *Ophthalmic & Physiological Optics: The Journal of the British College of Ophthalmic Opticians (Optometrists)* 25.4 (July 2005), pp. 328–339. ISSN: 0275-5408. DOI: 10.1111/j.1475-1313.2005.00293.x (cit. on p. 14).
- [44] W. Neil Charman and Hema Radhakrishnan. «Accommodation, pupil diameter and myopia». eng. In: *Ophthalmic & Physiological Optics: The Journal of the British College of Ophthalmic Opticians (Optometrists)* 29.1 (Jan. 2009), pp. 72–79. ISSN: 1475-1313. DOI: 10.1111/j.1475-1313.2008.00611.x (cit. on p. 14).
- [45] Matteo Boemio. «Development of an assistive device based on multisymbolic pupillary communication for patients in a state of complete paralysis». Master’s thesis. Torino, Italia: Politecnico di Torino, 2025 (cit. on p. 15).
- [46] Pupil Labs. *Pupil Core Hardware and Software Technical Documentation*. Available online: <https://pupil-labs.com/products/core/tech-specs>. 2014 (cit. on p. 19).
- [47] Gazepoint. *Gazepoint GP3 Eye Tracker Hardware Specifications and Open Gaze API Documentation*. Available online: <https://www.gazept.com/product/gazepoint-gp3-eye-tracker/>. 2026 (cit. on p. 20).
- [48] Pupil Labs. *Pupil Capture open software download page*. Available online: <https://docs.pupil-labs.com/core/software/pupil-capture/>. 2014 (cit. on p. 23).

- [49] Giada Bernabè. *Pupil-com: Python Graphical User Interface for PAR-based Communication*. GitHub repository. Available online: <https://github.com/giadabernabe01/Pupil-com>. [Accessed: July 10, 2026]. 2026 (cit. on p. 25).
- [50] Giada Bernabè. *Pupil-com_da: MATLAB Post-processing Scripts for Data Analysis*. GitHub repository. Available online: https://github.com/giadabernabe01/Pupil-com_da. [Accessed: July 10, 2026]. 2026 (cit. on p. 45).