School of Design
Master Degree in Digital and Interaction Design
Academic Year: 2025-2026



POLITECNICO
MILANO 1863

# Design of a Multi-Sensory VR Meditation Space Experience for Emotional Relaxation in Young Adults

SUPERVISOR:
Prof. Monica Bordegoni

AUTHOR:
Jiayi Tang

STUDENT ID:11150740

# 0. Abstract

Mental health challenges among young adults have become an increasingly pressing public concern, with stress, emotional dysregulation, and anxiety representing the most prevalent and underserved dimensions of this crisis. While mindfulness-based meditation has demonstrated robust clinical efficacy in stress reduction, its adoption among young, non-meditating populations remains limited by barriers of accessibility, prior training requirements, and environmental dependency. This study presents Cardia, a multi-sensory immersive virtual reality meditation system designed to deliver accessible, evidence-based emotional relaxation experiences to adults aged 18 to 34.

Grounded in theories of psychophysiological self-regulation, attention restoration, and multi-sensory integration, Cardia was developed within Unreal Engine 5 and integrates four sensory channels: visually immersive natural scenes calibrated to empirically validated parameters. A breathing guidance protocol synchronised with particle system animations.Real-time physiological biofeedback derived from BITalino PPG sensor data processed into RMSSD heart rate variability metrics and mapped to dynamic visual outputs via OSC protocol. And an Arduino-controlled olfactory stimulation module delivering lavender scent. Together, these components form a closed-loop environment–physiology–subjective experience system intended to lower the cognitive threshold for entering a meditative state.

A between-subject controlled experiment was conducted with 27 young adult participants to evaluate the relaxation outcomes of two conditions: an audio-visual VR experience (Condition A) and the same experience augmented with lavender olfactory stimulation (Condition B). Outcomes were assessed through both physiological measures (ΔRMSSD) and four validated subjective scales (PSS-10, PANAS, SUDS, RSQ). Wilcoxon signed-rank tests revealed statistically significant within-group improvements across both conditions on measures of immediate stress perception, relaxation state, and positive affect, with large effect sizes ($r \geq 0.85$), supporting the efficacy of VR multi-sensory meditation as a short-form emotional regulation intervention. Between-group comparisons indicated favourable trends for the olfactory-augmented condition, though statistical significance was not reached at the current sample size, pointing to directions for future investigation.

This research demonstrates that the integration of immersive VR, physiological biofeedback, and multi-sensory design can produce meaningful relaxation outcomes within a

single five-minute session, offering a scalable, non-clinical pathway for supporting young adult mental health. The findings contribute to the emerging intersection of interaction design and psychological wellbeing, and lay a methodological foundation for future work in multi-sensory VR health intervention.

# Index

0. Abstract ..... 1
Index ..... 3
1. Introduction ..... 6
1.1. Mental Health Risks Among Contemporary Young Adults ..... 6
1.1.1. Empirical Evidence and Statistical Trends ..... 6
1.1.2. Associated Risks and Consequences ..... 7
1.2. Causes of Mental Health Issues ..... 10
1.2.1. Macro-level Causes ..... 10
1.2.2. Physiological Causes ..... 12
1.3. Existing Intervention Approaches and Their Limitations ..... 15
1.3.1. Immersive Technologies: VR and AR ..... 15
1.3.2. Intelligent Interventions: AI Assistants and Chatbots ..... 15
1.3.3. Socially Assistive Robots ..... 16
1.3.4. Towards an Integrated Approach ..... 17
1.4. Objectives ..... 17
1.4.1. Research Necessity and Rationale ..... 17
1.4.2. Feasibility and Design Consideration ..... 20
1.4.3. Research Methodology ..... 23
1.5. Thesis Structure ..... 25
2. Technical Background ..... 26
2.1. Virtual Reality: Concepts and Hardware ..... 26
2.1.1. Definition and Core Dimensions ..... 26
2.1.2. Hardware Architecture and Key Parameters ..... 26
2.2. VR Applications in Stress and Emotional Regulation ..... 27
2.2.1. Development Background and Application Overview ..... 27
2.2.2. Applications in Mental Health and Stress Regulation ..... 27
2.2.3. Current Limitations and Research Positioning ..... 28
2.3. VR Content Creation and Development Platforms ..... 28
2.3.1. Overview of Content Creation Technologies ..... 28
2.3.2. Unreal Engine Development Platform ..... 29

3. Design and Development ..... 30
3.1. Cardia System Development Process ..... 30
3.1.1. Context ..... 31
3.1.2. User Profile–Driven Service System Design ..... 32
3.2. Narrative Framework and Scenario Design ..... 42
3.2.1. Story Arc and Scene Sequencing ..... 44
3.2.2. Breathing Guidance Structure ..... 47
3.2.3. Heart-Rate-Driven Environment Transitions ..... 49
3.3. Virtual Environment Design ..... 51
3.3.1. Landscape Categories ..... 52
3.3.2. Colors and Visual Language ..... 54
3.3.3. Viewing Distance and Spatial Perception ..... 57
3.3.4. Natural Environment Typologies ..... 60
3.4. Project Development ..... 63
3.4.1. Arduino-Based Physiological Data Processing ..... 63
3.4.2. Niagara Particle Parameters for Biometric Mapping ..... 70
3.4.3. Rendering Performance Optimization ..... 74
4. Olfactory Experiment ..... 86
4.1. Olfactory System Design and Implementation ..... 86
4.1.1. Rationale for Olfactory Augmentation ..... 86
4.1.2. Physiological Data Pipeline: BITalino to UE5 ..... 88
4.1.3. HRV-Driven Particle Visualization ..... 93
4.1.4. Olfactory Device ..... 98
4.2. Experimental Design ..... 100
4.2.1. Conditions and Procedure ..... 100
4.2.2. Signal Acquisition and HRV Computation ..... 103
4.2.3. Subjective Assessment Instruments ..... 109
4.3. Experimental Results and Analysis ..... 111
4.3.1. Physiological Data Results ..... 111
4.3.2. Subjective Questionnaire Results ..... 114
4.3.3. Integrated Discussion ..... 117
5. Conclusion ..... 123

5.1. Research Objectives ..... 123
5.2. Completed Work ..... 123
5.3. Achievements ..... 124
5.4. Future Work ..... 125
Bibliography ..... 127

# 1. Introduction

## 1.1. Mental Health Risks Among Contemporary Young Adults

### *1.1.1. Empirical Evidence and Statistical Trends*

In recent years, the mental health of young people—particularly university students—has shown a continuing trend of deterioration, and has become a significant issue in global public health. Drawing on data from multiple cross-national and regional empirical studies, the scale and complexity of this crisis becomes increasingly clear across three interconnected dimensions: the overall burden of stress and anxiety, the specific clinical conditions that predominate, and the ways in which these problems have shifted over time and across demographic groups.

In terms of prevalence, the overall psychological stress burden among young people is currently at a historic high, with marked escalation observed both before and after the COVID-19 pandemic. A survey conducted across nine countries found that rates of high stress, depression, and anxiety symptoms among university students reached 61.3%, 40.3%, and 30% respectively.[1] . In Europe, while most students hold a positive view of their physical health, approximately 37% to 58% report low levels of psychological well-being[2]. Moreover, in a global cross-sectional study, around 59% of respondents met clinically significant anxiety thresholds, and 77% reported experiencing moderate or higher levels of stress[3]. Related research also notes that mental health problems among university students have risen year on year, accompanied by significant structural and systemic barriers—such as the cost of accessing resources and the stigma associated with mental illness—further limiting the likelihood of receiving effective support.

Beyond these general indicators of distress, it is equally important to consider which specific conditions are most prevalent and how they interact. Among young people who report psychological distress, depression and anxiety disorders constitute the most common clinical presentations[4]. Within diagnosed student samples in Europe, both conditions accounted for as much as 71% of cases. Additionally, attention deficit hyperactivity disorder (ADHD, approximately 25%), eating disorders (approximately 20%), personality disorders (approximately 11%), and addiction-related issues (approximately 8%) occur at notably high rates[5]. Of particular concern is the prevalence of comorbidity: approximately 24.5% of students face simultaneous risk from both depression and anxiety, substantially increasing the

complexity of any intervention or treatment pathway. When multiple conditions co-occur, the challenges of identification, referral, and sustained care multiply accordingly, and the likelihood of a young person falling through the gaps of existing support systems increases considerably[6].

Viewed across time and across different population groups, the picture becomes still more concerning. From a longitudinal perspective, mental health problems among young people have intensified significantly in recent years. In Poland, for example, the risk of depression among university students rose from 21.98% before the pandemic to 49.20% during it—a striking near-doubling that illustrates how profoundly external disruption can accelerate psychological deterioration within this population. While the pandemic represented an exceptional period, the trajectory it exposed was not new: the underlying pressures that drive mental health difficulties among young people had been building for years before the crisis arrived, and available evidence suggests they have not substantially receded since. From a generational standpoint, multiple studies consistently demonstrate that individuals aged 18 to 34 score significantly higher on measures of stress, depression, and anxiety than middle-aged or older cohorts[7], reflecting a heightened degree of psychological vulnerability that appears to be structurally tied to this particular stage of life rather than to any single external event. Cross-sectionally, the proportion of university students experiencing high stress exceeds 60%, far surpassing the roughly 11% recorded in the general population—a gap that is not easily explained by individual differences alone, but points instead to the distinctive pressure structure that characterises the transition from adolescence into adult social and professional life[8]. Taken together, these patterns of prevalence, condition-specific burden, and temporal intensification converge on a clear conclusion: the mental health challenges facing young people today are neither marginal nor temporary, but represent a sustained and structurally embedded crisis that demands sustained and structurally informed responses.

### *1.1.2. Associated Risks and Consequences*

The widespread prevalence of stress, depression, and anxiety among young people represents far more than personal emotional difficulty. Its effects tend to ripple outward, producing long-term and far-reaching negative consequences for individual developmental trajectories and for the functioning of society as a whole. These consequences operate across two broad levels: the damage done to individuals as they move through formative stages of life, and the cumulative harm that accumulates at the level of social institutions, economic structures, and intergenerational opportunity.

### *Impaired Academic Performance and Cognitive Function*

At the individual level, the most immediately visible effects are those that bear on academic performance and cognitive function. Mental health problems erode young people's cognitive functioning in ways that are often insidious—attention fragments, memory falters, and the motivation to learn quietly recedes. Individually, none of these symptoms is catastrophic, but in combination they can cause a student who was previously managing well to gradually slip. Research demonstrates that behavioural problems in early childhood can lead to persistent academic underperformance in adolescence, while psychological distress during puberty is significantly correlated with failure to complete compulsory education and the loss of access to higher education. Some young people with anxiety disorders go further, developing school refusal, which severs their educational trajectory entirely.[9]

### *Deteriorating Physical Health and Increased High-Risk Behaviour*

The physical consequences of sustained psychological distress are equally serious, and tend to compound the academic and social difficulties already described. Prolonged psychological distress carries an almost inevitable physical toll: chronic insomnia—with PSQI scores above 5 recorded in up to 73% of participants in some studies—compromised immune function, and impaired wound healing. More troubling still are the associated behavioural changes: young people experiencing depression and anxiety are more likely to turn to smoking or alcohol as short-term emotional outlets, and their risk of self-harm and suicide is substantially elevated compared with the general population. This dual physical and psychological depletion systematically erodes an individual's capacity for social adaptation, creating compounding vulnerabilities that are increasingly difficult to address in isolation.[10]

### *Declining Social Competence and Deepening Loneliness*

The social dimension of individual impact is no less significant. Social withdrawal is one of the most common outward expressions of depression and anxiety, but its consequences extend well beyond simply not wanting to socialise. Young people with mental health difficulties are at significantly higher risk of peer victimisation, and experiences of rejection and isolation in turn exacerbate emotional difficulties and social avoidance, forming a self-reinforcing cycle that is difficult to break without external intervention. What begins as withdrawal can harden over time into chronic social exclusion, with consequences that outlast the original period of distress by years or even decades.[11]

### *Long-Term Displacement of Life Trajectory*

Perhaps the most underestimated dimension of individual impact is the long-term displacement of life trajectory that untreated mental health difficulties can produce. Adolescent depression does not fade naturally with age—approximately two-thirds of those affected will experience a recurrence before the age of 21, with lasting effects on their career development, intimate relationships, and basic social functioning in adulthood. Rising unemployment risk, the absence of professional identity, early unplanned pregnancy: these issues may appear unrelated on the surface, but they converge on a single underlying cause—an unresolved period of psychological adversity that has quietly reshaped a person's entire possible life course.[12]

### *Labour Market Damage and Rising Economic Costs*

Beyond the individual, the consequences of widespread youth mental health difficulties accumulate into broader social and economic costs that can no longer be treated as peripheral concerns. At the level of the labour market, declining employment rates and reduced productivity among young people, once scaled across an entire generation, generate macroeconomic effects that cannot be ignored. Lower lifetime earnings mean a shrinking tax base and constrained consumer spending, while a qualitative decline in labour supply exerts a long-term drag on overall economic vitality. The mental health crisis among young people is no longer solely a healthcare matter—it is gradually becoming a systemic risk with implications for national economic structure.[13]

### *Heightened Risks to Social Stability and Public Safety*

Large-scale collective anxiety and social alienation are not merely individual emotional states. They are latent variables that erode public trust and social cohesion. The pandemic offered ample evidence: psychologically vulnerable populations exhibit greater susceptibility to misinformation and are more prone to amplifying and spreading panic. In extreme circumstances this can escalate into mass psychological crises. This risk remains largely hidden in ordinary times, but once external shocks arrive, it tends to surface at a speed and scale that exceed all expectations[14].

### *Aggravating Intergenerational Poverty and Social Inequality*

The relationship between economic pressure and mental health is bidirectional: young people with fewer resources consistently report lower psychological well-being, while mental illness further compresses their access to education and their prospects for upward mobility. These two forces compound and reinforce each other, potentially becoming entrenched at the family and intergenerational level as a structural predicament that is extremely difficult to

escape,because the cycle of poverty and psychological vulnerability does not end simply because its subjects have grown up[15].

## 1.2. Causes of Mental Health Issues

The mental health challenges facing young people can be examined systematically from two perspectives: macro-level social environment (encompassing societal conditions, employment, academic pressures, and interpersonal relationships) and individual physiological and psychological mechanisms (including brain development and the endocrine system). The youth phase, as a critical transitional period from adolescence into adult society, is characterised by heightened psychological sensitivity, young people are uniquely susceptible to the simultaneous influence of external structural changes and internal physiological fluctuations.

The youth period is broadly regarded as a transitional stage into adult society, during which individuals are required to complete a range of demanding tasks: establishing identity, affirming self-worth, planning a career, and navigating a shift in social roles. For all these reasons, the mental health of young people is acutely sensitive to changes in the social environment and to external sources of stress.

### *1.2.1. Macro-level Causes*

The mental health challenges facing young people today cannot be attributed to individual psychological fragility alone, they are deeply rooted in a constellation of structural and societal forces that have intensified considerably in recent decades. Four interconnected macro-level stressors are particularly salient: the transformation of the social environment under conditions of global uncertainty, the economic pressures associated with labour market instability, the competitive demands of the educational system, and the erosion of interpersonal belonging in an increasingly digitalised society.

To begin with, the shift towards an increasingly online way of life has exposed young people to marked social isolation. Reduced face-to-face contact has weakened established social support networks, intensifying feelings of loneliness, social detachment, and emotional suppression. Research has shown that loneliness frequently serves as a significant mediating factor through which stress is transformed into depression and anxiety[16].Compounding this, the information overload that characterises the digital media environment has further amplified young people's psychological burden. As intensive users of social media and digital

platforms, young people are continuously exposed to a negative information landscape dominated by economic recession, unemployment risk, and social competition. This chronic exposure tends to generate persistent emotional vigilance and anxiety, heightening the risk of psychological exhaustion and emotional dysregulation[17].

Closely intertwined with these social-environmental shifts are the employment pressures and economic uncertainties that disproportionately affect the 18–34 age group—precisely because this is the critical period during which professional identity and economic independence are established. On one side, labour market instability is fuelling intense career-related anxiety among young people. The global economic slowdown following the pandemic, shrinking job opportunities, and structural shifts across industries have left young people confronting fewer employment prospects, heightened competition, and elevated risks of unemployment. Research indicates that professional uncertainty affects not only financial income but also erodes individuals' sense of social identity and self-worth[18]. On the other side, there is a well-documented relationship between financial strain and psychological well-being. Persistently rising costs of housing, education, and daily living have left a large proportion of young people under sustained economic pressure. Some studies show that a notable share of young people report significant financial difficulty, and that prolonged economic stress not only diminishes subjective well-being but may—through a cycle of stress, reduced effectiveness, and heightened vulnerability—exert lasting effects on future career trajectories and lifetime earnings.

The educational system represents a third structural domain through which macro-level pressures translate into individual psychological distress. Operating essentially as a highly competitive social sorting mechanism, the educational system engages with mental health through a bidirectional and complex dynamic that scholars describe using the frameworks of Social Selection and Social Causation. From the Social Selection perspective, early mental health difficulties can directly shape subsequent academic trajectories: anxiety, attention deficits, and behavioural problems reduce learning efficiency and social adaptability, increasing the risk of academic underperformance, reduced educational opportunity, and even dropout. Research has shown that young people with anxiety disorders are particularly prone to school avoidance, difficulty concentrating, and declining academic achievement[19]. The Social Causation perspective, meanwhile, emphasises that academic failure is itself a significant source of psychological stress. Prolonged performance pressure, university entrance competition, experiences of failure, and social comparison can trigger or exacerbate anxiety, depression, self-negation, and emotional dysregulation. Academic pressure and

mental health thus tend to form a mutually reinforcing cycle[20], in which poor well-being undermines performance, and poor performance in turn deepens psychological distress.

Finally, the domain of interpersonal relationships constitutes a fourth macro-level dimension through which structural conditions shape young people's mental health. As young people move through this life stage, their primary relational focus shifts from family attachment towards peer relationships and social identity, meaning that social experience exerts an increasingly profound influence on psychological well-being. Yet the very conditions described above—digital mediation of social life, economic precarity, and competitive academic environments—create fertile ground for interpersonal difficulty and social exclusion. Research has found that young people in a state of anxiety are more inclined to interpret social interactions negatively, feeling overlooked, dismissed, or unwelcome even in ordinary cooperative or conversational settings. This pattern of negative social perception further reinforces tendencies towards social withdrawal and loneliness[21],completing a feedback loop in which macro-level stressors are internalised as interpersonal failures, thereby deepening the individual's sense of isolation and psychological vulnerability.

Taken together, these four macro-level domains—social-environmental change, economic uncertainty, academic competition, and interpersonal disconnection—interact and mutually reinforce one another to produce the heightened rates of stress, anxiety, and emotional dysregulation observed among young adults today. Understanding this structural backdrop is essential for contextualising the need for evidence-based, accessible, and scalable mental health interventions targeting this population.

### *1.2.2. Physiological Causes*

The heightened sensitivity of young people to external stress is not explained solely by social environmental change. It is equally rooted in physiological and neural mechanisms that have not yet reached full maturity. The period from 18 to 25 in fact remains a critical stage of ongoing brain and nervous system development, during which emotional regulation, stress recovery, and cognitive control are not yet fully stable. Young people are consequently more susceptible to anxiety, depression, and emotional imbalance under conditions of high pressure. Two physiological dimensions are especially significant in this regard: the delayed maturation of the prefrontal cortex, and the role of the endocrine system in shaping the individual stress response.

The brain does not stop growing when a person becomes a teenager. It continues developing all the way through the youth years and into early adulthood. One of the most

important regions still undergoing development during this period is the prefrontal cortex, which sits at the front of the brain and governs executive control, rational decision-making, impulse inhibition, and emotional regulation. In most people, this region does not fully finish developing until around the age of 25[22], meaning that its incomplete maturation directly affects young people's capacity to handle stressful events. Because the prefrontal cortex has not yet built up its full ability to regulate the emotional centres of the brain, young people tend to react more strongly to stressors encountered in daily life. When they face situations such as competing for jobs, performing poorly in their studies, navigating interpersonal conflict, or confronting uncertainty about the future, their emotional responses can rapidly escalate beyond what they intended. Negative feelings such as worry, sadness, and frustration can accumulate over time and become increasingly difficult to resolve.

From a neurological standpoint, the prefrontal cortex's capacity to exert top-down inhibitory control over the limbic system—particularly emotional centres such as the amygdala—has not yet reached adult levels. This means that when something unexpected or unpleasant occurs, young people are more likely to experience a strong emotional reaction quickly, finding it harder to bring those reactions under control. This heightened reactivity can give rise to what researchers call State Anxiety—a feeling of intense tension, worry, and hypervigilance that emerges suddenly when a person encounters something they feel unable to manage. Compared to older adults whose brains have fully developed, young people often lack the stable and reliable cognitive tools needed to cope with ongoing stress and anxiety in a sustainable way. As a result, they are more likely to reach a state of emotional exhaustion in which even minor difficulties can feel overwhelming. When external pressure persists over an extended period, this hypersensitive neurological state may evolve into chronic anxiety disorder or depressive symptoms[23].

Beyond neurodevelopmental factors, the endocrine fluctuations characteristic of the youth period profoundly affect individual psychological stability and stress tolerance. Cortisol is widely regarded as one of the body's most important stress hormones. When an individual experiences prolonged stress, loneliness, or social isolation, the Hypothalamic–Pituitary–Adrenal (HPA) axis is continuously activated, triggering an intensified cortisol secretion response. Research has shown that sustained high cortisol levels not only heighten anxiety and emotional vigilance, but may damage brain regions associated with memory, learning, and emotional regulation—including the hippocampus—further impairing cognitive function and increasing the risk of depression[24].

In addition, changes in sex hormone levels represent another important physiological contributor to mental health disparities among young people. Epidemiological evidence consistently demonstrates that anxiety and depression are significantly more prevalent among young women than young men—a pattern that persists across diverse cultural and demographic contexts and cannot be adequately explained by differences in life circumstances alone. A substantial part of this disparity is attributable to the regulatory role of female sex hormones, particularly oestrogen, whose cyclical fluctuations exert a direct modulatory effect on the brain's stress response systems. Specifically, oestrogen interacts with the HPA axis in ways that heighten neuroendocrine sensitivity to psychosocial stressors, rendering the stress response more reactive under conditions of social evaluation, interpersonal conflict, and emotionally taxing situations. This does not reflect a deficit in coping capacity, but rather a biologically distinct stress-response architecture that amplifies both physiological arousal and subjective emotional intensity in response to certain categories of stressor.[25]

Taken together, mental health difficulties in young people are best understood not as the product of any single causal factor, but as the outcome of multiple converging influences operating simultaneously across external and internal domains. At the macro-level, structural pressures—including future uncertainty, labour market competition, academic demands, and complex interpersonal dynamics—accumulate to produce a sustained and compounding psychological burden. At the physiological level, incomplete prefrontal cortical maturation, endocrine fluctuations, and disrupted sleep architecture collectively reduce stress tolerance and impair the capacity for emotional recovery. These internal vulnerabilities do not operate independently of external conditions, rather, they interact with structural stressors in ways that amplify their psychological impact beyond what either dimension would produce in isolation.

This biopsychosocial understanding carries direct implications for intervention design. Approaches that address only the cognitive or behavioural dimensions of distress—such as psychoeducation or conventional talking therapies—are insufficient when the underlying vulnerability is partly physiological. Effective support for young people must therefore operate across multiple levels simultaneously: reducing exposure to structural stressors, while also directly targeting the nervous system through modalities capable of promoting parasympathetic activation, facilitating emotional recovery, restoring physiological equilibrium, and improving sleep quality. Only through such an integrated, multi-level approach can interventions achieve meaningful and sustained improvements in young people's mental health.

## 1.3. Existing Intervention Approaches and Their Limitations

In addressing the mental health needs of young people—particularly university students—current technological and therapeutic approaches are rapidly evolving from traditional clinical models towards digital, immersive, and intelligent modalities. Yet despite the breadth of innovation in this field, each category of intervention carries significant limitations that constrain its effectiveness as a scalable, accessible, and holistic solution.

### *1.3.1. Immersive Technologies: VR and AR*

*Virtual Reality Exposure Therapy (VRET):* Currently one of the most established applications, VRET involves gradually presenting anxiety-inducing stimuli within safe, controllable virtual environments to help young people manage phobias, PTSD, and social anxiety[26].

*Nature-Based Restoration (Nature Therapy):* VR is widely deployed to simulate natural landscapes such as forests and beaches. Research suggests that such immersive experiences can effectively reduce stress levels among university students and alleviate negative affect[27].

*Embodied Medicine:* By exploiting VR's capacity for 'body transfer'—allowing users to inhabit a different virtual avatar—this approach can help patients shift their self-perception, treat eating disorders, or cultivate self-compassion.

Despite these promising applications, immersive technologies face persistent structural constraints. Equipment costs remain high, and the specialist infrastructure required limits deployment outside clinical settings. A significant proportion of users experience adverse physical effects—including motion sickness, disorientation, and visual fatigue—that reduce tolerance and treatment adherence[28]. Furthermore, the clinical evidence base for VR-based interventions remains methodologically uneven: study sample sizes are frequently small, follow-up periods are short, and standardised protocols for therapeutic delivery are largely absent, making it difficult to draw generalisable conclusions about long-term efficacy[29]. Critically, current VR interventions tend to address stress and anxiety at the cognitive and perceptual level, without directly engaging the physiological dimensions of the stress response—such as autonomic nervous system dysregulation or disrupted sleep architecture—that underlie sustained psychological distress.

### *1.3.2. Intelligent Interventions: AI Assistants and Chatbots*

Automated Psychological Coaching (Chatbots): Applications such as Woebot employ AI to deliver real-time conversations grounded in Cognitive Behavioural Therapy (CBT), helping young people identify cognitive distortions, monitor their moods, and practise positive psychology techniques such as gratitude and stress reframing.

Personalised Precision Interventions: AI can tailor intervention content to users' personality profiles (such as levels of neuroticism or conscientiousness on the Big Five scale), thereby improving treatment adherence and effectiveness.

While AI-based chatbots offer advantages of accessibility, scalability, and low stigma, their clinical effectiveness remains contested. Early trials—including the frequently cited Fitzpatrick et al. (2017) randomised controlled study of Woebot—reported significant short-term reductions in depression and anxiety among college students, but were limited by brief intervention windows of two weeks and comparisons against passive controls rather than active therapeutic alternatives. Subsequent studies produced mixed results: a 2022 trial found no significant difference between a CBT chatbot and a control application for depression outcomes[30]. A broader limitation concerns the modality itself: text-based conversational agents engage users at the level of cognition and verbal self-reflection, but cannot deliver somatic or sensory input capable of directly modulating physiological arousal. For young people whose distress is partly rooted in autonomic hyperreactivity and endocrine dysregulation—as discussed in Section 1.2.2—purely cognitive interventions may address only part of the problem.

### *1.3.3. Socially Assistive Robots*

*Emotional Bonding and Companionship:* Robots such as Jibo are not passive instruments. Through prosocial behaviours—greeting users proactively, addressing them by name—they build a 'therapeutic alliance' with students. This rapport has been shown to significantly enhance the effectiveness of positive psychology interventions.

However, the evidence base for socially assistive robots in mental health contexts remains limited and methodologically fragmented. A recent umbrella review and meta-analysis found that while social robots showed some promise in improving non-psychiatric outcomes such as loneliness and positive affect, they demonstrated no significant pooled effects on clinical psychiatric outcomes including depression and anxiety[31]. Research in this domain has furthermore concentrated almost exclusively on older adult populations, and findings are not readily transferable to the distinct psychological needs and relational expectations of young adults. Beyond efficacy questions, barriers of cost, accessibility, and

user acceptance—particularly concerns around emotional authenticity and the ethics of human-robot attachment—represent substantial obstacles to widespread adoption[32].

### *1.3.4. Towards an Integrated Approach*

Taken together, these three categories of intervention—immersive technologies, intelligent conversational agents, and socially assistive robots—each address meaningful dimensions of the mental health challenge facing young people, yet none achieves comprehensive coverage of its complexity. A fundamental gap running across all three is the absence of direct physiological engagement: existing approaches tend to intervene at the level of cognition, perception, or social interaction, but stop short of delivering multimodal, embodied experiences capable of simultaneously calming the nervous system, regulating autonomic arousal, and restoring the physiological conditions necessary for emotional recovery. It is precisely this gap that motivates the design approach developed in the present project.

## 1.4. Objectives

### *1.4.1. Research Necessity and Rationale*

Between April and May 2026, this study conducted a survey targeting young people, primarily current students and recent graduates. A total of 18 valid responses were collected from participants aged between 22 and 29, with an average age of 24.8. The majority of respondents were female (67%), and the sample included current students (39%), recent postgraduate graduates (28%), and recent undergraduate graduates (28%). The survey employed three well-validated instruments: the Perceived Stress Scale (PSS-10), the Positive and Negative Affect Schedule (PANAS), and a self-rated physical stress scale. Together, these tools provided a multi-dimensional assessment of the mental and emotional state of the participants. The results aligned closely with findings from existing studies conducted globally, offering strong empirical support for the relevance and necessity of the intervention proposed in this research. Three dimensions of the survey findings are particularly instructive: the universally elevated stress levels that establish a clear basis for intervention, the emotional profile characterised by prominent negative affect and suppressed positive affect, and the patterns of stress perception and coping that reveal significant gaps in young people's self-regulation resources.

Regarding stress levels, the PSS-10 results revealed that the overall burden of stress among participants was both clearly and significantly elevated. Of the 18 respondents, 83%—that is, 15 individuals—scored within the high stress range, defined as a total PSS score of 25 or above. The remaining 17% fell into the medium stress range, and not a single participant scored in the low stress range. These figures considerably exceed what is typically observed when the same scale is administered to general adult populations, indicating that the young people in this sample are carrying a heavy and sustained psychological burden. Examining specific stress dimensions further reinforces this picture. Question 3—which most directly captures daily emotional tension through the item 'In the past month, how often have you felt nervous or stressed?'—Yielded a mean score of 3.0 out of 5, with approximately 72% of respondents selecting a rating of 3 or above, indicating that they experienced such tension sometimes or more frequently. This pattern suggests that daily emotional tension has become a pervasive and normalised experience among young people in this sample, rather than an occasional occurrence, and clearly points to the need for routine, non-clinical forms of emotional support.

The emotional profile emerging from the PANAS scale provided a more detailed and nuanced account of participants' affective states. On the negative affect subscale—covering emotions such as distress, upset, guilt, fear, hostility, and irritability—the mean score was 2.98 out of 5. On the positive affect subscale—covering emotions such as interest, excitement, energy, enthusiasm, and pride—the mean score was only marginally higher at 3.08 out of 5. The narrowness of this gap is telling: it indicates that negative emotions were not merely present in a peripheral sense, but had become substantially embedded in participants' everyday emotional experience. Complementing this, respondents were also asked to rate their subjective level of emotional pain on a scale from 1 to 10. The mean score across all participants was 5.28, already above the midpoint of the scale. More strikingly, 44% of respondents—8 out of 18 individuals—assigned themselves a score of 7 or higher, placing them in the high distress range. This means that nearly half the sample was experiencing a level of emotional pain that was sufficiently intense to be clearly felt and difficult to ignore in the course of daily life. Taken together, these findings build a compelling and urgent case for the implementation of accessible emotional support mechanisms capable of reaching young people before their difficulties escalate into more serious clinical conditions. The table below summarises the key indicators from this survey alongside relevant reference values from the existing literature.

| Scale Dimension | Key Indicator | Study Result | Reference Mean (Literature) |
|---|---|---|---|
| PSS-10 Stress Total Score | High stress rate (≥25) | 83%（15/18） | ~60% among university students |
| PSS-10 Q3 | Daily stress ≥3 on scale | 72% of respondents | — |
| PANAS – Negative Affect | Mean score (1–5 scale) | 2.98 | ~2.3 in normal adults |
| PANAS – Positive Affect | Mean score (1–5 scale) | 3.08 | ~3.2 in normal adults |
| Subjective Distress Rating | High distress (≥7/10) | 44%（8/18 respondents） | — |

Table 1.1 Comparison of Key Stress and Affect Indicators Between Study Sample and Literature Reference Values (N = 18)

Turning to stress perception and coping behaviour, the survey findings revealed further and equally significant gaps. When asked how they understood and framed their own experience of stress, 61% of respondents—11 out of 18—described stress as predominantly a negative experience in their lives, while only 17% reported being able to view stress in a more constructive or growth-oriented way. For the large majority of participants, stress was not something they felt they could learn from or channel productively. It was simply something that felt harmful and brought them down. Looking at how participants actually responded when under pressure, the two most commonly reported reactions were emotional outburst—releasing feelings in a sudden and uncontrolled manner—and simply pushing through, silently suppressing emotions while forcing oneself to continue. Both responses were selected by 50% of respondents. In contrast, only 11% reported actively engaging in deliberate, healthy emotion regulation strategies when under stress. This is a striking disparity: the great majority of young people in this sample did not possess a reliable set of tools or strategies for managing their emotional states when difficulties arose. Passive coping, rather than active and health-promoting self-regulation, was clearly the dominant pattern in this group.

When asked to evaluate the effectiveness of their current coping methods, participants gave an average rating of just 3.17 out of 5. Approximately 28%—5 individuals—felt their existing approaches were doing little to help, rating their effectiveness at 2 or below. The most commonly used methods were talking to friends (61%), listening to music (50%), and watching short videos or other entertainment (44%). While these are intuitive and readily accessible responses, they function primarily by diverting attention away from the source of

distress rather than engaging with the emotional system in a direct or regulatory way. It is also noteworthy that only 17% of respondents had any prior experience with meditation or breathing exercises, and as many as 39% had never tried any dedicated stress-relief method at all. This indicates that a substantial proportion of young people are navigating their daily emotional difficulties with very limited practical resources.

Taken together, three conclusions emerge clearly from these findings. First, the stress and negative affect levels among the young people in this study are already markedly elevated, and align closely with patterns documented by researchers across comparable populations internationally, establishing a strong empirical case for intervention. Second, the coping strategies currently employed by young people are predominantly passive and lack any direct engagement with emotional regulation processes, pointing to a significant gap in accessible self-regulation resources. Third, more than half of respondents had little or no experience with any structured approach to emotional regulation, demonstrating that there is genuine and practical value in offering an intervention that is easy to access and does not require a clinical context. The multi-sensory immersive VR meditation intervention proposed in this study is designed precisely to address these three converging gaps.

### *1.4.2. Feasibility and Design Consideration*

Based on the established necessity of this research, this section systematically examines the feasibility of the proposed VR meditation intervention from four perspectives: user acceptance of VR technology, physiological mechanisms supporting the intervention logic, the maturity of the technical implementation pathway, and evidence from previous related studies. Together, these aspects provide a comprehensive foundation for evaluating the practicality and potential effectiveness of the proposed design.

*Target Users' Acceptance of VR Technology*

The questionnaire results indicate that young adults generally have a basic understanding of and experience with VR technology. Among the participants, 56% (10 respondents) reported that they had used VR devices occasionally, while 11% (2 respondents) stated that they used VR frequently. Only 33% (6 respondents) had never used VR before. These findings suggest that the target user group is not completely unfamiliar with VR technology, which means that the learning curve and user training requirements are relatively low. As a result, the proposed intervention can be introduced with limited user education efforts, improving its practical applicability.

Regarding the willingness to use VR for stress reduction, participants demonstrated a notably positive attitude. The average interest score reached 3.89 out of 5. Furthermore, 61% (11 respondents) reported a high level of interest in VR-based stress relief applications, giving ratings of 4 or above. Among them, 39% (7 respondents) selected the highest possible rating of 5. These results indicate that young adults show considerable openness toward VR-assisted therapeutic experiences and are willing to explore such interventions as a method for managing stress and improving emotional well-being.

Overall, the findings suggest that the intended user population possesses both sufficient familiarity with VR technology and a strong interest in using VR for mental health support. This combination provides direct evidence for the user-side feasibility of the proposed VR meditation intervention and supports the likelihood of successful adoption and engagement during future implementation.

| VR Attitude Dimension | Agree / Positive (%) | Neutral (%) | Disagree / Negative (%) |
|---|---|---|---|
| VR is more immersive than mobile apps | 56% (10 participants) | 39% (7 participants) | 6% (1 participant) |
| VR can help me relax | 39% (7 participants) | 50% (9 participants) | 11% (2 participants) |
| I would be willing to try VR for stress reduction if given the opportunity | 61%（11 participants） | 39%（7 participants） | 0% (0 participants) |
| Using VR for stress reduction feels strange to me | 17%（3 participants） | 28%（5 participants） | 56%（10 participants） |

Table 1.2 Participants' Attitudes Toward VR-Based Stress Reduction (N = 18)

As shown in Table 1.2, 61% of respondents indicated willingness to use VR for stress reduction if given the opportunity, with no participant expressing outright refusal—a finding that reflects a broadly receptive attitude towards this modality. Additionally, 56% of respondents agreed that VR provides a greater sense of immersion than conventional mobile applications, suggesting that the medium's capacity for perceptual presence is already recognised by the target user group and may facilitate the attentional disengagement from daily stressors that effective stress reduction requires.

### *Physiological Mechanisms of Intervention*

From a neuroscientific perspective, multimodal immersive VR meditation rests on a well-established physiological rationale. As outlined in Section 1.2.2, the principal biological mechanisms underlying stress vulnerability in young people include the incomplete maturation of the prefrontal cortex—which reduces top-down inhibitory control over emotional centres such as the amygdala—and the sustained activation of the HPA axis, which drives prolonged cortisol secretion and heightens autonomic arousal. These neurophysiological conditions collectively impair emotional regulation and increase susceptibility to anxiety and psychological exhaustion.

Converging evidence from intervention research indicates that guided breathing, nature-based sensory immersion, and biofeedback-assisted relaxation are each capable of activating the parasympathetic nervous system, thereby attenuating cortisol secretion and reducing sympathetic arousal. The particular value of VR as a delivery platform lies in its capacity to integrate these distinct therapeutic modalities simultaneously within a single coherent experience. By combining visual nature scenes, auditory guidance, and olfactory stimulation through a scent diffusion module, the proposed system engages multiple sensory channels concurrently, inducing a depth of physiological relaxation that exceeds what any unimodal intervention could reliably achieve.

### *User Experience Preferences and Design Alignment*

The survey data concerning user experience preferences provide direct empirical support for the core structural decisions of the proposed design. With respect to preferred experiential outcomes, both *Relaxation in a quiet environment* and *Emotional release* emerged as the most frequently selected options, each chosen by 67% of respondents (12 individuals). *Distraction from real-life concerns* was selected by 39% (7 individuals). This distribution aligns closely with the project's design orientation, which centres on nature-based immersive scenes integrated with breathing guidance and narrative depth as primary vehicles for stress relief.

Regarding preferred interaction modalities, human voice guidance was the most favoured option, selected by 39% of respondents, followed by ambient environmental sound at 28%, with AI-generated voice and olfactory interaction each selected by 11%. Notably, two respondents provided unsolicited written comments in the open-response section, expressing a preference for naturalistic visual content and personalised adaptive settings—specifically requesting scene variations responsive to their real-time emotional state. This qualitative feedback offers additional support for the project's biofeedback-driven design logic, in which

the virtual natural environment is dynamically modulated according to the user's live heart rate data.

*Maturity of the Technological Path*

From an implementation standpoint, all primary technical components selected for this project are well-established and have demonstrated reliability in prior research contexts. Unreal Engine 5 (UE5) has been widely adopted for the development of high-fidelity VR environments, and its real-time global illumination and physically based rendering capabilities are well-suited to producing the naturalistic visual scenes required by this intervention. The BITalino hardware platform has been extensively validated in academic research settings for real-time acquisition of photoplethysmography (PPG) and electrodermal activity (EDA) signals. When integrated with a Python-OSC communication pipeline, the system enables sub-second transmission of physiological data to the UE5 environment, supporting responsive biofeedback modulation. The complete technical architecture therefore relies on no experimental or immature components, and is fully realisable within the scope of an academic research prototype.

### *1.4.3. Research Methodology*

This study adopts a progressive, three-phase framework that integrates research, design, and empirical testing. It combines three complementary methodological pathways—a quantitative survey, a design prototype, and a laboratory experiment using physiological data—into a coherent and mutually reinforcing research system in which each phase informs and supports the next.

The first phase focused on needs assessment and persona building, employing a mixed-methods approach to establish the empirical foundation for the study. An early-stage questionnaire was administered using standardised psychological instruments—specifically the PSS-10 for perceived stress and the PANAS for mood assessment—alongside a set of structured questions developed for this study. These tools were used to gather data from 18 participants concerning their mental health status, existing stress-coping behaviours, and attitudes towards VR-based interventions. The results of this phase served a dual purpose: they provided the empirical justification for the study's rationale, and they guided the selection of the design direction. By analysing the distribution of stress types and participants' preferred interaction styles, it was possible to construct representative user personas that informed the subsequent design process.

The second phase moved from research into design, applying a user-experience-centred design research methodology to develop the Cardia multi-sensory VR meditation system through an iterative prototyping process. This design work was organised around four integrated components. The first concerned the narrative framework and scene design: a cyclical breathing-guidance structure was developed and embedded within a series of distinct natural virtual environments—including deep-sea, forest, and cloud settings—constructed to facilitate emotional release and physical relaxation. The second component involved a heart-rate-driven environment: using the BITalino biosensor to monitor the user's heart rate in real time, the system triggers smooth, responsive transitions in the scene's colour temperature, ambient soundscape, and lighting, establishing a dynamic biofeedback loop between the user's physiological state and the virtual environment. The third component introduced a multi-sensory dimension beyond the audiovisual: an Arduino-controlled scent diffusion module was integrated to release olfactory stimuli corresponding to the virtual scenes, with the aim of deepening the felt sense of relaxation and immersion. The fourth component addressed the user interface and interaction design: informed by survey findings indicating that participants preferred human voice guidance and an uncluttered visual background, a minimal interaction layout was developed to avoid cognitive distraction and keep the user's attention directed towards the meditative experience.

The third phase constituted the empirical validation of the prototype through a quasi-experimental design incorporating both an experimental group and a control group. The testing procedure followed three sequential steps. In the pre-test stage, participants' baseline physiological data—including heart rate and skin conductance—were recorded, and mood scores were collected using PSS-10 and PANAS to ensure comparability between groups at the outset. During the intervention stage, the experimental group underwent the full Cardia system experience, including the dynamically responsive virtual scenes alongside synchronised lavender scent stimulation delivered via the olfactory module, the control group entered the same VR dynamic environment but without any olfactory component activated. In the post-test stage, physiological data and mood scores were collected again immediately following the experience, and statistical comparisons were conducted between groups to assess the differential effect of olfactory stimulation on overall relaxation outcomes.

These three phases connect to form a complete and self-sustaining research loop, moving from problem identification through solution development to effect verification. The data gathered in the first phase shaped the design decisions of the second, the prototype developed in the second phase provided the experimental instrument for the third. And the findings of

the third phase fed back to validate the overall system. Together, they advance step by step towards the central goal of this study: to determine whether a multi-sensory immersive VR meditation system can provide a meaningful, accessible, and physiologically grounded approach to emotional regulation for young adults.

## 1.5. Thesis Structure

To help readers follow the logic of this thesis, this section briefly outlines the content of the chapters that follow.

Chapter 2, "Technical Background," introduces the core concepts and hardware architecture of virtual reality, reviews its applications in stress and emotional regulation — including Virtual Reality Exposure Therapy and biofeedback-based relaxation — and outlines the VR development platforms and content-creation technologies used in this research, centred on Unreal Engine 5. Together, these provide the technical and theoretical foundation for the system design presented in later chapters.

Chapter 3, "Design and Development," describes the development process of the base version of the Cardia system, including the narrative framework and scenario design, the virtual environment design, and the heart-rate-driven visual feedback mechanism. The content of this chapter forms the base application that is further developed in the following chapter.

Chapter 4, "Olfactory Experiment," builds on this base application by introducing a new sensory dimension: olfactory stimulation. This chapter explains how the system was upgraded, covering the rationale for olfactory augmentation, the optimisation of the real-time physiological data pipeline, the implementation of heart-rate-variability-driven particle visualisation, and the hardware design of the olfactory device itself. Building on this upgraded system, the chapter then presents the design and execution of a controlled experiment comparing an audio-visual condition (Condition A) with an audio-visual-olfactory condition (Condition B), assessed through both physiological measures and subjective emotional scales. The results of this experiment are analysed and discussed in detail.\

Chapter 5, "Conclusion," restates the research objectives, reviews the work that has been completed, summarises the study's achievements, and reflects on its limitations and directions for future work.

# 2. Technical Background

## 2.1. Virtual Reality: Concepts and Hardware

### *2.1.1. Definition and Core Dimensions*

Virtual Reality (VR) refers to a computer-generated three-dimensional environment with which users can interact in real time, engaging multiple sensory modalities — principally vision and hearing — in ways that produce a compelling sense of physical presence within a simulated space.[33] The term was popularised by Jaron Lanier in the late 1980s, and the conceptual foundations of the field have since been articulated through four defining principles: Imagination, Immersion, Interaction, and Intelligence.[34]

Contemporary scholarship identifies three core dimensions that together constitute a functional definition of VR: immersion, the degree to which the technology supplants the user's perceptual access to the physical environment[35], presence, the subjective psychological state of "being there" within the virtual environment[36], and interactivity, the system's capacity to receive and respond to user input in real time.[37] These three dimensions provide the conceptual basis for evaluating the design choices of the Cardia system discussed in Chapter 3.

### *2.1.2. Hardware Architecture and Key Parameters*

The hardware architecture of a VR system functions as a tightly integrated input-output feedback loop, coordinating sensory output channels with real-time user input capture to sustain perceptual coherence.[38] On the output side, the head-mounted display (HMD) is the central device: current headsets range from wired systems requiring an external computer, such as the HTC Vive, to standalone wireless headsets with onboard processing, such as the Meta Quest series. Low-persistence OLED panels and optimised lens designs help reduce motion blur and keep the device comfortable for extended use.[39] On the input side, hand controllers typically support Six Degrees of Freedom (6DoF) tracking, achieved through either externally positioned base stations (outside-in tracking) or onboard cameras with computer-vision algorithms (inside-out tracking).[40]

A critical performance parameter governing both comfort and safety is Motion-to-Photon (MtP) latency — the interval between a physical movement and its corresponding update on the display. This latency must remain below approximately 20 milliseconds to prevent the perceptual conflict between the vestibular and visual systems that causes cybersickness. This threshold is directly relevant to the present research: any perceptible lag in the Cardia system's heart-rate-driven visual feedback would risk undermining, rather than supporting, the intended relaxation response (see Section 3.4.3 for the corresponding performance optimisation work).

## 2.2. VR Applications in Stress and Emotional Regulation

### *2.2.1. Development Background and Application Overview*

Since its origins in 1960s engineering prototypes, VR has progressed through successive waves of laboratory research and commercial development, reaching mainstream consumer adoption following the 2012 Oculus Rift crowdfunding campaign.[41] The technology is now applied across a wide range of domains, including education, entertainment, and industrial design, a systematic review of these applications falls outside the scope of the present research, which instead focuses on the domain most directly relevant to Cardia — the use of VR for psychological health and stress regulation.

### *2.2.2. Applications in Mental Health and Stress Regulation*

Virtual Reality Exposure Therapy (VRET) is among the most clinically established applications of VR in healthcare. Rather than asking a patient to confront a feared situation directly, VRET introduces it gradually, within a virtual environment, at a carefully controlled pace, through repeated, safe exposure, the nervous system recalibrates and the intensity of the fear response is reduced over time. VRET has been applied to specific phobias, social anxiety, and post-traumatic stress disorder.

For stress regulation specifically, VR systems typically pair immersive natural settings — forests, coastlines, open skies — with real-time biofeedback tracking signals such as heart rate and breathing. As the user moves through the environment, the system reads their physiological state and uses it to guide them towards slower breathing and deeper physical relaxation, actively shifting the autonomic nervous system away from a heightened stress response and towards a calmer baseline. This environment–physiology feedback loop is the direct conceptual precedent for the Cardia system, which extends the same principle by

mapping real-time heart rate variability to the visual and, in Condition B, olfactory qualities of the virtual environment (see Section 3.2.3).

A related and more forward-looking concept, sometimes termed Embodied Medicine, proposes that VR could regulate physiological and psychological states directly, by simulating internal bodily signals such as heart rate variability rather than only responding to them.[41] This framing is conceptually close to Cardia's closed-loop design, in which the visual and olfactory environment is not merely a passive backdrop but an active, physiologically responsive intervention.

### *2.2.3. Current Limitations and Research Positioning*

Two limitations in the existing literature are directly relevant to the positioning of this research. First, VR development has historically concentrated resources on visual fidelity — resolution, refresh rate, display precision — while comparatively little attention has been paid to other sensory channels, haptic feedback and, in particular, olfactory stimulation remain largely confined to experimental or prototype-stage implementations, without a mature, widely adopted technical approach.[42] This gap between visual richness and multi-sensory depth is precisely the space that the olfactory-augmented condition of the Cardia system, presented in Chapter 4, is designed to address.

Second, existing studies on VR for stress and mental health regulation are frequently constrained by small sample sizes, often fewer than 50 participants, short session durations, and an absence of long-term follow-up or randomised controlled designs.[43] These constraints limit the robustness of existing conclusions and are revisited in the discussion of this study's own sample size and design limitations in Chapter 4.

## 2.3. VR Content Creation and Development Platforms

### *2.3.1. Overview of Content Creation Technologies*

VR content creation has progressed from simple, low-detail three-dimensional models to computationally complex scenes incorporating visual, auditory, and increasingly tactile elements, alongside a broader shift from purely visual presentation towards multi-sensory experience design.[44] Rendering-side optimisation techniques such as foveated rendering, which restricts full-resolution rendering to the user's direct gaze region, have become an important strategy for balancing visual quality against the performance demands of real-time VR.

### *2.3.2. Unreal Engine Development Platform*

Unity and Unreal Engine remain the two dominant VR development platforms. Unity is widely used for its cross-platform flexibility and accessible C#-based scripting.[45] This research uses Unreal Engine 5 (UE5) as its development engine, primarily for its high-fidelity rendering capabilities: the Lumen system enables real-time global illumination and dynamic lighting, and the Nanite virtualised geometry system allows highly detailed assets to be used without manual optimisation, together achieving near-cinematic visual quality at a manageable performance cost.[46] UE5's Blueprint visual scripting system further lowers the barrier to building complex interaction logic, allowing designers without a programming background to participate directly in system construction — a workflow reflected in the development of Cardia's heart-rate-driven particle logic, which combines Blueprint with the Niagara particle system (see Section 3.4.2).[47]

# 3. Design and Development

## 3.1. Cardia System Development Process

Cardia is an immersive virtual reality relaxation system built around physiological signal input, designed to address the stress and emotional regulation needs that are widespread among young people today. The project is grounded in the theoretical frameworks of Positive Technology and mindfulness meditation, and its central aim is to take what has traditionally been a subjective, guidance-dependent form of psychological regulation and turn it into something that can be measured, perceived, and experienced with real-time feedback. For its primary environment, the system draws on the forest — a natural setting with well-documented restorative psychological effects — as its main experiential context. Through a combination of visual design, spatial audio, and rhythmic feedback, Cardia builds a space that feels enveloping and immersive, and that works on the user's emotional state not through instruction or narration, but through the sensory qualities of the environment itself.

On the technical side, Cardia collects physiological data from the user in real time through a heart rate sensor, monitoring indicators such as heart rate and heart rate variability trends. This data is passed through a communication channel built between Arduino and a virtual engine — in this case UE5 — where it is translated from abstract physiological signals into concrete, visible changes in the environment around the user.

The logic of this translation works in both directions. When the system detects that a user is in a heightened stress state, the environment responds accordingly — light flicker speeds up, particle movement becomes faster and more agitated, and the overall rhythm of the space quickens, making the user's internal tension something they can actually see and feel around them. As the user follows a guided breathing pattern — such as the slow, structured rhythm of Box Breathing — and their physiological signals begin to shift, the environment shifts with them. The light steadies, the particles slow, and the space settles into a calmer register. This creates a closed feedback loop between the body and the virtual world: the environment reflects the user's physiological state, and the user's physiological state is gradually shaped by the environment.

Taken as a whole, Cardia is more than a VR application prototype. It is a design exploration into what becomes possible when physiological computing and immersive interaction are brought together in a single system. Traditional forms of psychological support — however well-intentioned — tend to struggle with two things: immediacy and

individualisation. They often cannot respond to a person's state in the moment, and they frequently offer the same experience to everyone regardless of what that person actually needs right now. Cardia is built around both of these gaps.

### *3.1.1. Context*

In the context of accelerating social change and intensifying life pressures, psychological stress and emotional distress among the Chinese population have grown significantly in recent years. According to the *Blue Book of Mental Health: Report on the Development of Mental Health in China (2023–2024)*, published by the Institute of Psychology at the Chinese Academy of Sciences, depression levels peak among young adults aged 18 to 24 and decline progressively with age, with 38% of undergraduate students showing moderate or above depressive symptoms as of 2024[48]. Despite this growing need, access to professional support remains severely constrained: the utilisation rate of psychological counselling services among university students stands at only 2.8%, and among working adults, only 12% of those experiencing anxiety or depression actively seek help[49]. At the same time, public engagement with mental health services is undergoing a structural shift — from reactive crisis management towards proactive emotional self-care. The domestic psychological counselling market expanded from approximately 120 billion yuan in 2020 to nearly 300 billion yuan by 2024, with an annual compound growth rate exceeding 20%, and is projected to surpass 800 billion yuan by 2030 [50]. This trajectory reflects not only rising demand, but a broader cultural normalisation of mental health awareness among younger generations.

Taken together the findings from Chapters One and Two, the driving forces behind this design project converge clearly along two distinct dimensions.

The social context research and user survey presented in Chapter One reveal that young people — particularly students and recent graduates between the ages of 18 and 29 — are caught in a structurally embedded cycle of psychological stress. Survey data show that 83% of respondents fall within the high perceived stress range, and 44% report high levels of subjective distress. Existing stress-relief methods remain largely passive and surface-level in nature, with an average effectiveness rating of only 3.17 out of 5. At the same time, respondents demonstrated a strong willingness to engage with VR-based wellness tools — 61% expressed high interest, and none indicated explicit resistance. In terms of experience preferences, responses clustered around relaxation-oriented immersive environments and emotion-release interactions, each selected by 67% of participants. Taken together, these findings point to a clear design gap: given the limited accessibility of traditional mental health

resources, there is a meaningful social need and user demand for a low-barrier, self-directed, everyday emotional regulation tool capable of effectively triggering physiological relaxation responses.

The technical background research in Chapter Two provides a concrete pathway for addressing this gap. The unique capabilities of VR technology — spanning multi-sensory immersion, biofeedback integration, and real-time environment responsiveness — align precisely with the emotional regulation needs of young users. In particular, three core design mechanisms stand out: heart-rate-driven dynamic scene transitions, multi-sensory coordination through an olfactory module, and an immersive breath-guidance structure. Each of these can be implemented robustly within already validated technical frameworks, confirming both the feasibility and the engineering readiness of the proposed system.

Grounded in this dual research foundation, the design project — Cardia, a multi-sensory emotional relaxation VR meditation system — is now positioned to move into the system design and prototype development phase. It enters this stage supported by three pillars: a clearly defined problem awareness, thorough validation of user needs, and a reliable technical implementation pathway.

Chapter Three will document and elaborate on the full development process of the Cardia system. This includes service system design based on user personas, an immersive narrative framework and scene construction, the logic of heart-rate-driven dynamic environment transitions, and the visual language and spatial perception design of the virtual environment. Chapter Four will further present the introduction of the olfactory module and the system upgrade process, followed by the design and outcome analysis of a controlled experiment based on BITalino biometric data and the Unreal Engine 5 platform.

If Chapter One reveals why this design is worth doing, and Chapter Two demonstrates that it can be done, then from Chapter Three onward, this research will show exactly how it was designed and built.

### *3.1.2. User Profile–Driven Service System Design*

*Research Positioning and Scene Construction Logic*

This chapter focuses on constructing the future application scenarios for Cardia. Through a systematic design and description of its usage flow and service system within a private VR meditation and relaxation space, the chapter aims to demonstrate the feasibility of this research at the level of future commercialization and real-world deployment — rather than limiting its scope to experimental validation alone.

In terms of research logic, this chapter follows a progressive path structured around three stages: user, scenario, and system. The process begins by establishing core user personas grounded in questionnaire survey data and literature review findings. Building on this foundation, specific service application scenarios are then constructed. Finally, a user journey map is used to visualize the complete service flow and experience curve, providing a clear picture of how the system would unfold across time and touchpoints. This visualization serves to verify the internal consistency and practical operability of the design proposal under realistic conditions, confirming that the system holds together not only as a concept but as a workable experience that real users could move through.

### *Social Background and Current State of Psychological Stress Among the Target User Group*

Based on a combined analysis of questionnaire survey results and relevant literature, contemporary urban young adults — primarily between the ages of 25 and 35 — are facing a steady and ongoing rise in psychological stress. This stress does not stem from a single source. Instead, it is the product of multiple overlapping factors, including heavy workloads, interpersonal tensions, the demands of a fast-paced lifestyle, health-related anxiety, and the constant overload of digital information. The long-term accumulation of stress does not only affect an individual's immediate emotional state and daily quality of life — it also reaches deeper, undermining the sustainability of their physical and mental health as well as their broader social functioning.

Looking at how this group currently manages stress, common coping behaviors include scrolling through social media, listening to music, using meditation apps, spending time alone, or engaging in light physical activity. While these methods are widely adopted, they share a common limitation: they tend to lack sufficient depth of immersion, are easily disrupted by external distractions, and rarely enable a state of genuine deep relaxation. This gap between what existing tools offer and what users actually need forms the core starting point for the Cardia project.

Through a systematic review and synthesis of the survey data, the core target users of this project show several defining characteristics that together form the user foundation for this service design.

**High acceptance of emerging technology.** This group demonstrates strong willingness and prior experience with devices and systems such as VR/AR headsets, smart wearables, health-focused applications, and biometric technologies. Their overall level of digital literacy is high, and they are generally able to adapt to new interaction methods with relative ease.

**Sensitivity toward personal space and emotional boundaries.** This group has a strong sense of privacy and personal independence. They strongly prefer to manage their emotions in quiet, private settings free from external interruption, and show a clear aversion to the noise and exposure of shared public spaces.

**A focus on inner experience and physical wellbeing.** Compared to external entertainment consumption, this group is more willing to pay for products and services that enhance self-awareness and support physical and mental health. They are experience-oriented in a deeply personal sense and tend to exhibit a high level of self-reflection and body awareness.

**A degree of disposable income.** This group has the financial foundation to pay for high-quality experiential consumption. Their purchasing decisions tend to be rational and considered — they are more likely to make repeat purchases once they have clearly perceived the value a product provides.

Based on these four defining characteristics, the following user personas were developed. Together, the survey data and literature references cited throughout this section define the boundaries of the project's core target user group, and provide clear directional guidance for the design strategies applied in the subsequent service scenario development.

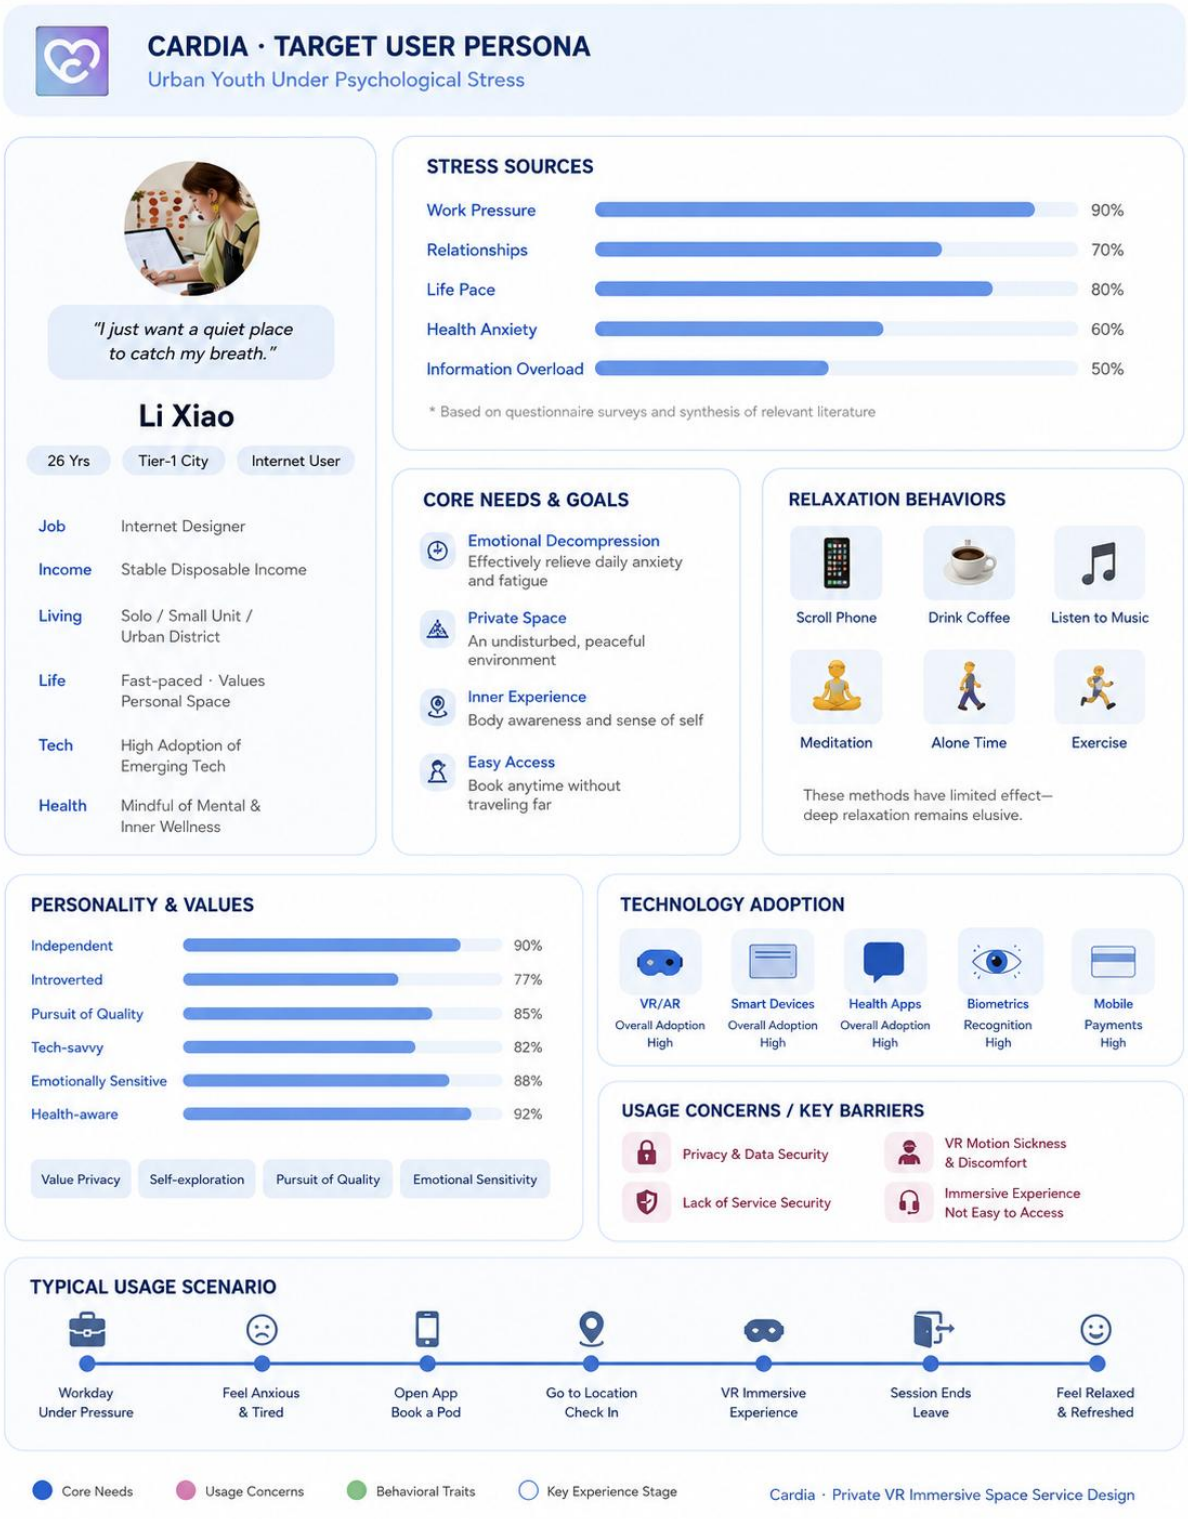


Figure 3.1 — Cardia Target User Persona

### *Core Needs and Usage Concerns*

At the level of core needs, the target users' expectations of the Cardia service center around four main dimensions. The first is effective emotional relief — the ability to meaningfully reduce everyday anxiety and fatigue within a relatively short period of time.

The second is the need for private space, meaning a physical environment that is independent, quiet, and free from interruption. The third is a pursuit of inner experience quality, where immersive means are used to achieve a state of deep relaxation in which the body and mind work in sync. The fourth is ease of use — the ability to book a session at any time through an online platform and access the service nearby, without the burden of a long commute.

At the same time, a number of significant usage concerns exist on the user side that cannot be overlooked when translating interest into actual adoption. These include worries about the security of personal and biometric data, a psychological resistance to real-name registration requirements, a sense of insecurity in fully unattended service environments, and concerns about the possibility of dizziness or physical discomfort caused by VR devices. If these barriers are not meaningfully addressed at the service design level, they will directly undermine users' willingness to try the service for the first time and reduce their likelihood of returning. For this reason, the service system design of this project establishes a specific response mechanism for each of these concerns, which will be described in detail in the following chapter.

With respect to the safety concerns that users may experience in an unattended service environment, this project further defines its service boundaries and risk management mechanisms. It is important to emphasize that the Cardia system is positioned as an emotional regulation and stress relief tool, not a medical diagnostic or therapeutic instrument. Its functional scope is focused on everyday, non-clinical emotional states such as general anxiety and fatigue — it does not constitute a substitute for professional psychological counseling or psychiatric care. Building on this positioning, the system introduces a simplified risk screening mechanism prior to booking, in which users complete a short questionnaire to provide an initial assessment of their current psychological state. When the system identifies signs that a user may be at elevated risk — such as severe anxiety, depressive tendencies, or other potential psychological concerns — it will not recommend solo immersive use. Instead, it will guide the user toward seeking professional medical or psychological support, reducing the potential risks associated with inappropriate use at the point of entry.

### *Future Application Scenario Design*

Building on the user persona described above, this project defines its future application scenario as a bookable, private VR immersive experience space system. In terms of spatial form, the system is embedded into urban neighborhoods through a distributed layout — similar to street-level convenience facilities or unmanned micro-service units. The goal is to

reduce the cost of reaching the service to a minimum, turning "always accessible" convenience into one of the product's core competitive advantages.

This research further analyzed device usage behavior and cost structure. Looking at real-world usage habits, VR devices in ordinary households tend to have low long-term usage frequency, showing a clear pattern of "novelty fade" , where usage drops sharply after the initial experience and the device eventually sits idle. This pattern is not unique to VR. It is also observed in other high-cost entertainment hardware such as home gaming consoles like the PlayStation 5, which similarly experience significant fluctuations in usage intensity over time. For a service scenario where emotional regulation is a low-frequency but high-value need, it makes little economic sense to ask users to bear the full cost of hardware purchase on their own.

By contrast, integrating VR devices with a spatial experience and deploying them in a shared service format significantly increases equipment utilization, lowers the cost per session, and reduces the upfront commitment required from users. Under this model, users do not need to absorb depreciation or maintenance costs — they simply pay per visit to access the complete experience. This "low barrier to entry, pay as you go" service structure aligns much more naturally with the consumption decision logic of the target user group.

For this reason, the distributed public experience space model adopted by this project is not, in any conventional sense, a "public venue." It is more precisely what might be called a "de-socialized private shared space" — a private-in-public model that preserves the individual's sense of privacy and personal experience while achieving a balance between resource efficiency and commercial sustainability.

Each experience unit is an independent, fully private space that supports end-to-end self-service operation. The interior is equipped with VR devices and basic comfort furnishings, including window views, soft seating, and carpeting. Through the coordinated combination of lighting atmosphere, material texture, and spatial narrative, the unit creates an environment that feels safe, comfortable, and genuinely conducive to physical and mental relaxation. The core logic behind this spatial strategy is to use both physical enclosure and environmental design as dual means of creating what is, for the user, a truly personal "emotional buffer zone" — a space that belongs entirely to them for the duration of their visit.



Figure 3.2 VR experience unit interior concept

Building on the foundational spatial design, the system further provides multi-dimensional personalization capabilities to enhance the targeting of emotional regulation and deepen the level of immersive experience. Specifically, users can make customized selections across two dimensions based on their current emotional state and personal preferences.

**Scent environment customization.** Users can choose from a range of olfactory atmosphere options including woody, fruity, and floral scents. These work through the neural association mechanisms between smell and emotion, helping users enter a relaxed state more naturally and effectively.

**Immersive theme scene selection.** A variety of VR meditation environments are available, including starry sky, ocean, and abstract artistic settings. Each is paired with differentiated lighting design and spatial narrative structures that guide users into varying depths of meditative experience.

These personalization mechanisms are not simply a layering of added features. They are grounded in a multi-channel intervention logic for emotional regulation — using the coordinated interaction of visual, olfactory, and tactile inputs (including soft furnishing textures) to maximize both the depth of immersion and the precision of emotional intervention.

In terms of spatial design, the system adopts a simplified and continuous structure that deliberately avoids complex visual information, which could otherwise create cognitive load and interfere with the relaxation process. Fine-grained control over lighting, spatial depth, and

environmental dynamics is used to build a stable and soothing atmosphere throughout the session. At the same time, the design places strong emphasis on rhythmic and gradual change, minimizing sudden stimuli in order to sustain the user's continuous attention and support the ongoing process of emotional regulation. This approach allows the virtual space to work in coordination with the user's internal state, laying the foundation for the physiological feedback interactions that follow in subsequent stages of the experience.

*Booking Mechanism and Identity Verification System*

On the service system and booking side, users are required to complete real-name verification and sign a privacy agreement, linking their basic information to the backend system. Once registered, users can check the availability of experience spaces through the online platform, and the system will automatically recommend the nearest available venue that matches their preferred time slot based on location data. After confirming a booking, users make an online deposit payment to receive a time-limited and unique access code.

The design logic behind introducing a real-name verification mechanism operates on three levels.

First, a dual verification approach combining a unique access code with biometric recognition — such as facial identification — ensures that access to the space is strictly limited to the verified individual, reducing the risk of unauthorized entry through code sharing or leakage.

Second, the system places limits on how frequently a single user can make bookings within a given period, preventing resource hoarding and the resale of reserved slots to third parties.

Third, in the absence of on-site staff management, real-name binding establishes a clear accountability mechanism for device usage, ensuring the safety and traceability of high-value VR equipment. It is important to emphasize that the backend system throughout this process follows strict data privacy protection principles. The personal information collected is used solely for the purposes of necessary identity verification and operational management, and is not applied to any other use.

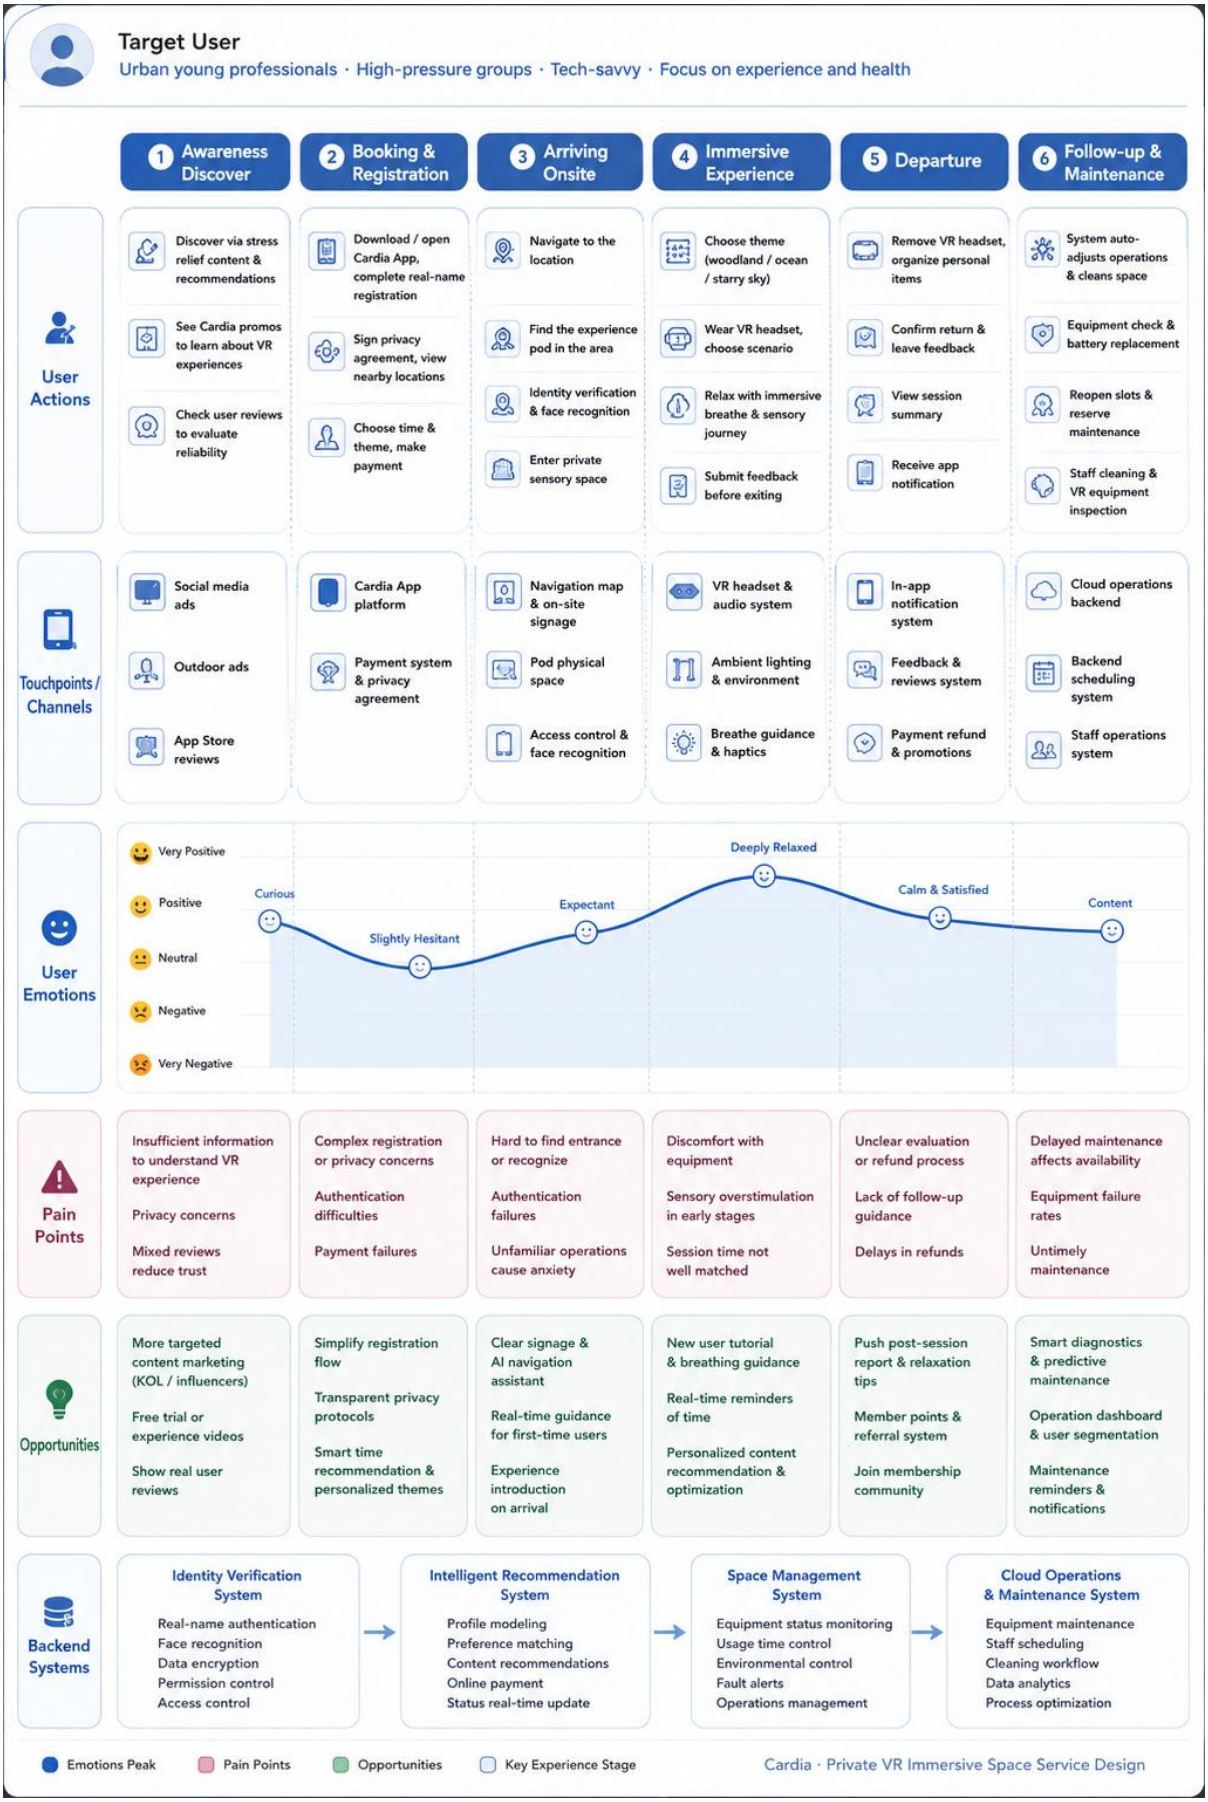


Figure 3.3 Cardia Service User Journey Map

*Unmanned Operations and Cloud-Based Scheduling and Maintenance*

In terms of space maintenance and operational management, the system adopts a coordinated model of "cloud scheduling plus zone-based operations," enabling efficient management of multiple distributed experience units without the need for traditional on-site staff.

Specifically, after each user session ends, the system automatically reserves a maintenance time window and pushes task instructions to the operations staff responsible for the corresponding area through a backend intelligent scheduling system. Operations staff are organized by community zone, with each person responsible for the routine maintenance of multiple experience units within a defined area. The system centrally plans the most efficient work route and time schedule for each staff member based on the usage status of individual units, improving operational efficiency and reducing overall running costs.

At the level of practical execution, maintenance tasks cover three areas. The first is a physical space integrity check, which includes verifying that furniture and fixtures are undamaged and complete. The second is a VR device operational status check, covering controller battery replacement and cable connection verification. The third is environmental cleaning and reset of the space, ensuring that every user receives the same quality of starting experience. Once all maintenance tasks have been completed and verified by the system, the experience unit is automatically reopened on the booking platform.

If damage or an abnormal condition is discovered during the maintenance process, the system will initiate an accountability tracing procedure against the corresponding user based on their real-name registration data, and apply the relevant compensation or liability mechanism within the applicable legal framework. This mechanism is designed to protect the safety and integrity of shared assets through institutional accountability, while also safeguarding the normal usage rights of other users.

In addition, to address the risk of unexpected incidents in an unstaffed environment, the system incorporates a set of basic safety intervention mechanisms. For example, each space is equipped with a one-tap exit and emergency assistance interface, allowing users to immediately end the session and leave the space if they feel any discomfort. At the same time, the system performs basic anomaly detection by monitoring session duration and fundamental interaction behavior, and can trigger a remote notification mechanism when necessary. Through these measures, the system provides users with a minimal but meaningful layer of safety protection without compromising the overall sense of private, uninterrupted experience.

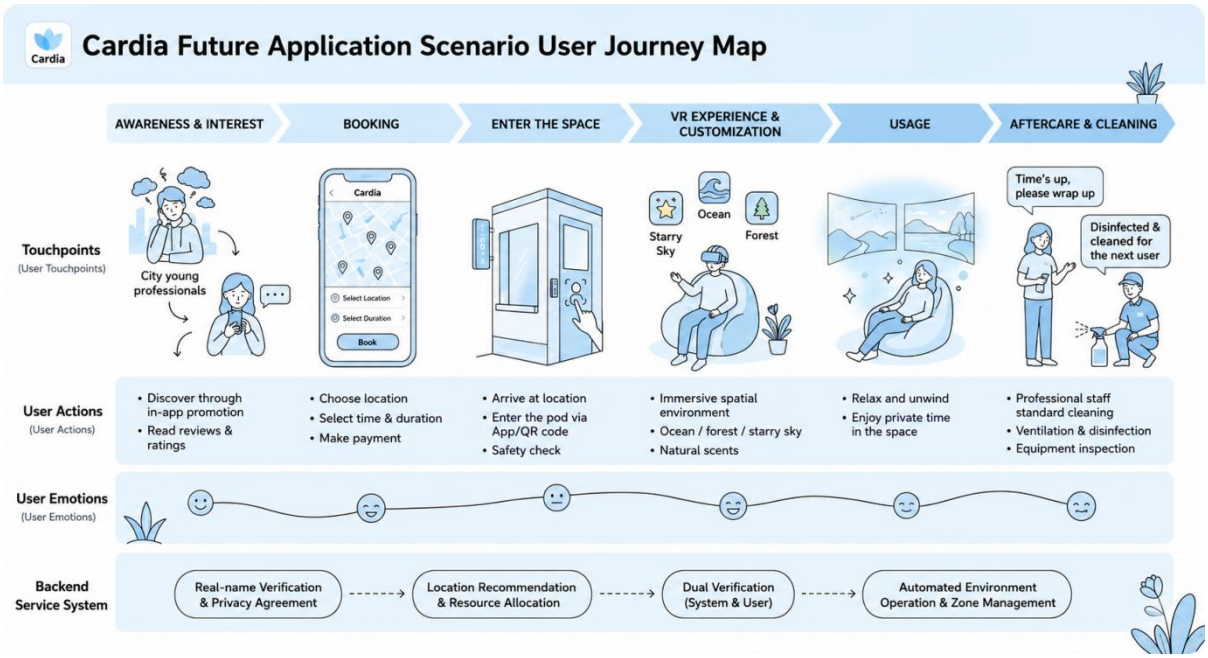


Figure 3.4 Cardia Future Application Scenario User Journey Map

## 3.2. Narrative Framework and Scenario Design

This system builds a closed-loop interaction logic driven by the user's physiological signals. Through real-time collection, processing, and feedback of heart rate data, the system translates the user's internal bodily state into dynamic changes within the virtual environment, gradually guiding them toward a state of calm and relaxation.

The process begins with the user putting on a heart rate sensing device, such as a wristband, to enable physiological data acquisition. The sensor continuously monitors changes in heartbeat and transmits the data to a microcontroller — an Arduino Uno — for basic processing and signal conversion. The converted signal is then fed into the virtual reality system running on UE5. Throughout this process, the system accomplishes a fundamental translation: it takes what is originally an imperceptible internal physiological state and converts it into a computable digital signal, which then serves as the core driving variable for all subsequent interaction. This translation is what allows the virtual environment to respond meaningfully to the user's body in real time, rather than playing out a fixed, pre-scripted sequence of events. The result is a feedback loop in which the user's own nervous system, in a sense, becomes the controller — and the act of calming down itself shapes the experience.

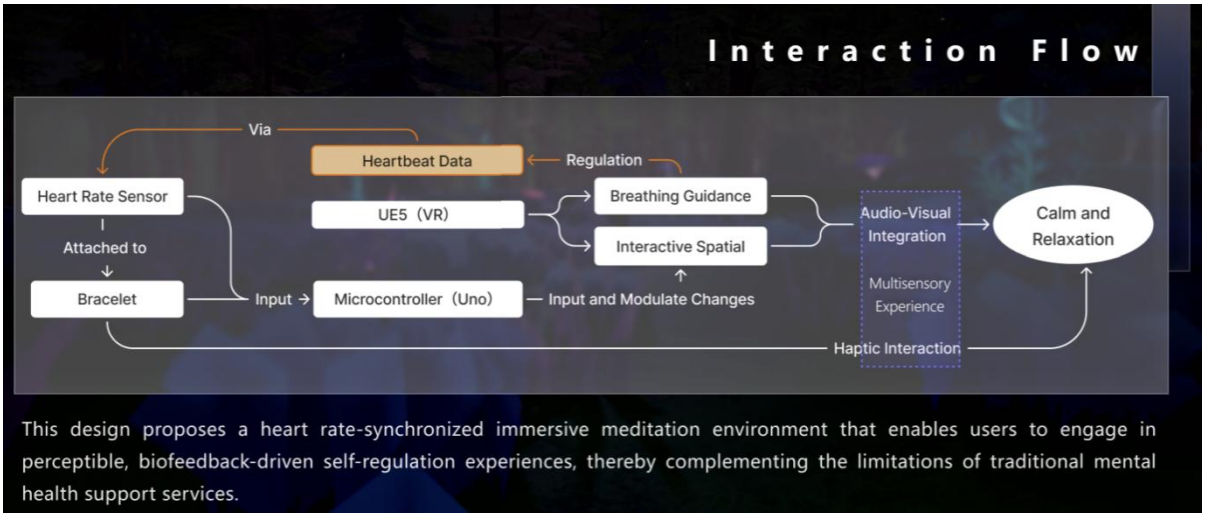


Figure 3.5 Interaction Flow

Within the virtual environment, heart rate data is embedded into the system logic as a key parameter and applied to spatial feedback in two ways. On one hand, the system guides users through regular breathing using visual cues or rhythmic prompts, lowering physiological arousal at a bodily level. On the other hand, changes in heart rate are mapped onto a range of environmental variables within the virtual scene — such as particle movement speed, spatial rhythm, and the intensity of visual feedback — allowing users to directly and intuitively "perceive the changes in their own state." This mechanism of externalizing physiological feedback into a visible, felt experience enables users to develop awareness of their internal condition during immersion, and to actively participate in self-regulation under the system's guidance. The result is a real-time interactive loop structured around physiological data, environmental feedback, and behavioral adjustment.

At the same time, the system further strengthens immersive depth and regulatory effectiveness through the integration of multimodal sensory inputs. At the visual level, the VR environment serves as the primary carrier, using natural or abstract spatial settings to build a relaxing atmosphere. Breath guidance provides structured rhythmic control, functioning as the main pathway through which users regulate their behavior. Where conditions allow, haptic feedback can also be introduced to reinforce physical presence and body engagement. The coordination of these different sensory channels across time and rhythm gives users a continuous and consistent compound sensory experience, improving both the effectiveness and the stability of emotional regulation. Over the course of sustained interaction, users gradually transition from a high-arousal state to a low-arousal state — expressed as a decrease in heart rate, a steadying of the breath, and a progressive release of emotional tension.

Taken as a whole, this system builds a closed-loop interaction mechanism centered on heart rate sensing, environmental feedback, and breath regulation, enabling users to carry out a physiological feedback-driven process of self-regulation within an immersive virtual space. This data-driven interaction approach not only strengthens the user's capacity to perceive and understand their own internal state, but also addresses the limitations of traditional psychological interventions in terms of real-time responsiveness and experiential quality. In doing so, it opens a new design pathway for building more personalized and embodied systems of psychological support.

### *3.2.1. Story Arc and Scene Sequencing*

This section defines the experiential narrative structure and usage scenarios of the system, with a focus on explaining how immersive environments with genuine emotional appeal are created through narrative construction and contextual design, so as to guide users into a process of relaxation and meditation.

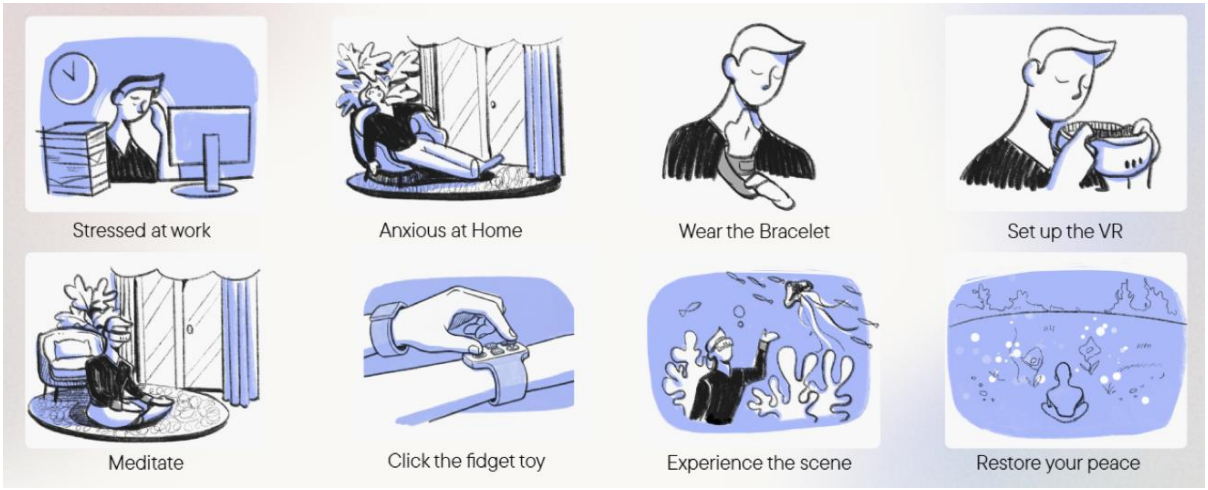


Figure 3.6 — Storyboard of the Cardia User Experience Narrative: From Stress Onset to Physiological Restoration

This storyboard presents the complete interactive narrative path of the project, tracing the journey from the user's initial problem all the way through to the experiential outcome. Its purpose is to translate the abstract concept of psychological stress into a concrete experience flow that can be addressed through deliberate design intervention.

**Frame 1 — High work pressure.** The narrative opens with a scene that will be immediately recognizable to the target user group: a person sitting at a desk, visibly weighed down by the demands of work. The clock on the wall, the posture of exhaustion, and the surrounding environment together communicate a state of sustained cognitive and emotional

overload. This is not a dramatic crisis but a quiet, chronic condition — the kind of low-grade stress that accumulates across a working day and does not simply disappear when the task is finished. This opening frame serves a deliberate design function: it anchors the entire narrative in a lived reality that the user already knows, establishing an immediate sense of recognition and relevance.

**Frame 2 — Anxiety during leisure time.** Rather than transitioning into rest, the user moves from the desk into a leisure setting and yet remains caught in a state of persistent inner tension. Lying down, physically still, but mentally unable to switch off — this frame captures one of the most common and underaddressed aspects of modern stress: the inability to truly rest even when the external demands have temporarily paused. This stage clearly defines the core pain point of the target user, which is the absence of an effective and sustainable means of emotional regulation. It also establishes the realistic context and demand basis from which all subsequent design intervention flows.

**Frame 3 — Putting on the wristband.**The first point of system contact is introduced calmly and without friction. The user puts on a wristband sensor — a familiar, unobtrusive form factor — and in doing so begins the transition from passive suffering to active engagement with a regulatory process. The wristband collects physiological signals, primarily heart rate data, which will serve as the foundational input variable for the feedback mechanism that follows. This moment is not merely a technical step. It is a behavioral turning point. The act of putting something on — of choosing to begin — marks a shift from "enduring stress" to "entering a state of self-regulation." The design of this interaction deliberately keeps the barrier as low as possible, requiring no configuration, no learning curve, and no prior technical knowledge.

**Frame 4 — Putting on the headset.** The second hardware interaction completes the transition from the physical world into the digital immersive space. As the VR headset is placed over the eyes, the user is separated from the visual environment of reality and brought into a virtual space that is fully controlled, carefully designed, and attuned to their current physiological state. This frame represents the crossing of a threshold — the moment at which the design system takes over as the primary environment. The headset is shown being worn with a sense of calm deliberateness, reinforcing the idea that this is a chosen, intentional act rather than an imposition.

**Frame 5 — Finding a comfortable space.** Once inside the virtual environment, the first stage of the experience focuses not on stimulation but on settling. The user finds a comfortable position within the virtual space — a moment of psychological and physical

placement that prepares the body and mind for what follows. This phase embodies a core design principle of the project: low cognitive load and low interaction demand at the point of entry. By removing the need for complex navigation, menu interaction, or decision-making, the system allows the user to arrive in the experience rather than perform it. The environment itself does the work of receiving the user, communicating safety and stillness through spatial design, soft materials, and gentle ambient elements.

**Frame 6 — Tactile relaxation.** A close-up interaction captures the user engaging with a tactile element — a button, a surface, a subtle physical control. This frame reinforces the multi-sensory dimension of the experience and highlights the role of haptic feedback in grounding the user in their body. Touch, in this context, is not decorative. It serves a physiological function: gentle, repeated tactile input has been shown to support parasympathetic nervous system activation, contributing to the reduction of arousal. In the design system, haptic cues are synchronized with breathing guidance and visual rhythm, creating a coherent multi-channel signal that speaks directly to the body's regulatory systems.

**Frame 7 — Experiencing the scene.** The user is now fully immersed in the core meditation environment. In this representative example, the chosen scene is a forest — an environment with strong psychological grounding in attention restoration theory and stress reduction research. The visual language of the scene uses particle systems, flowing organic forms, and a color palette tuned to low-arousal calming states. Heart rate data continues to drive the environment in real time: as the user's physiological state shifts toward calm, the scene responds — particles slow, light softens, and spatial rhythm deepens. The user is not watching the environment. They are in a feedback relationship with it. Their own body, through the data it produces, is shaping the world around them. It is worth noting that the forest scene is one example within a broader, expandable system of immersive environments. The full design framework includes multiple scene types with distinct psychological profiles — deep ocean, starfield, and abstract perceptual spaces among them — each corresponding to a different emotional regulation pathway. Users may choose their environment based on personal preference or current emotional state, enabling a more targeted and individualized experience.

**Frame 8 — Returning to calm.** The final frame shows the user in a state of visible stillness — heart rate lowered, breathing steady, emotional tension released. The figure is seated quietly, surrounded by a soft, diffuse environment that reflects the internal transition that has taken place. This is the goal state the entire system is designed to support: not euphoria, not excitement, but a return to baseline — a gentle restoration of equilibrium. The

experience ends not with a sudden cut back to reality but with a gradual transition, allowing the regulated state to consolidate before the user re-engages with the world outside.

Taken together, the eight frames of this storyboard trace a coherent arc: **situation introduction → device engagement → immersive experience → state transformation**. Each frame corresponds to a specific design decision — about interaction friction, sensory input, environmental response, and narrative pacing. The storyboard does not merely illustrate what the system does. It demonstrates why each element exists and what role it plays in the larger logic of physiological feedback-driven emotional regulation. By grounding the design in a recognizable human story, the system establishes emotional credibility before asking the user to trust it with something as personal as their own nervous system.

### *3.2.2. Breathing Guidance Structure*

To achieve the physiological regulation goals of the system, a cyclical breath guidance structure based on Box Breathing is introduced and embedded within the immersive environment, transforming breathing rhythm into a continuous experience of space and time. Rather than presenting breath guidance in the form of explicit verbal instructions, the system communicates the breathing pattern through subtle environmental cues — such as expansion, contraction, and rhythmic shifts within the surrounding space. This approach keeps the user's cognitive load low while allowing the body to respond naturally to the environmental rhythm, making the act of regulated breathing feel less like following a command and more like being gently carried by the world around them.

The specific breathing pattern adopted is the Box Breathing protocol, which divides each breath cycle into four equal phases: inhale for 4 seconds, hold for 4 seconds, exhale for 4 seconds, and hold again for 4 seconds. This symmetrical 4-4-4-4 structure has been selected for two reasons. First, its regularity makes it cognitively accessible — users do not need to memorise complex instructions or maintain conscious attention to the counting process once the rhythm has been internalised. Second, the equal duration of each phase produces a balanced activation of both the sympathetic and parasympathetic nervous systems, supporting a gradual shift towards a calmer physiological baseline without inducing the abrupt state changes that asymmetric breathing patterns can sometimes trigger.

The temporal structure of the session is organised as a linear narrative arc, as illustrated in the storyboard timeline embedded within the system design. The session opens with a human-voice meditation guidance segment, during which the user is gently welcomed, invited to find a comfortable position, and guided through the initial relaxation of the face, neck, and

shoulders. This opening segment serves a dual function: it establishes the tone of the experience and begins the process of directing the user's attention inward and away from external concerns. The background throughout this phase consists of soft natural soundscapes—wind, leaves, and birdsong—which provide an acoustic anchor that is both undemanding and subtly restorative.

Following the introductory phase, the core breathing guidance loop begins. The script unfolds in stages, first drawing the user's attention to the movement of light within the virtual scene—"Now look at the light in front of you while focusing on your body"—before introducing the breathing rhythm explicitly: "Inhale for 4 seconds 1 2 3 4, and now hold it 1 2 3 4, exhale for 4 seconds 1 2 3 4, and hold it again 1 2 3 4." This loop is repeated a designated number of times (Loop × N), with the voice gradually stepping back in subsequent cycles and encouraging the user to follow the environmental rhythm independently: "Now try on your own—just follow the light, it will guide you." This progressive withdrawal of explicit instruction is intentional: it mirrors the psychological process of internalisation, in which external guidance is gradually replaced by an embodied, automatic response.

The session closes with a brief re-orientation phase, during which the user is gently brought back to awareness of their physical surroundings: "When you feel ready, you can remove the VR headset." This transitional moment is designed to support smooth re-entry into ordinary consciousness, avoiding the abrupt perceptual discontinuity that can occur when immersive experiences end without a structured closure sequence.

This cyclical structure is both repetitive and predictable by design, which helps reduce anxiety and supports the gradual formation of a relaxed state. As the interaction continues, users progressively synchronise with the system's rhythm, moving from a phase of conscious, deliberate control toward one of unconscious, effortless regulation. By weaving breath guidance directly into the environment rather than presenting it as a separate layer of instruction, the system reduces cognitive interference and deepens the sense of immersion. The result is a unified experience in which physiological regulation and perceptual engagement are no longer two distinct processes but a single, continuous whole.

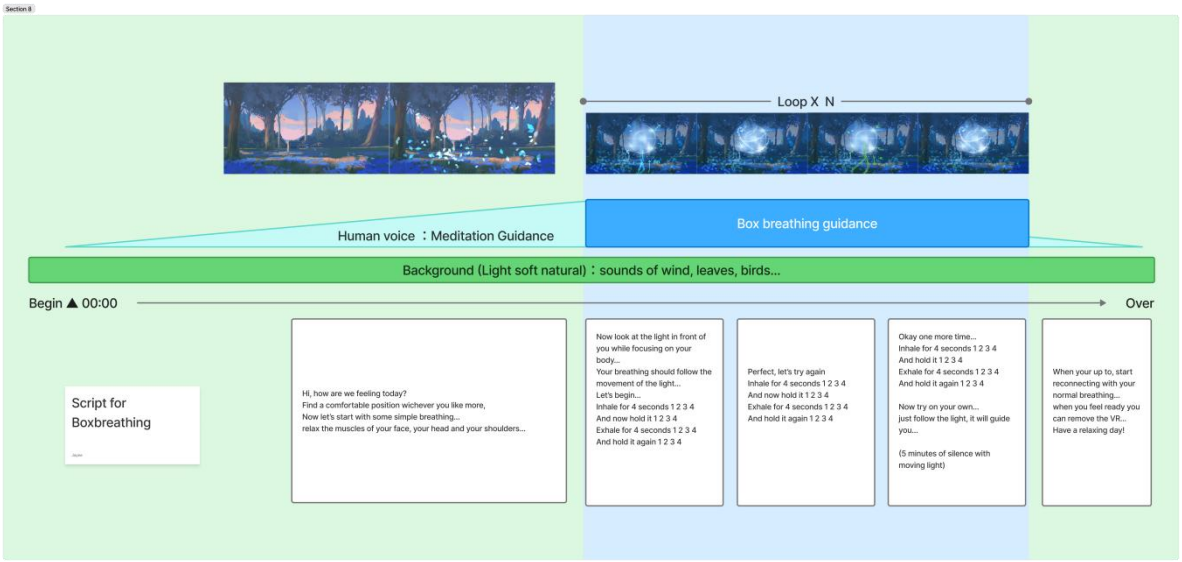


Figure 3.7 — Session Timeline and Script Structure of the Box Breathing Guidance Protocol

### *3.2.3. Heart-Rate-Driven Environment Transitions*

This section focuses on the construction of the immersive spatial experience, treating the virtual environment as the core medium of psychological intervention rather than a space for visual display alone. The space is designed as a dynamic experience system that responds to the user's physiological state in real time. By integrating sensory input, bodily engagement, and interactive feedback, it shifts the user's role from that of a passive observer—someone who watches—to that of an active participant who is fully present within the environment, whose internal state shapes the world around them, and whose experience of that world in turn shapes their internal state.

One of the core mechanisms underlying this dynamic is a heart-rate-driven scene transition logic that establishes a direct and continuous mapping relationship between physiological data and environmental change, enabling real-time biofeedback between the user's body and the virtual world. The system continuously monitors the user's heart rate via the BITalino biosensor and converts the incoming data stream into a set of key environmental parameters that actively shape the virtual space—including movement intensity, visual density, particle activity, colour temperature, and rhythmic variation. As the user's heart rate rises or falls, these variables respond in kind, creating a living environment that is neither static nor pre-scripted, but genuinely and continuously reactive to the physiological state of the person inhabiting it.

The scene transition logic operates according to a clear and physiologically grounded set of principles, as illustrated in the system design diagram. When the user's heart rate is

elevated—indicating a state of heightened arousal or residual stress—the environment remains relatively calm and visually stable, deliberately minimising the introduction of additional stimulation that might compound the user's existing activation level. This design decision reflects an understanding of how the nervous system processes environmental input under stress: when arousal is already high, rich or rapidly changing visual stimuli can intensify rather than reduce the stress response. By keeping the environment quiet and undemanding at this stage, the system creates the conditions for the user's autonomic nervous system to begin its natural descent towards a lower arousal state without interference.

As the user's heart rate gradually decreases—a physiological signal indicating that the parasympathetic nervous system is becoming more dominant and that the body is entering a recovery phase—the virtual environment begins to respond with increasing richness and dynamism. Dynamic visual elements such as softly glowing light particles, rhythmically swaying foliage, and gently pulsing luminescence gradually emerge within the scene, their movement synchronised with the pace of the user's slowing breath and heartbeat. Colour temperatures shift towards warmer, softer tones. Spatial density becomes lighter and more open. And the overall rhythmic quality of the environment settles into a slower, more fluid cadence. The environment, in this way, mirrors and reinforces the user's internal movement towards calm—not as a passive reflection, but as an active amplifier of the relaxation process already underway within the body.

In terms of transition strategy, the system deliberately adopts gradual, incremental change rather than abrupt environmental shifts. Sudden transitions would risk introducing sensory disruption at precisely the moment when the user most requires continuity and perceptual stability. Smooth transitions, by contrast, preserve the integrity of the immersive experience and allow the body to register environmental change without being startled out of its ongoing regulatory process. This commitment to continuity is central to the system's therapeutic logic: immersion is not merely an aesthetic quality but a functional precondition for effective physiological regulation.

Taken together, this mechanism builds a closed-loop physiological feedback system in which the user's own body becomes the primary driver of the experience. Rather than watching a fixed meditation scene unfold at a predetermined pace, the user is engaged in an ongoing and reciprocal dialogue with the environment—their nervous system and the virtual world continuously informing and responding to one another. Crucially, this feedback loop also serves a perceptual and educational function. By externalising invisible internal physiological changes into tangible, perceivable shifts in the surrounding environment, the

system strengthens the user's capacity for interoceptive self-awareness. Users begin to notice the connection between their own physiological state and what the environment is doing around them, which in turn encourages more intentional and conscious engagement with the regulation process. Over repeated sessions, this awareness has the potential to deepen into a genuine self-regulatory skill—an ability to recognise and consciously work with one's own physiological signals that extends well beyond the duration of the VR experience itself.

## 3.3. Virtual Environment Design

This section focuses on how psychological research directly shaped the VR scene design. Working in close collaboration with students from a psychology background, their findings on perception, emotion, and stress recovery were treated as design constraints rather than simply background references — meaning they were not consulted after the fact to validate decisions already made, but were built into the design logic from the start, actively determining what the environment could and could not do.

This distinction matters. In many design projects, research plays a legitimising role: it is cited to support choices that were already aesthetically or intuitively driven. In this project, the relationship was reversed. Psychological findings on attentional restoration, autonomic nervous system regulation, and the role of environmental predictability in reducing anxiety were introduced at the earliest stages of the design process, before any visual or spatial decisions had been finalised. They functioned not as endorsements but as parameters — setting the boundaries within which creative decisions were then made.

For instance, research on stress recovery consistently points to the restorative value of natural environments characterised by low complexity, soft movement, and the absence of sudden or unpredictable stimuli. These findings directly constrained the visual language of the scenes: they ruled out sharp contrasts, rapid transitions, and visually dense compositions, and guided the team towards open spatial structures, gentle particle motion, and warm, diffused lighting. Similarly, findings on the relationship between environmental controllability and perceived safety informed the decision to make scene transitions gradual and physiologically responsive rather than pre-scripted and fixed. The user's sense that the environment was moving with them — rather than imposing a predetermined rhythm upon them — was understood from the outset as a psychological necessity, not an optional design feature.



Figure 3.8 — UE5 Rendering of the Cardia Nocturnal Forest VR Environment

### *3.3.1. Landscape Categories*

This section classifies and defines the different environment types used in the VR experience—including natural landscapes, abstract spaces, and others—and analyses the distinct effects each type has on the user's psychological state. The selection of landscape categories in this project was not made on aesthetic grounds alone. It was grounded in an accumulating body of empirical research on the differential restorative effects of virtual natural environments across a range of psychological and physiological outcomes.

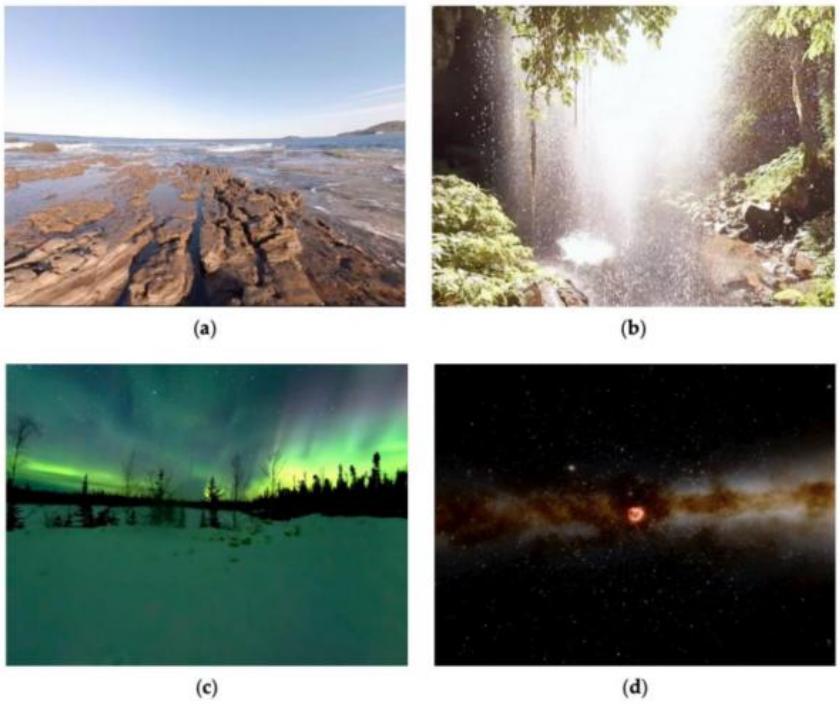


**Figure 2.** In-game view screenshots of the different 360-degree scenarios. (**a**) Beach. (**b**) Cascade cave. (**c**) Aurora Borealis. (**d**) Space.

Figure 3.9 Source: Vaquero-Blasco, M. A., Perez-Valero, E., Morillas, C., & Lopez-Gordo, M. A. (2021). Virtual reality customized 360-degree experiences for stress relief.Sensors, 21(6), 2219.

Research comparing multiple landscape types has consistently demonstrated that not all natural environments produce equivalent stress-relief effects, and that the specific configuration of a scene—including its category, spatial openness, the presence of water elements, and colour palette—plays a significant role in determining its restorative potential. One large-scale study examining the perceived restorative levels of different VR landscape types found a clear hierarchy of effectiveness: grassland environments ranked highest, followed by forest, water, undergrowth, and urban square settings in descending order[51]. The same study found that environments characterised by higher openness, greater presence of green and blue elements, and more exposure to sky and sunlight consistently produced stronger reductions in perceived anxiety, depression, and stress. In terms of colour preference, green was identified as the most relaxing landscape colour for the majority of participants, followed by blue and purple—a finding that aligns with chromotherapy research showing that cool-toned environments, particularly blue-dominant ones, are associated with measurable reductions in heart rate and blood pressure[52].

Within forest environments specifically, research has shown that different sub-types of woodland setting produce meaningfully different outcomes. A study exposing participants to seven distinct forest resting environment types via VR video found that purely naturalistic woodland—containing only plants, soil, and sky—did not outperform environments that included a degree of designed or semi-artificial intervention. Forest scenes incorporating wooden platforms, benches, or dynamic water features such as waterfalls[53] demonstrated superior stress-relief effects on both physiological and psychological measures. Dynamic water landscapes in particular emerged as especially effective, producing significant reductions in stress both physically and psychologically.




**Figure 1.** The videos of seven different types of forest resting environments: (**1**) structure, (**2**) wood, (**3**) wood with bench, (**4**) wood with platform and bench, (**5**) platform with trees, (**6**) waterfall with trees, and (**7**) pool with plants.

Figure 3.10 The videos of seven different types of forest resting environments: (1) structure, (2) wood, (3) wood with bench, (4) wood with platform and bench, (5) platform with trees, (6) waterfall with trees, and (7) pool with plants.

Source: Wang, X., Shi, Y., Zhang, B., & Chiang, Y. (2019).

These findings collectively guided the landscape selection strategy adopted in this project. Rather than defaulting to the most visually striking or intuitively "natural" environments, the design prioritised scene types with demonstrated restorative efficacy: open spatial structures, soft green and blue colour temperatures, the presence of gentle dynamic movement such as light particles , and a low degree of visual complexity. Abstract chromotherapy-influenced environments were also considered, with research confirming that dynamic animated elements—such as particle systems and fluid motion—enhance users' sense of presence and immersion relative to static visual conditions, thereby strengthening the overall therapeutic effect of the experience[54].

### *3.3.2. Colors and Visual Language*

The colour system and visual language of the Cardia environment were designed with a single overarching objective: to enhance the user's relaxation experience, emotional regulation capacity, and sense of immersion through deliberate and evidence-informed chromatic and stylistic choices. Rather than treating colour as a purely aesthetic variable, this section

approaches it as a functional design parameter—one whose psychological and physiological effects on the user are as consequential as the spatial or interactive dimensions of the experience.

The empirical basis for colour selection in therapeutic VR environments has been established across several converging bodies of research. Chromotherapy studies have consistently demonstrated that blue-dominant environments produce measurable calming effects on the autonomic nervous system, including reductions in heart rate and blood pressure. Research comparing empty, static, and dynamic virtual chromotherapy room conditions found that blue environments—whether presented as a plain monochromatic field or enriched with dynamic animated elements such as star particle systems, fractals, and ocean flow—reliably induced relaxation responses in participants, with the dynamic condition additionally enhancing users' sense of presence and immersion relative to static alternatives[55]. Particularly relevant to the visual design of the Cardia breathing orb is the dynamic condition illustrated in Figure 1 of that study, in which a luminous circular halo rendered against a blue background and surrounded by radiating particle arrays served as the central visual element of the virtual chromotherapy room. The structural similarity between this element and the Cardia breathing orb is not coincidental: both employ a glowing circular form positioned at the centre of the user's field of view, using particle-based animation to create a sense of rhythmic expansion and contraction. This design precedent provided direct visual evidence that a luminous orb form within a blue-dominant environment is not only aesthetically viable but functionally effective in sustaining user attention and deepening immersion within a therapeutic VR context**.**A complementary study further confirmed that immersive 360-degree VR nature experiences with blue-toned visual environments produced significantly greater reductions in total negative affective state compared to standard relaxation methods (Vaquero-Blasco et al., 2021). Landscape restoration research similarly identifies blue as the second most commonly preferred restorative colour among participants after green, with cool-toned environments in general associated with greater ease of emotional relief and relaxation[56].

These findings collectively informed the decision to adopt a blue-violet palette as the primary chromatic register of the Cardia scene. As discussed in the preceding section, the shift away from the green-dominant palette recommended by landscape restoration research was necessitated by the functional requirements of the breathing orb—the central interactive element of the system. In a nocturnal, blue-violet environment, the luminous breathing orb achieves the contrast and visual hierarchy necessary for it to function as an unambiguous focal point, while the surrounding scene recedes into a supportive and non-competing

background. The colour temperature of the environment was calibrated to remain in the cool range throughout the experience, with subtle shifts towards warmer violet tones introduced during moments of deeper relaxation to signal and reinforce the user's transition to a calmer physiological state.




Figure 3.11 — Three Phases of Breathing Orb Visual State in the Cardia Blue-Violet Environment

Beyond the base palette, the visual language of the Cardia environment is characterised by three additional principles derived from the research literature. The first is luminance stratification: the scene is structured around a clear hierarchy of brightness levels, in which the breathing orb functions as the primary light source, mid-ground environmental elements such as softly glowing particles occupy a secondary luminance register, and the background landscape—rendered in deep, low-saturation blues—provides a visually recessive foundation. This stratification ensures that the user's attentional resources are naturally drawn and held by the regulatory element without requiring conscious redirection. The second principle is dynamic restraint: while the chromotherapy literature confirms that dynamic visual elements enhance presence and engagement, it also cautions that excessively complex or rapidly changing animations can increase cognitive load and undermine the parasympathetic conditions necessary for relaxation. The Cardia environment therefore employs dynamic elements—particle drift, gentle luminous pulsation, and slow background motion—that are perceptible enough to sustain immersion but slow and soft enough not to compete with the breathing rhythm for the user's attention. The third principle is chromatic coherence: all visual elements within the scene, including particles, lighting, and environmental textures, are drawn from a unified cool-toned palette, ensuring that no element introduces a colour contrast strong enough to trigger involuntary attentional capture or disrupt the overall affective tone of the experience.



Figure 3.12 — UE5 Rendering of the Cardia Forest Scene Demonstrating Luminance Stratification

Taken together, these decisions produce a visual environment that is simultaneously legible, immersive, and therapeutically coherent. The colour system does not merely set a mood. It actively structures the user's perceptual experience in ways that support the physiological and psychological regulation goals of the Cardia system, functioning as an integral component of the intervention rather than its decorative surface.

### *3.3.3. Viewing Distance and Spatial Perception*

Through optimising viewing distance and spatial layout, this section addresses how to improve user comfort, reduce cognitive load, and enhance the stability of the immersive experience. While colour and visual language define the affective tone of the Cardia environment, the spatial parameters of the scene—specifically how far visual elements are positioned from the viewer and how open or enclosed the perceived edges of the environment appear—exert an equally significant influence on the user's physiological and psychological comfort within the virtual space.

Empirical research into the spatial configuration of restorative VR landscapes provides direct guidance for these design parameters. A systematic study examining the effects of viewing distance and edge permeability on anxiety, depression, and stress constructed a matrix of VR landscape scenes varying viewing distance across three conditions (20 m, 100 m, and 200 m) and edge permeability—the degree to which the visible perimeter of the scene is open versus enclosed by vegetation or other boundary elements—across three corresponding

levels (30%, 70%, and 100%). The results of this study indicated that the optimal viewing distance for restorative virtual landscapes lies between 140 and 160 metres, a range that produced significantly lower perceived anxiety than either very close (20 m) or distant (200 m) viewing conditions (Shu et al., 2022). For edge permeability, the same study found that the most restorative environments maintained between 52% and 57% visible perimeter openness, with statistically significant differences in anxiety perception observed between low-permeability (30%) and high-permeability (100%) conditions, as well as between moderate (70%) and fully open (100%) conditions. Critically, the research also established that fully enclosed spaces—those approaching 0% permeability—were consistently less effective for emotional restoration than environments offering at least a moderate degree of openness, suggesting that excessive spatial confinement undermines rather than enhances the calming effect of a virtual landscape.

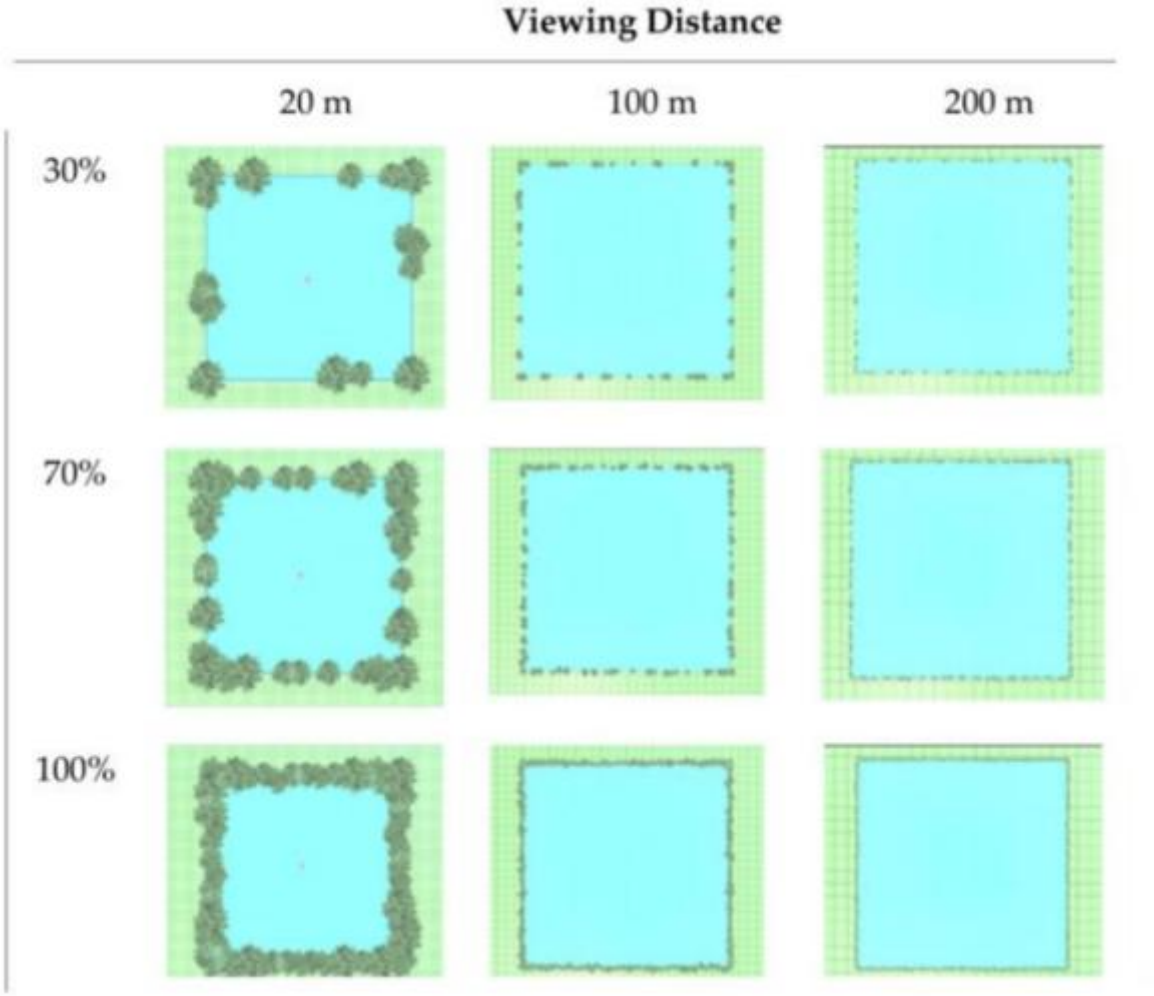


Figure 3.13 VR landscape scene design plans: schematic diagrams of viewing distance (20 m, 100 m, 200 m) and edge permeability (30%, 70%, 100%) conditions.

Source: Shu, Y., Wu, C., & Zhai, Y. (2022).

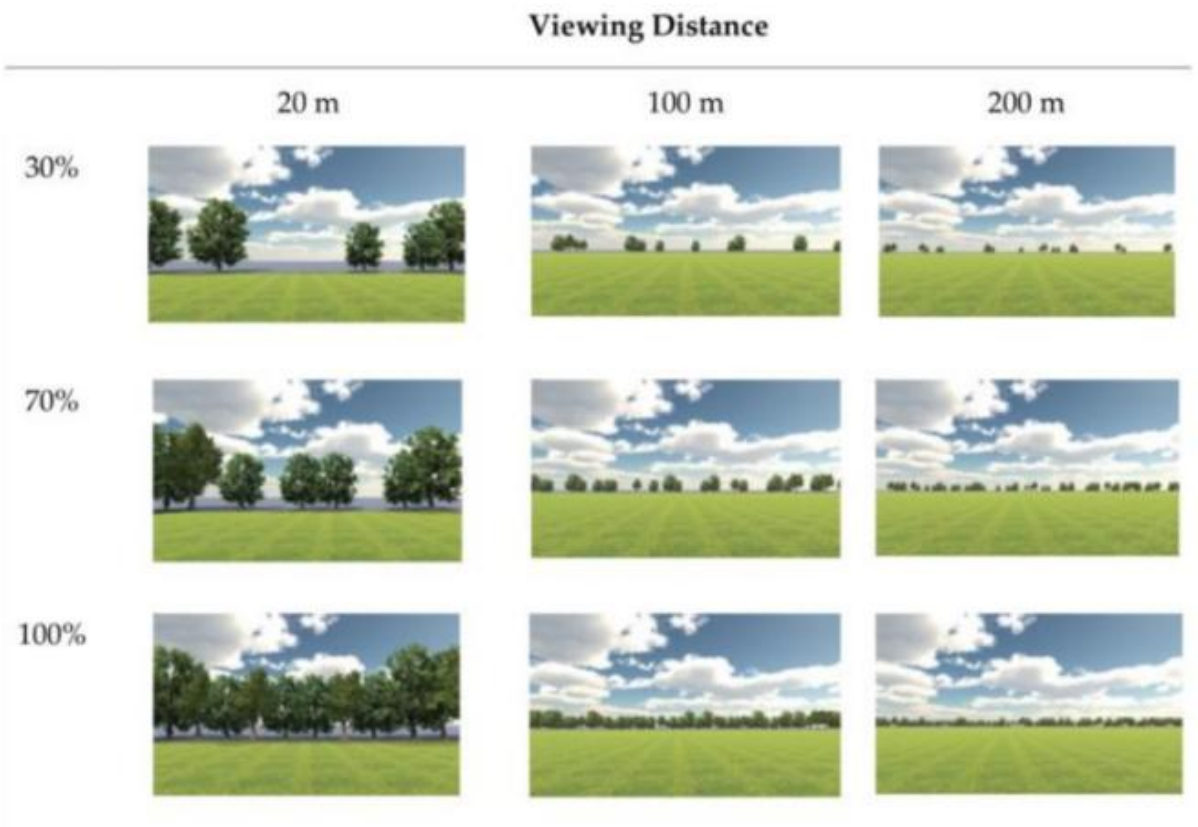


Figure 3.14 VR landscape scene renderings: photographic stimuli corresponding to viewing distance and edge permeability conditions. Source: Shu, Y., Wu, C., & Zhai, Y. (2022).

These findings carry direct implications for the spatial design of the Cardia environment. Although the Cardia system employs a forest setting rather than the grassland landscapes used in the source study, the underlying spatial principles—moderate viewing distance and partial rather than complete edge openness—were adopted as guiding parameters for scene construction. The virtual forest was therefore designed with a sufficiently extended sightline to avoid the sense of visual crowding associated with close-distance viewing, while retaining enough surrounding tree density at the periphery to provide a perceptible, partially enclosed boundary rather than an entirely open or entirely walled environment. This intermediate configuration was selected specifically to avoid two failure modes identified in the literature: on one hand, an overly open or distant scene risks producing a sense of spatial disorientation or emotional flatness, weakening the user's sense of presence. On the other hand, an overly enclosed scene risks inducing a subtle sense of confinement that could counteract rather than support relaxation, particularly in a seated, head-mounted VR context where the user's physical mobility is already constrained.

Beyond viewing distance and edge permeability, spatial perception within the Cardia environment was further informed by considerations specific to the breathing guidance mechanism. Because the central breathing orb occupies a fixed position at the centre of the user's field of view throughout the session, the surrounding landscape was deliberately structured to recede gently towards the periphery, with foreground elements—such as the rhythmically swaying grass discussed in Section 3.3.1—positioned at a closer perceptual distance to establish depth, while background tree lines were placed further away to approximate the empirically supported optimal viewing range. This layered spatial arrangement ensures that the user's peripheral vision remains gently populated with low-intensity environmental detail, sufficient to sustain a stable sense of immersion, without introducing visual elements close or dense enough to compete with the central breathing orb for attentional resources.

Taken together, the spatial design strategy adopted for the Cardia scene reflects a deliberate calibration informed by empirical restoration research: a moderately distant, partially open forest environment that balances the user's need for a stable, comfortable, and non-claustrophobic sense of space with the system's functional requirement to maintain the breathing orb as the unambiguous perceptual and attentional centre of the experience.

### *3.3.4. Natural Environment Typologies*

By incorporating natural elements and the principles of Biophilic Design, this section addresses how the categorisation and configuration of natural environment types within the Cardia scene can enhance users' psychological comfort and sense of engagement. While the preceding sections have established the colour system, viewing distance, and edge permeability that structure the Cardia environment, the underlying typology of natural elements — which categories of landscape are present, and in what degree of openness — constitutes a further layer of evidence-informed design decision-making that shapes the overall restorative quality of the experience.

Empirical research comparing the restorative effects of different urban park landscape categories provides direct support for this approach. A large-scale study cBy incorporating natural elements and the principles of Biophilic Design, this section addresses how the categorisation and configuration of natural environment types within the Cardia scene can enhance users' psychological comfort and sense of engagement. While the preceding sections have established the colour system, viewing distance, and edge permeability that structure the Cardia environment, the underlying typology of natural elements—which categories of

landscape are present, and in what degree of openness—constitutes a further layer of evidence-informed design decision-making that shapes the overall restorative quality of the experience.

Biophilic Design, as a guiding theoretical framework, holds that human psychological wellbeing is enhanced through direct and indirect contact with natural elements and patterns, reflecting an evolved affinity for environments that historically signalled safety, resource availability, and habitability. This framework provilassified natural environments into six representative categories—grassland, water, forest, undergrowth, urban square, and a residual "other" category encompassing less common landscape types—and further subdivided each category according to degree of spatial openness (high, medium, and low). Across these typologies, participants consistently reported the strongest restorative effects in three categories: grassland, forest, and water landscapes, identifying these as producing the greatest sense of psychological relief and emotional engagement[57]. When asked to indicate which landscape category they associated most strongly with relaxation and recovery, respondents ranked grassland highest (31.2%), followed by forest (27%), water (23.4%), undergrowth (12.5%), and urban square settings (5.8%), establishing a clear hierarchy of restorative effectiveness across landscape types.

With respect to spatial openness, the same study found that high and medium degrees of openness were perceived by participants as equally restorative, while fully closed or low-openness environments were consistently less preferred, reinforcing the principle—already established in Section 3.3.3 regarding edge permeability—that excessive spatial enclosure undermines rather than enhances the restorative quality of a natural environment. Approximately 56% of participants expressed a preference for high-openness environments, with a further 34% favouring medium openness, suggesting that the great majority of users derive psychological benefit from landscapes that avoid a sense of confinement. Regarding colour, the study further confirmed that cool colour tones—particularly green, blue, and purple—were the hues most commonly associated with relaxation, corroborating the chromatic principles already adopted in the Cardia colour system discussed in Section 3.3.2.

These findings carry direct implications for the categorical structure of the Cardia natural environment. The forest typology selected for the primary Cardia scene falls within the top tier of restorative landscape categories identified in this research, providing strong empirical justification for its selection over less restorative alternatives such as urban square or undergrowth settings. Moreover, the inclusion of dynamic water elements—as discussed in Section 3.3.1 regarding the particular restorative power of moving water—aligns the Cardia

environment with water landscapes, the third most highly rated category in this typology, allowing the scene to draw simultaneously on the restorative strengths of two of the highest-performing landscape types rather than relying on a single category alone. The degree of openness adopted within the forest scene was likewise calibrated towards the medium-to-high range identified as most broadly preferred, ensuring that the environment avoids the perceptual confinement associated with low-openness or fully enclosed forest configurations.

Taken together, the typological design of the Cardia natural environment reflects a synthesis of evidence-based landscape categorisation, deliberately combining forest and water elements within a moderately open spatial configuration and a cool-toned colour palette. This integrated approach allows the scene to draw simultaneously on multiple converging sources of restorative effect—landscape category, spatial openness, and colour—rather than depending on any single design variable in isolation, thereby strengthening the overall therapeutic coherence of the Cardia experience.

| Landscape category | Degree of openness | | |
|---|---|---|---|
| | High | Medium | Low |
| Grassland |  A1 |  A2 |  A3 |
| Water |  B1 |  B2 |  B3 |
| Forest |  C1 |  C2 |  C3 |




Figure 3.15 Classification matrix of natural landscape categories (grassland, water, forest, undergrowth, urban square, others) and degrees of spatial openness (high, medium, low).

Source: Ma, J., Li, H., Zhao, M., & Williams, J. (2025).

## 3.4. Project Development

### *3.4.1. Arduino-Based Physiological Data Processing*

This section introduces the physiological data acquisition and processing pipeline based on Arduino sensors — covering heart rate data collection, signal filtering, normalization, and parameter mapping methods within Unreal Engine. These data serve as the core input for driving the real-time dynamic feedback mechanism within the immersive system.

To link the user's body state with the feedback in the virtual world in real time, this study uses an Arduino Uno R3 with a Heartbeat Sensor to build a heart rate module. The heart rate data is then sent to Unreal Engine 5 (UE5) in real time through the SerialCom5 plugin. SerialCom5 is a UE plugin that allows serial port communication. Its main job is to open a real-time data channel between the game engine and outside hardware. When the system starts, UE5 first uses the Open Serial Port With Flow Control node to open the serial port, and connects with the Arduino at a speed of 115200 Baud Rate. Once the connection works, the

Arduino keeps sending the heart rate values from the Heartbeat Sensor to UE5 as text strings. The blueprint system then reads and processes this text, and uses it to drive the Niagara particle system to create visual feedback.

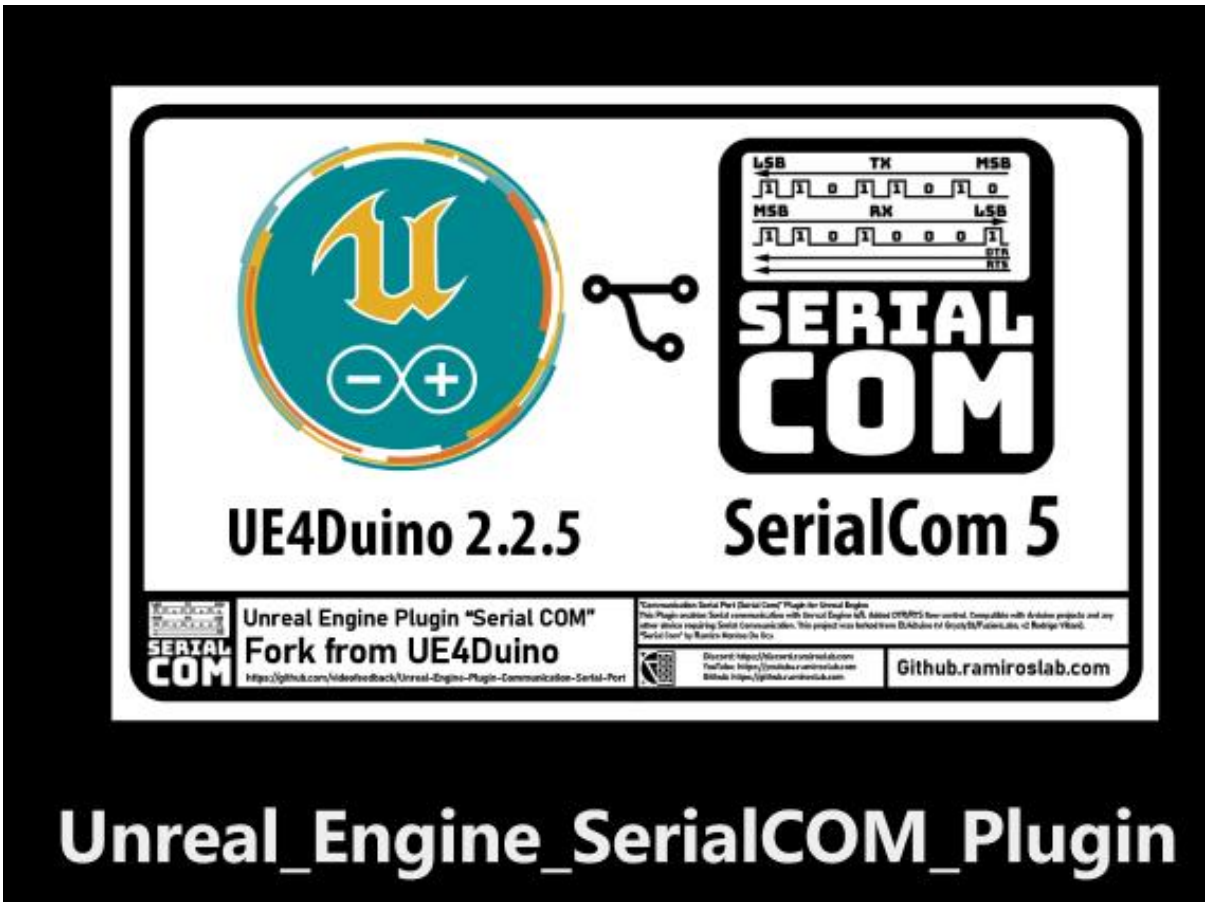


Figure 3.16 SerialCom5 Plugin Interface for Arduino–UE5 Serial Communication

However, real tests showed that the raw data from the Heartbeat Sensor was quite unstable. Compared with medical-grade equipment, a low-cost photoplethysmography (PPG) sensor is more affected by things like finger pressure, room light, small body movements, and sampling errors. Because of this, the system often showed sudden jumps, strange peaks, and missing data during long monitoring sessions. For example, even when a user stayed still, the heart rate readings could jump a lot in a short time, and some readings even dropped to zero or came back empty. These odd values do not show the user's real body state. They come from sensor noise and small errors in the serial connection.

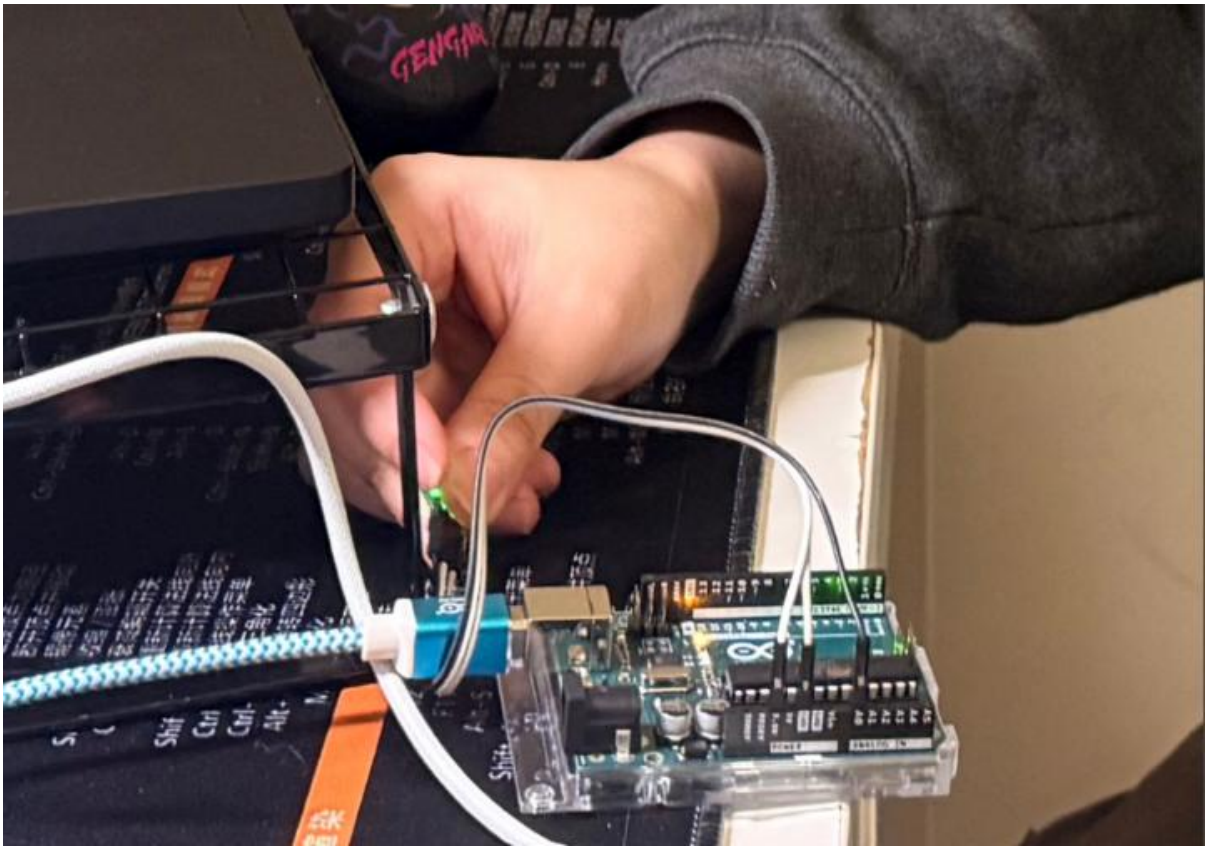

Figure 3.17 — Photograph of Arduino Uno R3 Connected to the Heartbeat Sensor

This kind of unstable data has a direct effect on the VR meditation system built in this study. The goal of the project is to use nature scenes and particle effects to create a calm, immersive, and healing mood, so the visual feedback needs to change in a soft and gradual way. If the raw heart rate data were sent straight to the particle system, any sensor noise would cause sudden changes in the number of particles, their movement, or their glow. These sudden changes could easily break the user's focus and lower the sense of immersion and comfort during meditation. For this reason, it is important to smooth out the raw data before it reaches the visual feedback system, so that noise has less effect on the virtual world.

To solve these problems, this study built a data buffer and smoothing system inside the UE5 blueprint for body feedback. The main goal is not just to make data collection more exact. Instead, while still keeping the body signal real and honest, the goal is to turn the noisy raw heart rate data into smooth visual control values that work well for an immersive meditation space. Because this project is a VR system built for mental healing, users do not need to know their exact heart rate number. What matters more is that they feel their own body state changing through soft, steady, and rhythmic changes in the environment. So unlike

medical monitoring, which needs exact numbers at every instant, this study cares more about how usable and continuous the body data feels during the experience.

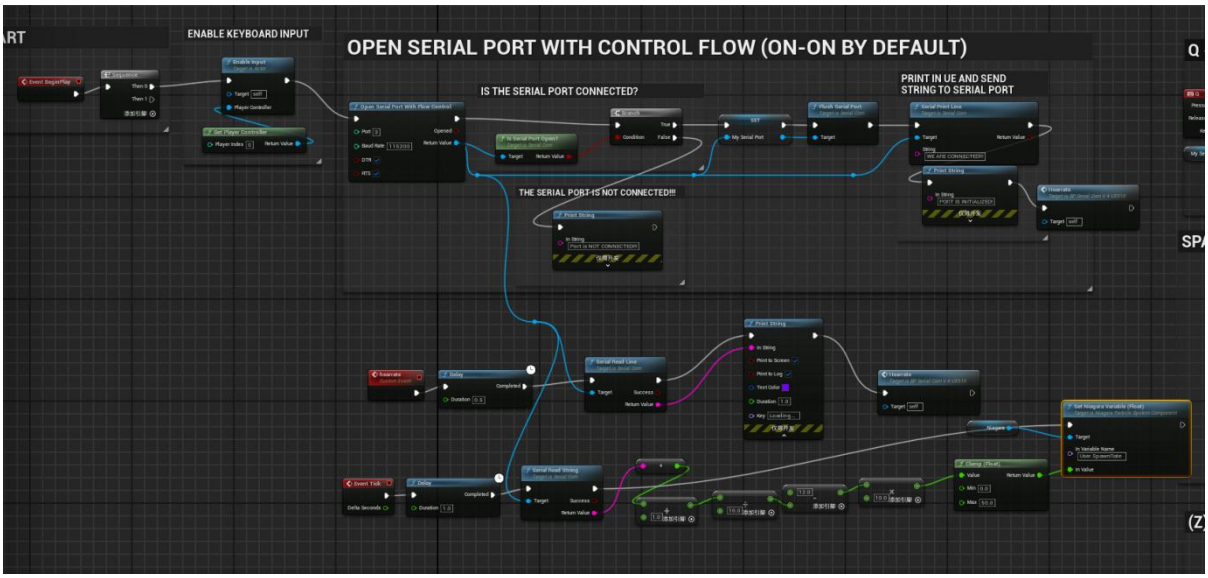


Figure 3.18 — UE5 Blueprint: Serial Port Initialization and Heart Rate Data Reception Logic

While the system runs, the Arduino keeps receiving the pulse signal from the Heartbeat Sensor and sends it to UE5 through the serial port. Since the Arduino sends data as plain text, UE5 first uses the Serial Read String node to read the current data from the serial buffer, then turns it into a float number for later math. Unreal Engine usually updates at 60 frames per second or more, but the heart rate sensor updates much slower than that. If UE5 tried to read the serial port every frame, it would add extra work for the system and could read the same value more than once, which would make the noise even worse. To avoid this, a Delay node was added to the blueprint to control how often new data is read. By setting a fixed time gap, the system reads new body data at a steady rate. This works like a buffer layer between the data source and the visual feedback, so the data update speed matches the pace of the scene, and high-frequency noise has less effect on the user.

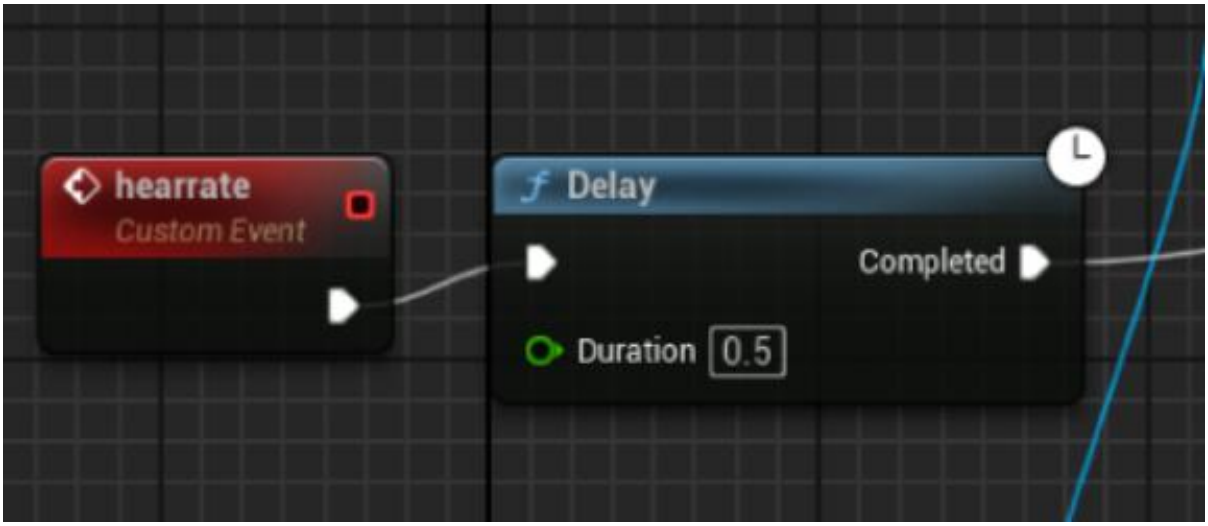


Figure 3.19 — UE5 Blueprint: Delay Node Configuration for Sampling Rate Control

After reading the data, the system normalizes the raw heart rate value. The number range from the sensor does not match the range needed by the particle system, so mapping the heart rate value straight to a visual variable makes it hard to control how sensitive the feedback is. For example, a normal resting heart rate is usually between 60 and 90 BPM, and it can rise above 100 BPM under light stress. If the raw numbers were used to drive the particle system directly, small body changes might be too small to notice, while big swings could create feedback that looks too extreme. To fix this, the study uses a simple linear normalization method to remap the data. Subtraction, division, and multiplication are used together to turn the raw heart rate range into a standard control range that fits the Niagara system. After this step, the heart rate differences between different users are scaled to the same range. This lets the system respond more steadily to changes in the user's state, while keeping the visual feedback predictable and consistent.

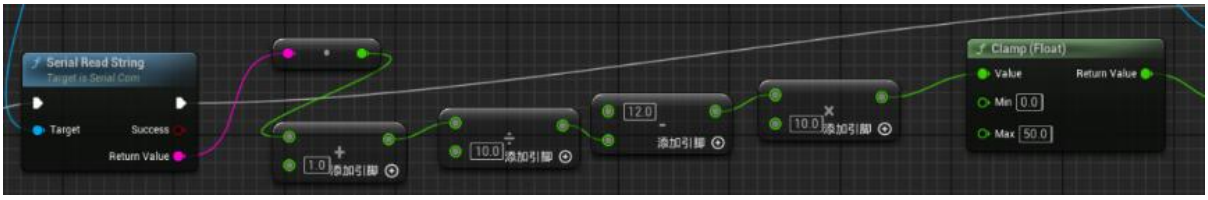


Figure 3.20 — UE5 Blueprint: Linear Normalization of Raw Heart Rate Data

Still, normalization alone could not fix the problem of strange jumps in the body data. During real tests, extreme values sometimes showed up when the contact between the user's finger and the sensor changed. For example, if the sensor briefly lost contact, the output could

suddenly drop to 0. If the signal was misread, it could also show a peak far above the normal range. These bad values only lasted a moment, but if they went straight into the visual calculation, they would still cause a sudden, obvious change on screen. To stop this, a Clamp node was added to the blueprint to limit the range of the input data. This node finds and filters out any value outside a reasonable range, and pulls it back inside a set limit. This way, even when the sensor gives a wrong reading once in a while, the error does not pass straight through to the visuals, which makes the whole interaction much more stable.

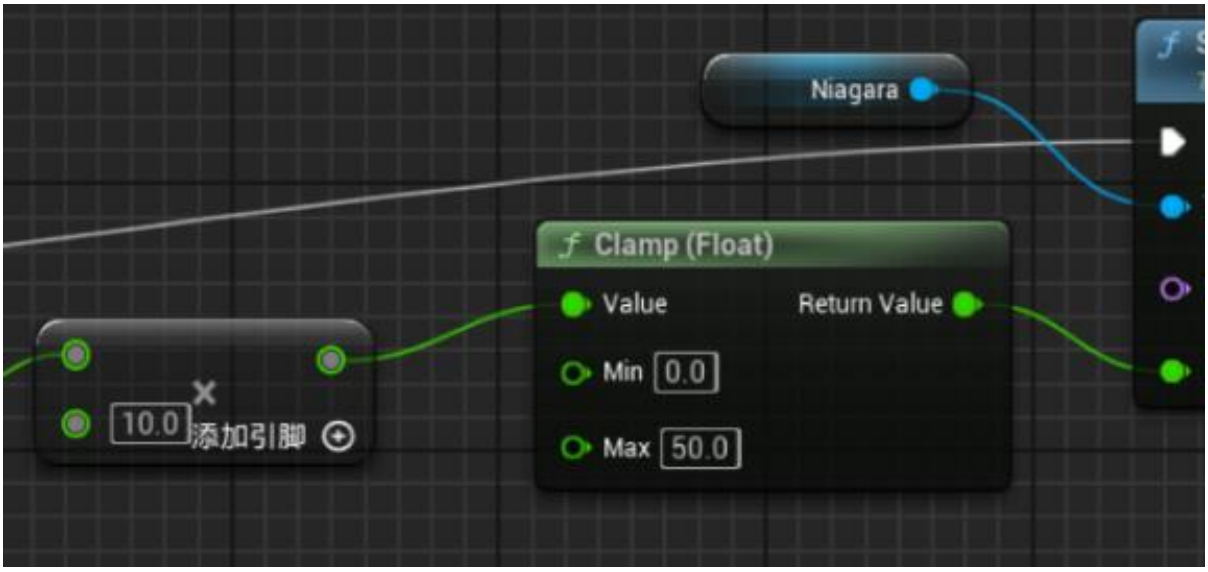


Figure 3.21 — UE5 Blueprint: Clamp Node for Outlier Filtering

Besides filtering out bad values, this study also uses the built-in interpolation feature of the Niagara particle system to add a second layer of smoothing on the visual side. When a new heart rate value is written into the particle system, the number of particles, how fast they spawn, and how they move do not switch to the target value right away. Instead, they change slowly over a short period of time. In other words, the system smooths the raw signal at the data level, and also smooths the feedback at the visual level. This two-layer smoothing stops sudden, jarring changes that would otherwise come from the body signal, so the feedback the user sees always feels smooth and natural.

Apart from connecting the data through the main blueprint, a Print String node was also added so that, during testing, it would be easy to see right away whether the data was reaching UE5. The Serial Read Line node keeps reading one line of data from the serial port, and its Return Value output is also linked to the In String pin of the Print String node, with both Print to Screen and Print to Log turned on. This shows the raw value as text on the UE editor screen

and in the output log in real time, so during development it is easy to check by eye whether the data keeps arriving, and whether there are any jumps or gaps. After printing, the flow continues through the Heart Rate custom event to drive the rest of the visual feedback.

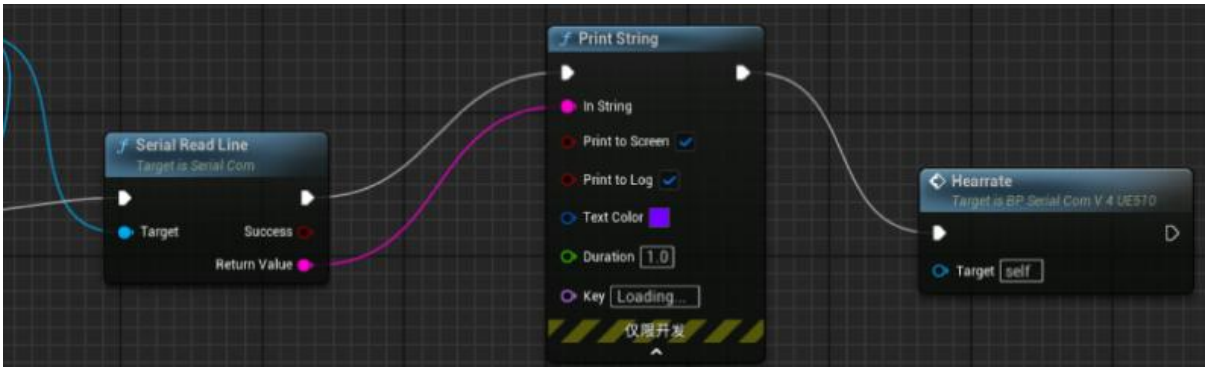


Figure 3.22 — UE5 Blueprint: Print String Node for Real-Time Data Debugging

In this study, the processed heart rate data is ultimately mapped to the Niagara particle system, where it is used to control key visual parameters such as particle spawn rate. The system adopts a progressive environmental feedback mechanism based on heart rate change: that is, as the user's heart rate gradually decreases and stabilizes, the number of particles in the scene increases step by step, and the dynamic elements within the space also become progressively richer. As particles continue to accumulate within the environment, the user can intuitively sense that the scene is responding to their own physiological state, thereby forming a process of environmental growth that is driven by bodily state. Compared with traditional biofeedback systems that present physiological information through numerical values or charts, this study translates physiological data into spatialized visual change, allowing an abstract bodily signal to be presented within the virtual environment in a more natural and immersive form. Users do not need to understand the specific heart rate value. They need only observe the environment gradually becoming visually richer in order to perceive the trend of change in their own state. This design transforms the physiological feedback process into an interactive experience with a sense of exploration and companionship, so that calmness is no longer merely a passively received outcome, but rather a process that can be responded to and reinforced by the environment.

From the perspective of interaction design, this mechanism not only achieves real-time visualization of physiological data, but more importantly, constructs a closed-loop interactive process of "perception–feedback–re-perception." When the user observes a change in the environment, they will actively adjust their breathing rhythm and bodily state, and these new physiological changes will, in turn, further influence the environmental feedback, thereby

forming a continuously cycling mechanism of self-regulation. This feedback logic differs significantly from traditional biofeedback systems that rely on textual prompts or numerical displays. The latter tend to emphasize the transmission of information, whereas this study places greater emphasis on the shaping of experience — that is, guiding the user, through environmental change, to gradually become aware of their own bodily state, and to complete the process of emotional regulation at a largely unconscious level.

From a technical implementation perspective, this module essentially functions as a Physiological Data Interpreter. Rather than simply mapping the sensor's output directly to a visual effect, it establishes a complete data-processing pipeline between data acquisition and visual presentation, encompassing multiple stages including serial communication, data reading, sampling control, normalized mapping, outlier filtering, and visual interpolation. Each processing step is designed to reduce the uncertainty inherent in the raw physiological signal and to improve the continuity and comfort of the final feedback effect. This is especially true in the VR meditation scenario, where users' perception of environmental change is often more sensitive than their perception of the numerical values themselves. The system therefore needs to prioritize the stability of feedback over the pursuit of absolute real-time numerical synchronization.

In summary, by constructing a multi-level data smoothing mechanism, this study successfully addressed the issues of data fluctuation, abnormal peak values, and signal loss present in the practical application of the Heartbeat Sensor, and achieved a stable mapping process extending from physiological signal acquisition to immersive visual feedback. Compared with an approach that drives the scene directly using raw heart rate data, this method significantly reduces visual disturbance and improves both the operational stability of the system and the continuity of the user experience, allowing physiological feedback to be integrated into the virtual environment in a manner that is more natural, gentle, and aligned with the requirements of a meditative scenario — thereby providing a reliable technical foundation for subsequent interaction design in psychological healing applications.

### *3.4.2. Niagara Particle Parameters for Biometric Mapping*

This section explains how the system was built to choose the right Niagara particle settings, so that sudden changes in body data would not cause problems in the visuals. Following the earlier design rules about color, viewing distance, and scene type, this study chose to link body data closely with the visual look of the virtual world. The main reason is this: the Cardia system's core healing idea relies on a real-time feedback loop between the

body and the environment. If the body data only existed as numbers stored in the background, with no clear and readable form for the user to see, then the idea of a “visible body,” which biofeedback depends on, would be weakened, and users would lose a way to see their own body state through the changing environment. So turning heart rate, the most useful body signal, into a visual change the user can notice in the virtual scene became a key part of making the whole system work as a closed loop.

To be exact, the design goal works like this: when the heart rate goes up, meaning the heart beats faster and the user feels more tense or alert, the number of particles in the scene should go down, and they should look less active, so the visuals do not add extra stimulation. On the other hand, when the heart rate goes down, showing that the user's body is becoming calmer and more relaxed, the particles in the scene should become more active and show richer movement. This richer visual movement matches and supports the positive change happening inside the user.

As for why a particle system was chosen to carry this visual feedback, it follows the earlier theory that weak, low-strength visual movement helps people relax more easily. Based on this theory, the study used the Niagara system in UE5 to create fireflies flying through the scene. From a design point of view, the particle system is spread out around the edges of the scene, unlike the glowing breathing guide ball and breathing guide line, which are the main focus of interaction. This way, the particles stay in the background and do not pull the user's attention away from breathing, which is the main interaction. Also, in terms of the story of the scene, fireflies fit naturally into a night forest setting. They give a reason for small bits of light in the environment, match the natural feel of a forest at night, and also make the whole scene look nicer and feel more atmospheric. In this way, the fireflies work well in terms of function, story logic, and looks, all at the same time.

Starting from how the Niagara particle system works, this study looked closely at a few key settings that could carry the body data, including the Spawn Rate setting in the Emitter module, the Lifetime Mode setting under Initialize Particle in the Particle Spawn module, and the Wind Force and speed-related settings in the Particle Update module. Together, these settings are the main controls that shape how particles look in Niagara, and they became the main candidates tested in the experiments that followed.

Of these, Spawn Rate sets how many particles appear in each unit of time. It directly affects how many particles are on screen and how rich the scene looks. A higher number means more fireflies appear, making the mood feel more active. A lower number means fewer particles, making the scene feel quieter. Lifetime Mode sets how long a single particle lasts,

from when it appears to when it disappears, and this can be a fixed number or a random range. This setting matters a lot for keeping the scene looking steady and natural. Wind Force and the speed-related settings together control how particles move and how fast, including their direction, how much they sway, and the overall pace of movement. These settings are usually seen as the ones most linked to how “active” a scene feels, since the human eye is quite sensitive to changes in movement speed. At first, the research team leaned toward mapping the body data straight to the speed settings, hoping that faster or slower particle movement would clearly show how high or low the user's heart rate was.

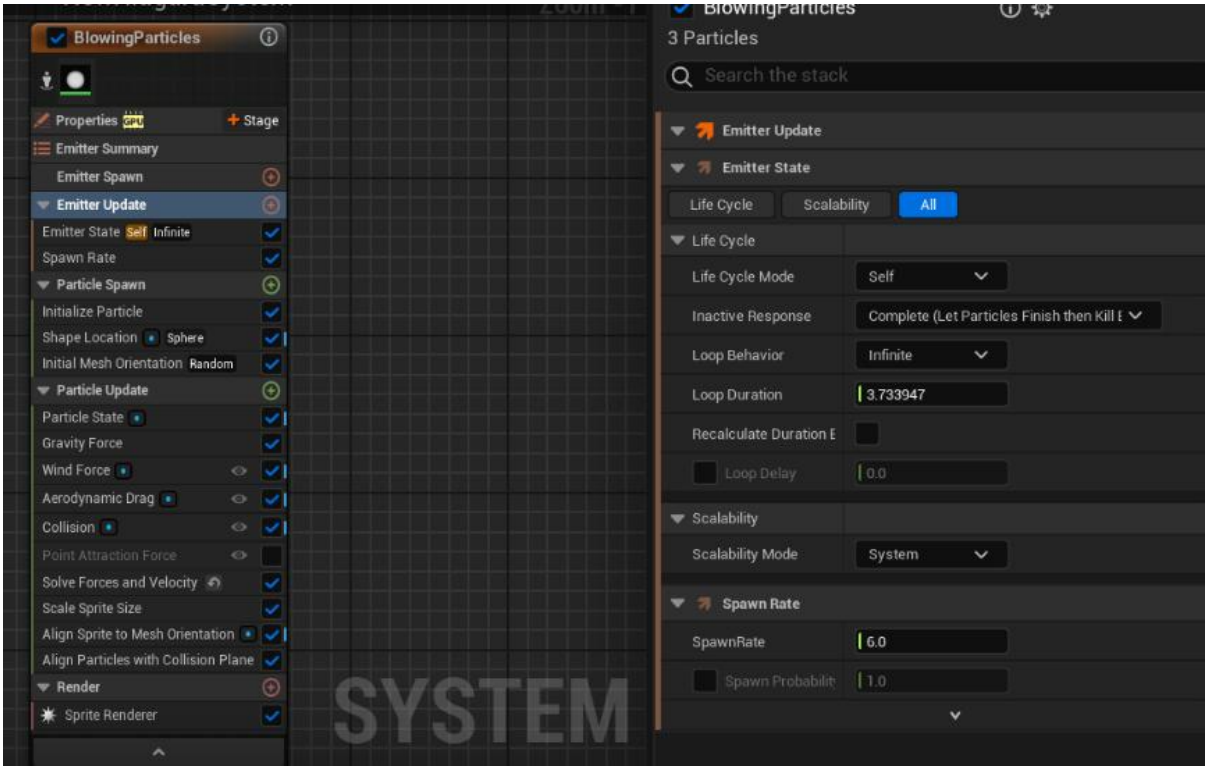


Figure 3.23 — Niagara Editor Interface Showing BlowingParticles Emitter Configuration and Spawn Rate Settings

However, repeated tests showed that the body data often changed suddenly and unevenly. The heart rate value did not change smoothly. Instead, it often jumped or shifted suddenly within a short time. Even though the goal was to show body state changes through how “active”the particles looked, and activity was most closely tied to the speed setting, linking this kind of sudden, jumpy data straight to speed caused two problems. First, the human eye usually cannot pick up small changes in how fast a particle moves within a short window of time. This means that even when the data changed, the user might not see a clear, readable signal, which weakened the whole point of the biofeedback. Second, and more important, if the data stream itself had a glitch, a delay, or an odd error, which is common with real sensor

data, then letting this unstable data drive particle speed directly could easily cause particles to suddenly freeze or jump during testing. This kind of broken, jumpy visual is the opposite of the steady, immersive feeling the system is meant to create, and it could even catch the user off guard, causing mild discomfort or breaking their focus without warning.

To solve these problems, this study chose to link the body data to the Spawn Rate setting instead of the particle speed setting. The core idea behind this change is simple: by keeping the Lifetime of every particle the same, and using a logic similar to math integration, sudden changes in the body data no longer cause a sharp, instant change in how a single particle moves. Instead, they turn into a slower, building-up change in the number of particles over time. In other words, when the heart rate shifts, the effect is not a particle suddenly speeding up or slowing down. Instead, it shows up as a gradual change in how fast new particles are created, and the slow rise or fall in particle count shows the change in the scene's overall “activity” in an indirect way. Mapping body data to spawn rate instead of directly to gravity or speed gives much more visual stability and room for error. Even if the underlying data has a short glitch or error, because the Spawn Rate change only shows up through new particles being added, rather than changing the motion of particles that already exist, the scene never freezes or jumps suddenly. This keeps the biofeedback working well while also keeping the whole visual experience smooth and steady, which fits the core need of a meditation experience for steadiness and calm.

On the technical side, as shown in the blueprint logic, the Event Tick node keeps listening for and reading the text data coming from the body sensor. A Delay node controls how often this data is read. The raw text is then turned into a float number, and a Clamp node keeps it within a safe range, such as 0 to 50, so that any odd value cannot push the particle system too far or make it look broken. The processed number is then passed through the Set Niagara Variable (Float) node into User. SpawnRate, a custom float variable already exposed in the Niagara particle system. This creates a real-time, dynamic link from the body data to the particle spawn rate. This technical path lets the body data keep shaping the firefly particles in the virtual forest in a way that looks soft and feels natural, completing the full loop from body signal to immersive visual feedback.

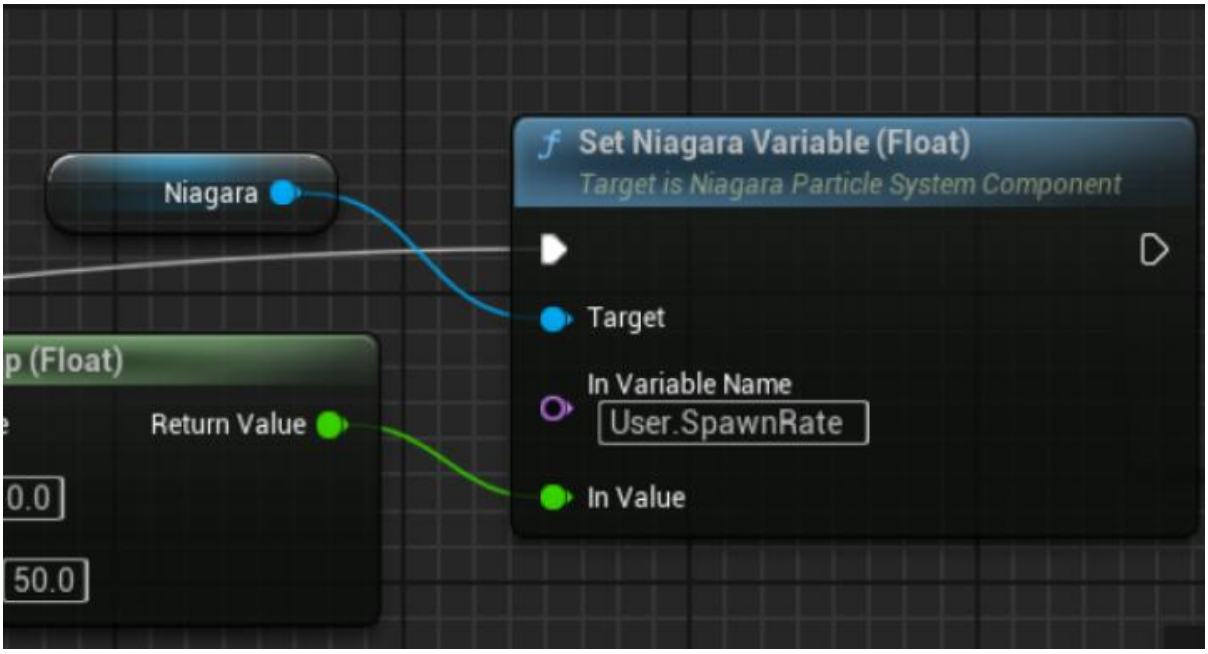


Figure 3.24 — UE5 Blueprint: Clamp Node Connected to Set Niagara Variable (Float) Node for Real-Time Spawn Rate Control

### *3.4.3. Rendering Performance Optimization*

During development and testing, putting on the headset and entering the scene showed clear stutter and dropped frames right away. This goes directly against the steady, calm feeling that a meditation experience needs. Stutter does not just break the sense of immersion. In a VR scene it can also cause dizziness and nausea, which goes against the core goal of the Cardia project: helping users reach deep relaxation while keeping the body's stress response as low as possible. Because of this, checking and improving the scene's rendering performance in a careful, step-by-step way became a necessary part of making the Cardia system usable and comfortable.

This chapter first looks at what frame rate means and how it relates to the quality of a VR experience, and lists the main things that affect frame rate in a UE5 project. It then uses the real stutter seen in the Cardia scene to find where the performance problems come from, which turn out to be mostly the Niagara particle system, the plant-swaying material, and the rendering and collision cost of the large number of plant objects in the scene. Finally, it lays out a set of reusable steps for improving performance in a scene like Cardia, and explains the trade-offs made, based on the design rules for meditation interaction.

Frame rate, often shown as FPS (Frames Per Second), is how many frames the rendering engine sends to the display each second. Making one frame needs three steps working together like a pipeline: the CPU runs the scene logic (the Game Thread, which handles things like Actor Tick, physics, and user input), the Render Thread sets up the drawing instructions,

and finally the GPU does the pixel-level shading work (the GPU Thread). Whichever of these three steps takes the longest sets the time needed for one frame, and this sets the upper limit on frame rate. Frame rate and frame time are opposite numbers: the higher the frame rate, the less time is allowed for each frame. For example, 60 FPS gives about 16.7 milliseconds per frame, while 90 FPS only gives about 11.1 milliseconds. In VR, the scene must be rendered twice, once for each eye, known as stereo rendering, so the actual workload is almost double that of a normal desktop scene. This means VR projects must manage their per-frame time budget much more strictly than typical desktop games or applications. Unlike desktop games, frame rate in VR is not just about how smooth things feel. It directly affects physical comfort. When someone wearing a headset turns their head, the vestibular system in the inner ear, which senses balance, picks up the movement right away. If the screen does not refresh fast enough to keep up with the head turn, the images the eyes see no longer match what the inner ear senses. This mismatch, called sensory conflict, is one of the main causes of VR motion sickness. Because of this, the VR industry holds frame rate to a higher and stricter standard than regular software. Most headsets lock the rendering frame rate to the display's refresh rate and expect the app to hit that native rate. If it cannot, the system falls back on methods like Asynchronous Timewarp or Spacewarp to fill in the gaps, but these can only soften the problem, not fully replace a truly stable frame rate.

Different VR platforms expect different frame rates. PC-based VR headsets such as the HTC Vive and Valve Index usually need 90 FPS or higher, giving about 11.1 milliseconds per frame. The Meta Quest 2 targets 72 FPS, or about 13.9 milliseconds per frame. The Meta Quest 3 targets 90 FPS, or about 11.1 milliseconds. Pico headsets target either 72 or 90 FPS, giving a frame budget of about 13.9 or 11.1 milliseconds.

For the headset platform actually used in the Cardia project, the rendering frame rate should ideally stay steady at the device's native refresh rate, with no ongoing frame drops. When the frame rate falls below the device's refresh rate, for example below 72 FPS, three common problems tend to show up. The first is stutter, where the picture jumps unevenly and movement no longer looks smooth. The second is motion sickness, where the sensory conflict causes nausea, dizziness, or sweating. The third is a felt delay, where the user clearly notices a gap between turning their head and the screen catching up, creating a sense that the picture cannot keep pace with their movement. For an app like Cardia, which is meant to help users relax and lower body arousal signals such as heart rate and skin response, any of these problems can cancel out or even reverse the intended effect. Instead of settling into a calm,

low-arousal state, the user might be pushed into a stress response by stutter alone. This is the main reason behind the performance work covered in this chapter.

To understand how to improve the scene in a careful, step-by-step way, it helps to first understand the basic structure of the rendering pipeline. The rendering pipeline is the series of steps that the GPU and the engine carry out, in order, to turn scene data into the final image on screen. The output of each step feeds into the next, flowing through the same fixed order every single frame. It mainly includes the stages described below.

The first stage is the geometry stage, or vertex processing. The mesh vertex data for every object in the scene first passes through the vertex shader, which transforms each point from model space to world space, then to camera space, and finally to screen space. If a material includes a World Position Offset, such as leaves swaying in the wind, this is also where each vertex gets shifted in real time. Next comes rasterization, where the transformed triangles are projected onto the screen's pixel grid to work out which pixels each triangle covers. After that comes the pixel, or fragment, stage, where every covered pixel runs through the pixel shader once, sampling textures and calculating lighting. Most of the nodes in a material blueprint, such as hue shifts or desaturation, run pixel by pixel at this stage. Since a screen usually has far more pixels than a mesh has vertices, this stage is where rendering cost adds up the fastest, and VR doubles this cost because of stereo rendering. Lighting and shadows come next, where the engine decides, based on the light type, whether to calculate lighting fresh every frame, called dynamic lighting, or simply read from a lightmap that was baked in advance, called static lighting. Then comes post-processing, where screen-wide effects such as bloom, depth of field, and tone mapping are layered on top of the finished image. Finally, the result is written to the frame buffer and sent to the display.

Within this pipeline, there are two main ways to calculate lighting: forward rendering and deferred rendering. Forward rendering works out the lighting from every light source for each object during the pixel shading stage, one object at a time. Deferred rendering instead first writes geometry and material data into several G-Buffer textures, then calculates the lighting all at once in screen space. Deferred rendering suits scenes with many dynamic lights better, but it costs more memory and bandwidth. Most VR projects use forward rendering instead, because it works better with multi-sample anti-aliasing, handles see-through materials like leaves and hair more naturally, and is easier to keep within a rendering budget. The trade-off is that forward rendering cannot use real-time global lighting systems like Lumen, which depend on the G-Buffer, so it must rely on baked lighting to get indirect light. This is why forward rendering is treated as one factor affecting frame rate, and why the light-baking

process is explained separately later on. The UE5 rendering pipeline is fairly complex, and the things that affect the final frame rate can be grouped into a few main areas.

The first area is geometry and draw calls. The total number of triangles the scene must process in one frame, known as the triangle count, matters a lot: more triangles mean more work for both vertex processing and rasterization. The number of draw calls also matters, since each separate draw call creates extra communication overhead between the CPU and GPU, and too many small draw calls, especially from many separate static meshes, can slow the render thread noticeably. How much instancing is used also matters: if the same mesh can be combined into one draw call through instancing, such as UE5's Instanced Static Mesh or Hierarchical Instanced Static Mesh, the number of draw calls can drop sharply.

The second area is shaders and materials. Shader complexity matters: the more nodes a material blueprint has, the more textures it samples, and the more math it runs, especially trig functions or noise functions, the more expensive it is to shade each pixel. World Position Offset, used for things like wind-driven swaying, shifts geometry in real time during the vertex shader stage. It does not add more triangles, but it does add more work at the vertex stage, and it often breaks some rendering shortcuts, such as certain acceleration structures used by Nanite. Two-sided rendering and transparent materials also add cost, since two-sided lighting fixes and alpha blending greatly increase overdraw, meaning the same pixel gets drawn more than once.

The third area is lighting and shadows. Dynamic lighting and real-time shadows cost much more than static, pre-baked lighting, since they must be recalculated every single frame. Forward rendering also brings a limit: while VR projects generally choose it for performance and anti-aliasing reasons, it cannot use dynamic global lighting systems like Lumen, so it must rely on baked lightmaps for indirect lighting, which adds extra requirements for setting up static lights and the baking workflow.

The fourth area is effects and particle systems. With the Niagara particle system, the number of particles, how often the emitter updates, whether GPU simulation is turned on, and whether collision or lighting interaction is enabled, all directly affect the load on both the GPU and the CPU. Stacking several post-processing effects, such as depth of field, bloom, and screen-space reflections, also adds cost, and this cost doubles in VR because of stereo rendering.

The fifth area is logic and CPU work. Tick events that run every frame without limits, or heavy loops inside blueprint logic, take up time on the Game Thread. Communication with outside data sources also matters. For example, in the Cardia project, receiving body signal

data from the BITalino device over OSC or a serial port can add extra CPU load if it is read too often or in an inefficient way. Physics and collision calculations add cost too: the more objects in the scene that can collide, the more work the physics engine must do checking and solving these collisions every frame.

There are two common patterns when looking for where a bottleneck sits. When the GPU time is too high, the usual causes are the Niagara particle system, stacked post-processing effects, Lumen dynamic lighting, real-time shadows, or volumetric clouds and fog, and the typical sign is that the stutter gets worse as resolution goes up, with large red or pink areas in the Shader Complexity view. When the Game Thread is too heavy, the usual causes are Tick events with no limit, frame-by-frame loops in blueprints, frequent parsing of UDP or serial port data, or too many overlap collision callbacks, and the typical sign is that the stutter stays even after lowering visual quality, while CPU usage stays high.

It is worth being clear that frame rate is a result of how the system is performing, not the cause itself. In other words, the goal of optimization work is not to push the frame rate number up directly. It is to lower the actual time spent in each stage of the rendering pipeline, whether that is CPU logic, the render thread, or GPU shading, and a higher frame rate naturally follows from that. UE5 offers console commands such as stat fps, stat unit, and stat rhi, which show the frame rate number, how time is split across the Game, Draw, and GPU stages, and finer details like memory use, draw call count, and triangles drawn. These are the core tools for finding where a bottleneck sits. The Shader Complexity and Quads view also helps, since it shows the pixel shading cost of each part of the scene using color, making it easy to see visually where the cost is coming from. The optimization work described in the rest of this chapter is based on reading the output of these tools.

When the Cardia scene was tested in the headset, clear frame rate swings and stutter showed up right from the start, especially in areas with a lot of plant life in view. Using the Shader Complexity and Quads view in the UE5 editor together with stat rhi data, the check found three main sources of performance cost: the Niagara particle system in the scene, the wind material used for the plants, and the rendering and collision load from the large number of plant and rock objects themselves.

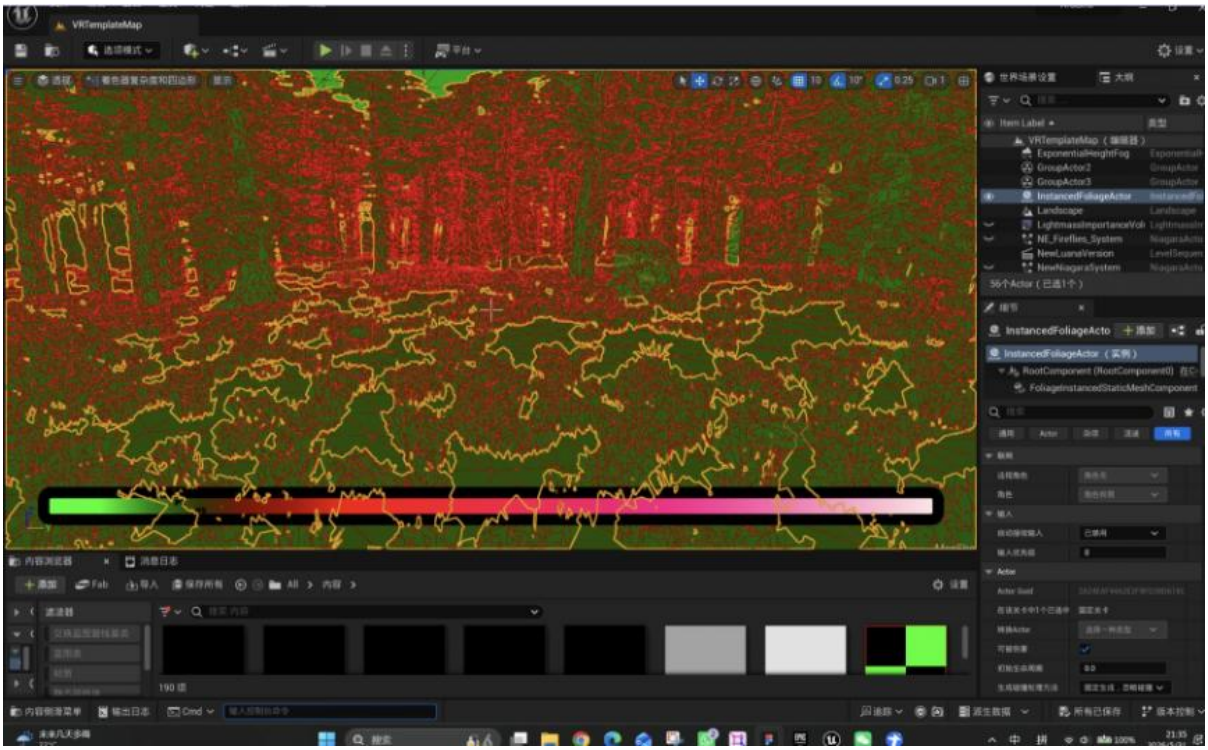

Figure 3.25 — UE5 Shader Complexity and Quads View of the Cardia Scene Prior to Optimization, Highlighting High-Cost Plant Regions in Red and Pink

As the picture shows, in the Shader Complexity view, the areas of the scene filled with trees and bushes turn mostly red or even pink-white, where the color bar runs from green through red to pink-white to show shading cost going from low to high, while large, low-poly areas like the terrain stay green, meaning low cost. This shows that the main cause of stutter is not big geometry such as the terrain or buildings. It comes from the huge number of plant instances, each carrying a fairly heavy shading cost on its own.

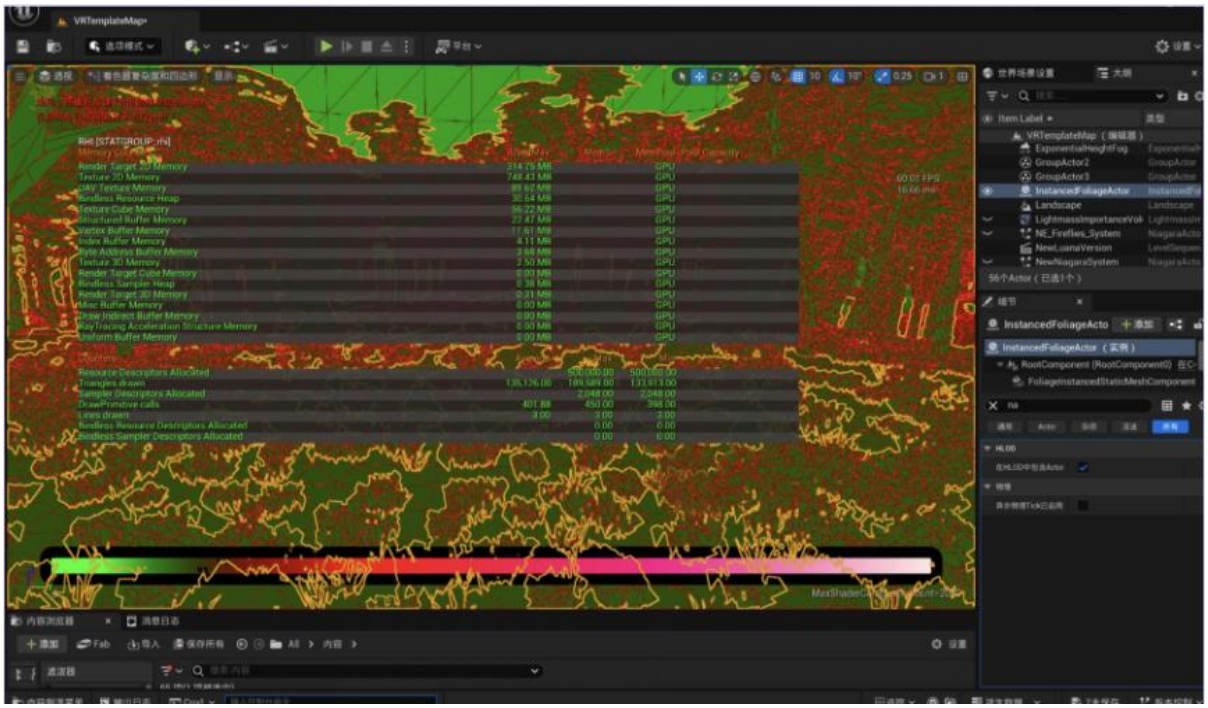

Figure 3.26 — UE5 stat rhi Output Displaying Triangle Count, Draw Call Count, and Texture Memory Usage Before Optimization

The data from the stat rhi command gives more exact numbers. Triangles drawn peaked at around 189,000, Draw Primitive calls peaked at about 450, and 2D texture memory reached as high as 748.43 MB, making it the single largest user of graphics memory. Together, these numbers point to the same conclusion: the number of plant-related meshes, the amount of texture data, and the cost of shading them are the main parts of the current frame rate bottleneck.

The swaying of leaves in the wind is not made with skeletal animation or physics simulation. Instead, it comes from a vertex-level shift driven by the World Position Offset node inside the material blueprint, built around a custom material function called Simple Grass Wind. This function takes three inputs: Wind Intensity, which sets how strong the wind is. Wind Weight, which sets how much a leaf responds to the wind, usually set through vertex color or a texture mask so that leaf tips and stems sway by different amounts. And Wind Speed, which sets how fast the swaying repeats over time. It uses the engine's global time as input, runs it through a sine or noise function to get a movement vector that changes in a repeating pattern, and sends that vector to the WorldPositionOffset pin of the MakeMaterialAttributes node. Because this calculation runs in the vertex shader stage, every single vertex on every leaf mesh in the scene has to be recalculated every frame. This means the more natural the swaying looks, which usually means more vertices and a more complex wind function, the heavier the load on the vertex stage becomes.

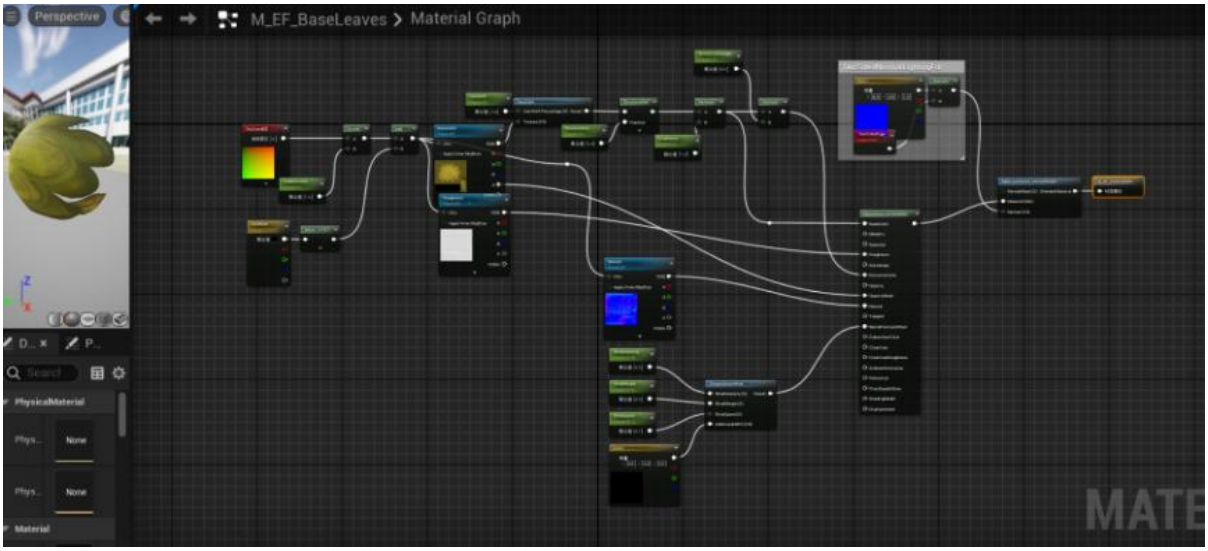


Figure 3.27 — Material Graph of M_EF_BaseLeaves Showing Wind Animation and Two-Sided Lighting Nodes

At the same time, this material adds several more calculations during the pixel shader stage. A hue shift, plus two separate desaturation steps controlling fraction and brightness, work together to give the same leaf texture random changes in color depth and warmth, so a large repeated texture does not look flat and repetitive. A TwoSidedNormalLightingFix node also corrects the normal direction for two-sided leaves: since leaf meshes are usually thin, single-layer shapes rendered as two-sided to save on triangle count, the back face would show clearly wrong lighting if it simply reused the front face's normal, so an extra fix is needed. Every one of these calculations has to run once for every pixel covered by the leaf texture, and since the plants in the Cardia scene are laid out in huge numbers using the Instanced Foliage system, even calculations that seem cheap on their own add up in a straight line, or worse, as the instance count grows. This is what shows up as the large, intense red and pink areas in the Shader Complexity view, and it is the main source of GPU load.

Beyond the shading cost of the leaf material itself, the Niagara particle systems used for atmosphere, such as fireflies and floating light dust, also use a fair amount of GPU and CPU resources. The number of particles, how often the emitter updates, and whether lighting interaction or collision detection is turned on, all directly affect how expensive these effects are. The large number of trees, bushes, rocks, and other ground objects in the scene also add cost, since they generally use fairly detailed meshes along with normal and roughness texture maps. Once these are stacked up through instancing, both memory use, as seen in the texture memory numbers earlier, and the number of draw calls grow a lot. Together, these three things, the leaf material, the particle systems, and the large number of plant instances, are the root cause of the clear stutter felt the moment the headset enters the scene.

To address these problems, and keeping in mind Cardia's design goal of low interaction cost and stable immersion, a step-by-step optimization process was put together. It covers four areas: trade-offs in interaction design, texture optimization, model and collision optimization, and light baking.

The first step, a trade-off in interaction design, was to remove collision from the scene.

Before working on the technical side of rendering, the team first took something away at the level of interaction design. Cardia's core interaction goal is to let users enter and stay in a calm, meditative state with as little physical and mental effort as possible. This is different from typical VR content built around moving, exploring, and touching objects, such as walking, grabbing, or collision feedback in VR games. Cardia is designed instead as a mostly still, seated or standing meditation experience, where the user mainly interacts with the system through breathing and body signals, not by moving through the scene or physically touching objects in it. Because of this, most of the static meshes in the scene, such as cliffs, tree trunks, and rocks, still carried collision that was originally meant to support movement-based interaction, even though Cardia's actual experience never uses it.

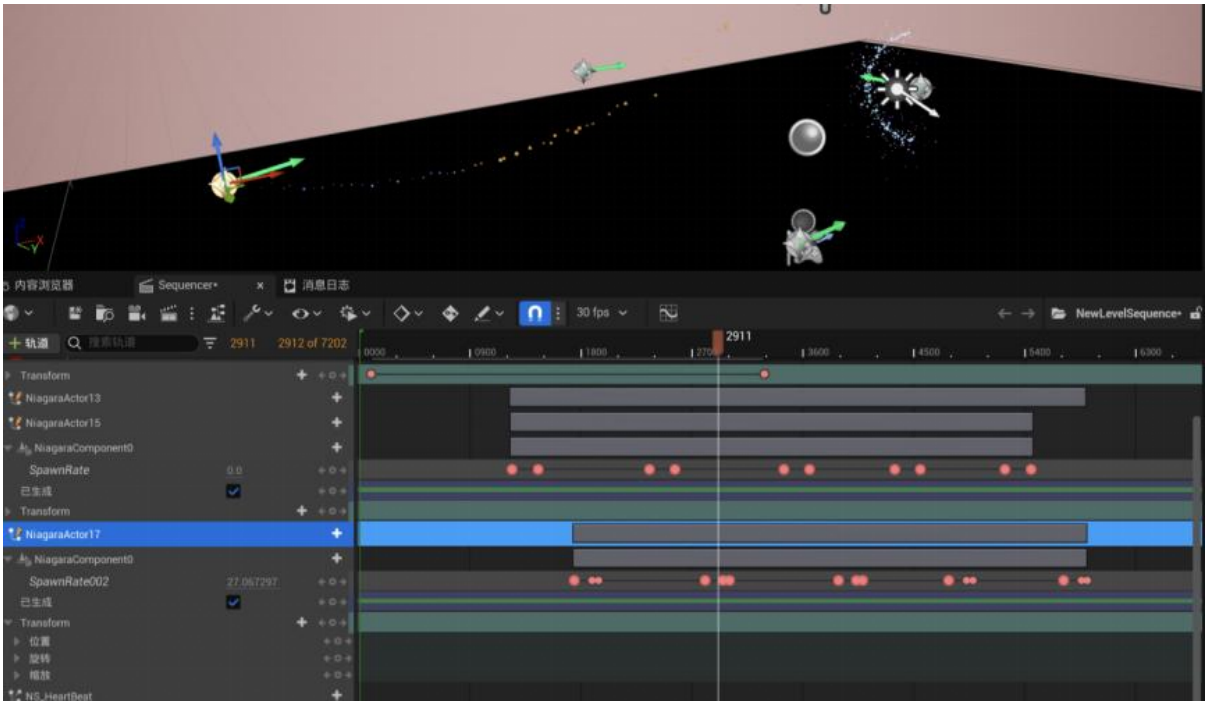


Figure 3.28 — UE5 Sequencer View of Niagara Particle Actors and Collision Wireframes Before Collision Removal

The team selected the terrain and cliff-type Actors in the World Outliner, such as EF_Cliff1, EF_Cliff2, and EF_Cliff3, all at once, and changed their Collision Preset from the default setting, which usually includes full physics and query responses like BlockAll, to

NoCollision. The reason this helps is that UE5's physics engine has to keep a bounding structure, such as a BVH, for every Actor with collision turned on, and check it for overlaps and blocking every single frame. This work happens on the CPU's physics thread, and the more collision objects there are, and the more complex their shapes, the more it costs. Once an object is set to NoCollision, the engine removes it completely from the physics scene's detection queue, making it invisible at the physics level. It still renders normally and still responds to lighting, but it no longer takes part in any ray tracing, overlap events, or collision checks.

Since users in the Cardia scene never walk into or physically touch static objects like cliffs and tree trunks, there is no movement controller that needs to be blocked, and no physics interaction event that ever needs to fire. Keeping the collision on these objects added nothing useful to the experience, while still costing CPU time on the physics thread every frame. Setting them all to NoCollision is an optimization with zero loss to the experience and a real gain in performance. It does not rely on lowering visual quality or simplifying any models. Instead, it starts from the interaction design itself, finds the parts that no longer match the design goal, and removes them. This is the optimization idea of taking away whatever does not serve the experience, and it was the first step taken in this chapter, and the one given the highest priority.

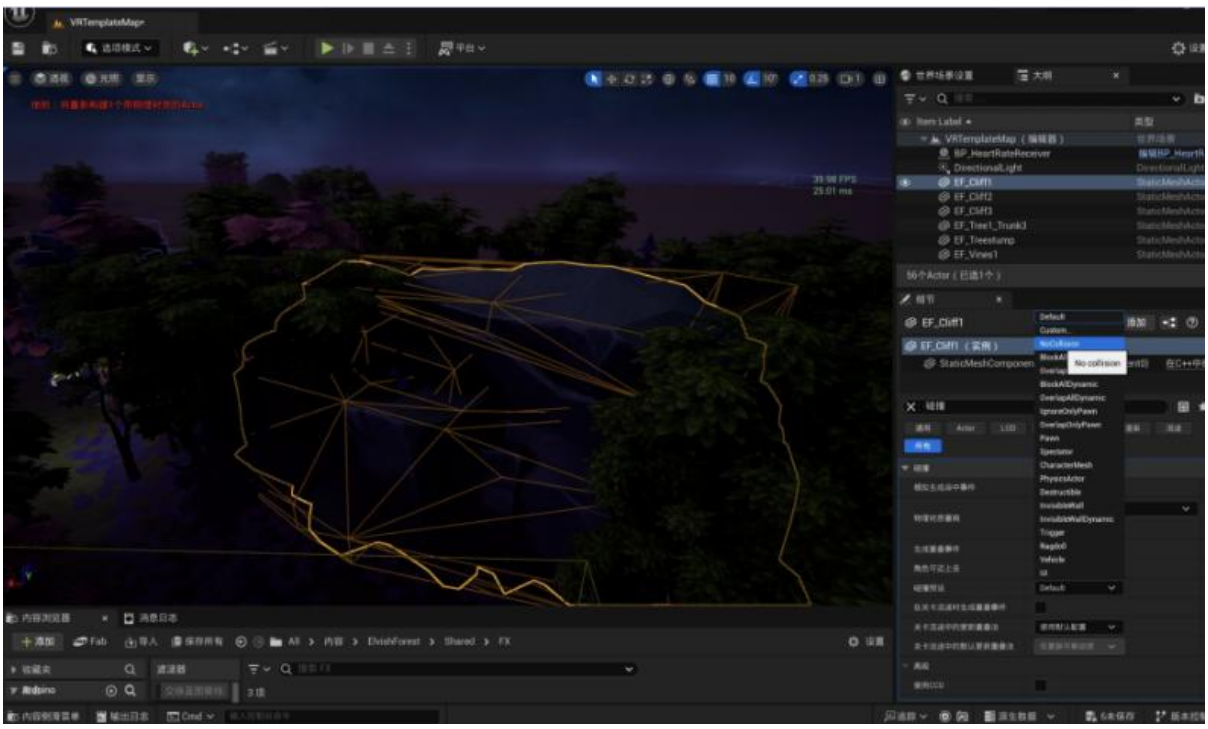

Figure 3.29 — UE5 World Outliner Showing Collision Preset Changed from Default to NoCollision for Terrain-Type Actors

The next step looked at model and collision precision. The team checked the on-screen triangle count and draw call count using the stat rhi command, aiming to keep triangles under 5 million and draw calls under 1,000, and used the optimization view to find models with too many triangles. For plant and rock models that went over budget, the team turned on the Nanite virtual geometry system and set a custom triangle reduction percentage. After confirming that the reduced version looked almost the same as the original, the reduced model was exported and put back in the same place in the scene, and Nanite was then turned off again to avoid conflicts with the forward rendering pipeline's shadow and lighting steps. One thing worth noting is that the collision export option should be turned off when exporting these models. This follows the same logic as removing collision earlier: since most static objects in the scene no longer need to take part in physics checks, there is no reason to keep generating extra collision data at the model level either.

After finishing this round of interaction-level trimming, led by removing collision, along with the model and texture work that followed, the overall color in the Shader Complexity view shifted clearly from red and pink toward green. The stat rhi data showed the frame rate rising from about 40 FPS before optimization, around 25 milliseconds per frame, to close to the device's target of 60 FPS, around 16.7 milliseconds per frame. It is worth noting that the density and material complexity of the plants themselves were not heavily cut back in this round. The current work mainly removes performance loss that came from unneeded physics checks and texture memory pressure. Going forward, while still keeping the visual density of plants, which matters for how natural and immersive the meditation scene feels, the next stage could further trim the shader instruction count for plant materials and the particle density of the Niagara effects.

This chapter looked at the stutter problem found during real headset testing of the Cardia VR meditation system. Starting from the basic link between frame rate and VR experience quality, it laid out the main things that affect frame rate in a UE5 project. Using the Shader Complexity view together with stat rhi data, it found that the main performance bottleneck in the current scene comes from the per-pixel and per-vertex cost of the leaf wind material, combined with the load from the Niagara particle system and the large number of plant instances. Based on this, a step-by-step optimization process was put together, covering trade-offs in interaction design, texture optimization, and model and collision optimization. Among these steps, removing the collision that no longer matched Cardia's low-interaction-cost design goal shows the project's underlying idea clearly: performance work should serve the experience, not just serve technical numbers. Optimization should start by asking what the

interaction design actually needs, and only then move deeper into the technical details of the rendering pipeline.

# 4. Olfactory Experiment

## 4.1. Olfactory System Design and Implementation

Building on the established positive effects of multi-sensory immersive experiences on psychological stress relief, this study extends the existing Cardia system prototype by introducing an olfactory feedback module and integrating a physiological signal acquisition system, through which the original interaction logic has been optimized and restructured. To more rigorously examine how different interaction mechanisms influence users' psychological relaxation outcomes, this research employs a comparative study design, systematically investigating the effect of olfactory stimulation on the efficacy of VR-based meditation. In parallel, to minimize the potential confounding influence of extraneous variables on experimental outcomes, certain haptic interaction components present in the earlier prototype were deliberately removed, thereby ensuring greater comparability across experimental conditions and strengthening the overall validity of the results.

### *4.1.1. Rationale for Olfactory Augmentation*

Following the development and preliminary testing of the first Cardia prototype, the research team conducted a systematic reflection on the overall interaction logic of the system. While the initial version successfully demonstrated the technical feasibility of combining a VR environment with biofeedback mechanisms, notable opportunities for improvement remained across three dimensions: the richness of sensory channels, the naturalness of the interaction flow, and the intuitiveness of physiological data visualization. Guided by these reflections, the second iteration of Cardia introduced upgrades along three core axes: the addition of an olfactory channel, the replacement of the biometric acquisition device coupled with a reduction in interaction burden, and a comprehensive restructuring of the biometric data visualization modality.

At the level of interaction design, the system preserved and deepened the original dynamic visual feedback logic. The ring-shaped visual elements within the scene continuously adjust in response to the user's real-time physiological state: when the user's stress level is elevated, the motion velocity of the visual elements increases accordingly. As breathing rhythm becomes more measured and HRV status improves, the overall pace of scene movement slows in parallel, forming an immersive feedback relationship that remains synchronized with the user's internal physiological condition. Furthermore, the system

provides implicit breathing rhythm guidance through dynamic visual forms such as particle flows and breath-guiding trajectories, enabling users to naturally regulate their breathing frequency without relying on explicit textual prompts, thereby reducing the sense of cognitive intervention.

In terms of technical selection, BITalino Belt was ultimately chosen as the core physiological data acquisition device. Compared to certain consumer-grade wearable solutions evaluated during earlier testing phases, BITalino demonstrated superior performance in real-time data transmission stability and signal quality, making it particularly well-suited to the high-frequency continuous physiological data acquisition demands of a VR experimental environment. Concurrently, the research team removed the physical wristband interaction module present in the first prototype, with the aim of reducing the additional operational burden imposed on users during the wearing process and rendering the overall interaction flow more lightweight and natural.

The introduction of the olfactory channel represents the most exploratory design decision of this upgrade cycle. Within the theoretical framework of multi-sensory immersive experience, olfaction has long been one of the sensory dimensions least systematically incorporated into VR experience design. Nevertheless, a substantial body of research in psychology and neuroscience has established a direct neurophysiological connection between olfactory perception and emotional regulation. Olfactory signals are projected directly through the olfactory bulb to the amygdala and hippocampus — two brain regions that play central roles in emotional processing and memory encoding — enabling scent to rapidly trigger shifts in emotional state at a subconscious level, with a response speed that can surpass even that of visual and auditory channels. Among the wide range of aromatic compounds studied, lavender (*Lavandula angustifolia*) has had its calming effects substantiated across multiple experimental investigations. Its primary active constituents, linalyl acetate and linalool, are understood to suppress excessive activation of the central nervous system, reduce cortisol levels, and promote parasympathetic nervous system activity, thereby producing pronounced relaxation and anxiety-relieving outcomes (Koulivand et al., 2013). In both laboratory and clinical contexts, lavender scent has been shown to effectively lower participants' heart rate and systolic blood pressure, while also significantly reducing self-reported anxiety and tension (Goel et al., 2005). A number of studies have additionally highlighted lavender's positive effects on sleep quality improvement and preoperative anxiety intervention, further supporting its potential as a non-invasive emotional regulation tool.

Drawing on this body of evidence, the second version of Cardia introduced a lavender olfactory module into the VR scene, delivering a controlled concentration of scent near the user's headset through a compact diffusion device, thereby establishing a multi-sensory intervention experience that operates in coordination with the visual and auditory channels. The incorporation of olfactory stimulation is designed to test a central hypothesis: whether the simultaneous introduction of a natural scent within a VR meditation experience can produce measurable benefits — at both the physiological level (e.g., HRV, heart rate) and the subjective level (e.g., perceived stress, emotional valence) — that exceed those achievable through audio-visual experience alone.

Regarding biometric data visualization, this upgrade reconstructed the original Niagara particle system-based visual representation into a continuous wave-form material ring whose semantic qualities more closely approximate those of an electrocardiographic waveform. This adjustment was intended to enhance the perceptual clarity of the mapping between visual feedback and physiological intuition — allowing users to more directly establish a perceived connection between the undulating graphics on screen and their own cardiac state, rather than relying on the metaphorical abstraction of particle-form representations.

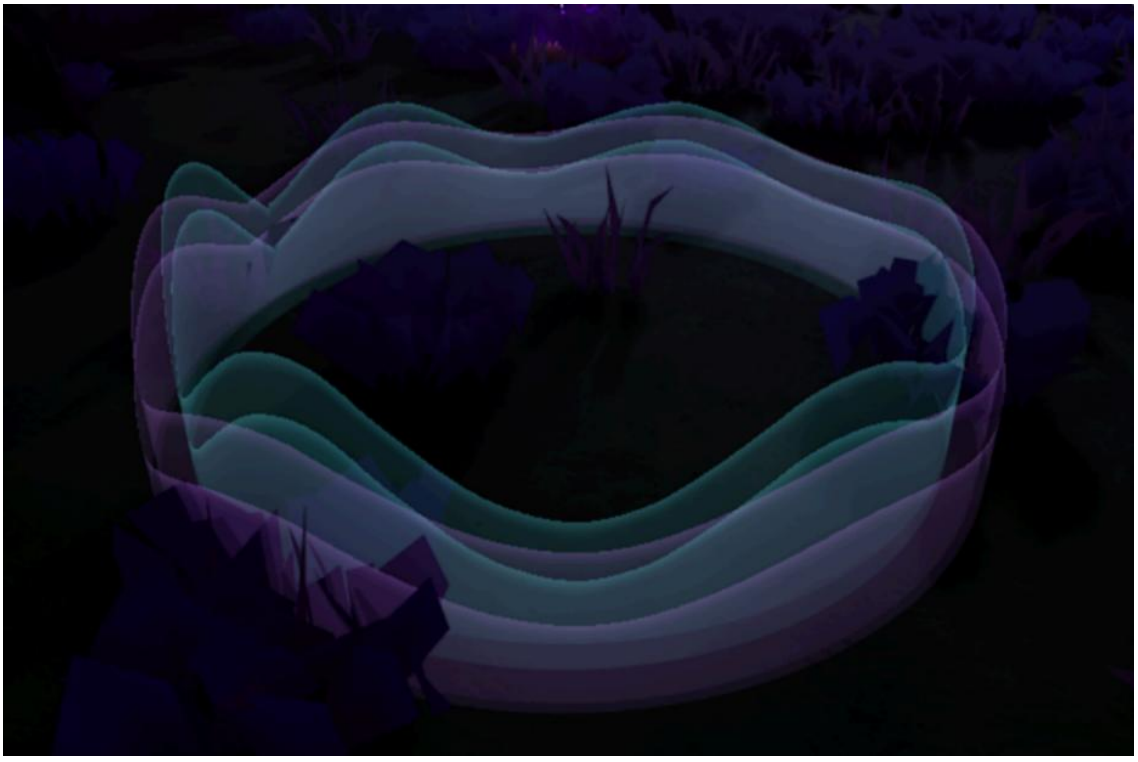

Figure 4.1. Updated Biometric Data Visualization Interface

### *4.1.2. Physiological Data Pipeline: BITalino to UE5*

The biometric data transmission pipeline of this system consists of four sequentially connected stages: the BITalino Assembled Pulse Sensor (PPG) captures raw photoplethysmographic signals. OpenSignal software handles device configuration and data stream initialization. A Python script performs real-time signal processing and HRV metric computation. And finally, data packets are transmitted via the OSC protocol to the Unreal Engine 5 Blueprint system, driving the biofeedback visualization within the scene.

*BITalino Assembled Pulse Sensor*

BITalino Assembled is a modular physiological signal acquisition platform designed for biomedical research and rapid prototyping, developed by PLUX Wireless Biosignals in Portugal. This study employs its PPG (photoplethysmography) sensor module, which detects changes in blood flow volume at the skin surface through a photosensitive element, enabling continuous non-invasive recording of pulse waveforms. The device supports Bluetooth wireless transmission and offers a maximum sampling rate of 1,000 Hz, making it well-suited for high-frequency, low-interference real-time physiological data acquisition within VR experimental environments. Compared to consumer-grade wearable devices, BITalino offers clear advantages in signal stability and data accessibility, as it is capable of outputting raw analog signals for downstream algorithmic processing, establishing it as the core hardware foundation of this study's biofeedback pipeline.

*OpenSignal Connection Configuration*

OpenSignal is the official visualization configuration and data monitoring software provided by BITalino. In this study, after the device was paired with the experimental computer via Bluetooth, its MAC address (**98:D3:A2:FD:8B:F9**) was automatically recognized and displayed as connected within OpenSignal. The sampling rate was configured at 100 Hz — a frequency that provides sufficient temporal resolution for computing HRV-related metrics such as RMSSD and pNN50, while avoiding the processing delays associated with data redundancy. In the channel configuration, **Channel A2**was enabled and assigned to the PPG sensor, with all remaining channels kept inactive. Under the integration settings, "Lab Streaming Layer" (LSL) was enabled, allowing the subsequent Python script to subscribe to the real-time data stream via the LSL protocol and achieve seamless cross-software data interfacing.

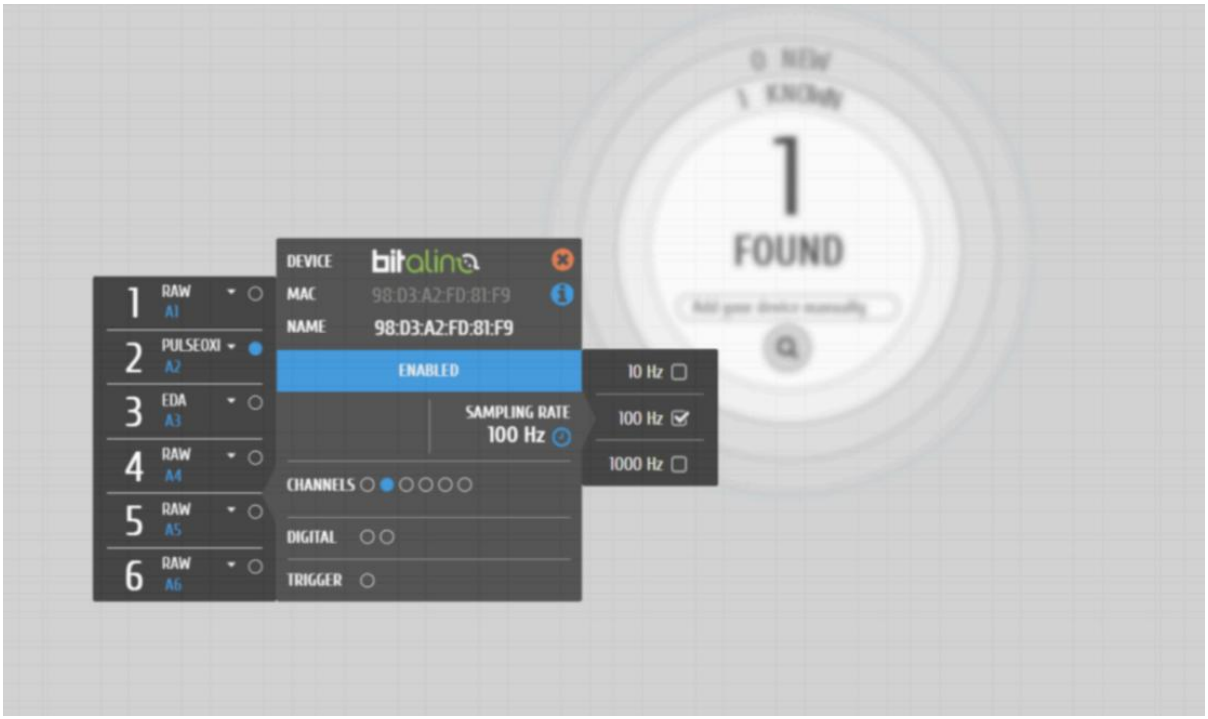


Figure 4.2. BITalino OpenSignal Device Configuration and Channel Setup

*LSL Open Listening*

Lab Streaming Layer (LSL) is a lightweight network protocol specifically designed for real-time laboratory data transmission. Once LSL is activated, OpenSignal broadcasts the data stream over the local network, and the Python script subscribes to this stream through the StreamInlet interface of the pylsl library, achieving immediate data acquisition from the device without requiring any additional port mapping configuration.

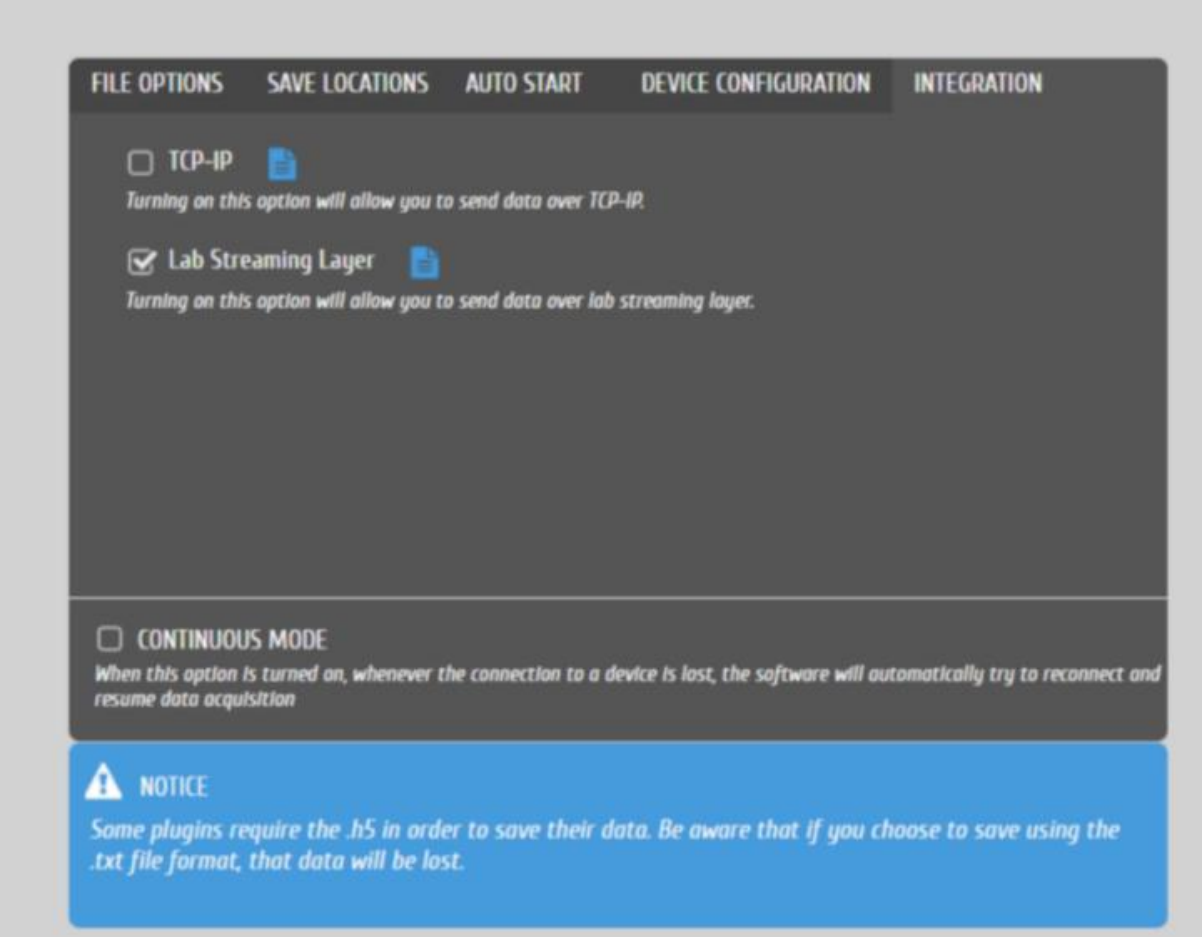


Figure 4.3. BITalino Lab Streaming Layer Integration Settings

*Python Script Implementation Logic*

The Python script (see Appendix) is responsible for three core tasks: data reception, signal processing, and OSC forwarding. The script continuously subscribes to the PPG data stream from OpenSignal via pylsl, pulling raw sample values in real time at 100 Hz, while maintaining a 30-second sliding window to buffer the inter-beat interval (IBI) sequence. During the peak detection stage, the script invokes the scipy.Signal.Find_peaks function, identifying cardiac peaks from the PPG waveform according to a configured minimum inter-peak distance and amplitude threshold, and subsequently computing the RR interval sequence between adjacent peaks. Based on this sequence, the script computes RMSSD (the root mean square of successive differences between adjacent RR intervals) in real time as the core quantitative HRV metric, and encapsulates the instantaneous heart rate (BPM) together with the RMSSD value into an OSC message, forwarding it to local port **5005** to drive real-time scene responses in UE5. In addition, the script saves the complete session data in CSV format to a local research data directory for subsequent offline analysis. Prior to each experimental

session, a 60-second data discard period and a 300-second valid baseline acquisition window are configured to ensure that physiological data is recorded under stable conditions.

*OSC Protocol and UE5 Plugin*

OSC (Open Sound Control) is a lightweight messaging protocol based on UDP network transmission, originally developed for real-time control in music and multimedia applications, and subsequently adopted widely in interactive system development due to its low latency and cross-platform compatibility. In Unreal Engine 5, OSC functionality is implemented through the official built-in "OSC Plugin". Once enabled, the engine can directly invoke nodes for OSC server creation, listening, and message parsing within the Blueprint system, without requiring any additional low-level network code.

Communication between the Python script and UE5 is routed through the local loopback address (**127.0.0.1**), ensuring low-latency and stable data transmission suited to real-time biofeedback-driven scene control.

*UE5 Blueprint Logic*

The data reception logic on the UE5 side is encapsulated within the Actor Blueprint BP_HeartRateReceiver, whose Event Graph presents the complete control chain from OSC message reception to scene parameter updates.

When the Event BeginPlay node is triggered, the Blueprint executes its initialization sequence: it first calls the Create OSC Server node, setting the reception address to the local loopback IP (**127.0.0.1**), assigning the listening port to **5005**, and enabling the Multicast Loopback option to ensure that locally originated messages are correctly routed. Once the server instance is created, it is assigned to the Blueprint variable OSCServer via the SET node, completing the reference binding. The Blueprint then calls the Bind Event to On OSC Message Received node, registering the custom event OnDataReceived as the callback handler for incoming OSC messages, thereby establishing an event-driven asynchronous listening mechanism.

When the Python script dispatches an OSC message, the OnDataReceived custom event is triggered, with its input parameters carrying the complete OSC message object (Message), the source address (IPAddress), and the port number (Port). The Blueprint immediately calls the Get OSC Message Float At Index node, extracting the heart rate float value from the message body at index **0** and storing it in the variable NewVal_0. Once extracted, this value is converted to string format via the Float to String node and output in real time to the viewport via the Print to Screen node for data verification during the debugging phase. Finally,

NewVal_0 is passed as a driving parameter to the heart rate visualization system (the material ring waveform), controlling its dynamic deformation and animation rhythm, thereby achieving real-time mapping from physiological data to visual feedback.

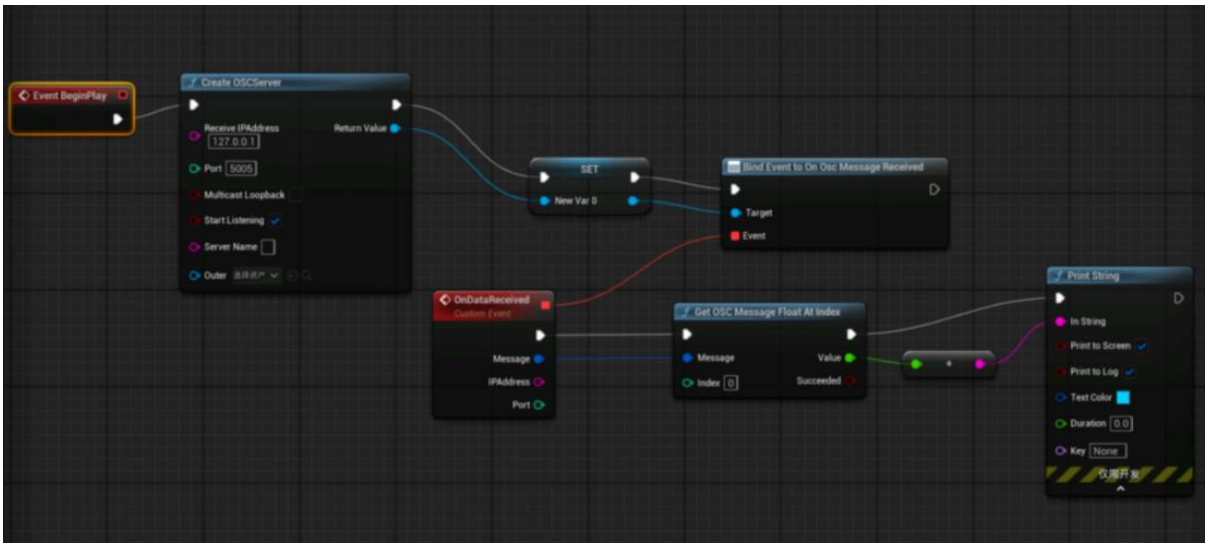


Figure 4.4. HeartRateReceiver Blueprint Event Graph: OSC Message Reception and Parameter Routing

Following the configuration and debugging of each stage described above, the raw pulse waveform data acquired by the BITalino PPG sensor is capable of being stably transmitted into the UE5 Blueprint system via the complete pipeline of OpenSignal → LSL → Python → OSC, with an end-to-end latency of approximately 100 milliseconds. The heart rate values printed in real time to the viewport correspond closely to the values measured directly by the sensor, confirming that the entire biometric data pipeline possesses reliable real-time transmission capability under VR experimental operating conditions, and providing a stable technical foundation for the biofeedback-driven immersive experience in subsequent phases of the study.

### *4.1.3. HRV-Driven Particle Visualization*

This section describes how the Cardia system integrates real-time heart rate data with the wave-form material ring visual elements within the scene, constructing a dynamic biofeedback visualization layer that carries the semantic qualities of an electrocardiographic waveform. The overall implementation pathway consists of three hierarchical levels: the material node graph defines the visual generation logic of the waveform. The Niagara particle system serves as the hosting and rendering container for the material. And the Blueprint

control logic is responsible for mapping the physiological data received via OSC in real time onto the material driving parameters.

*M_Flowing Material Node Graph*

The material M_Flowing constitutes the visual core of the entire heart rate visualization system. Its material graph generates a continuously flowing ripple morphology through the layered superposition and modulation of multiple noise functions. The overall construction principle of the material graph is as follows: a time-driven noise field serves as the foundational source of deformation, with its amplitude, frequency, and phase offset controlled parametrically, so that the resulting waveform contour can respond in real time to changes in external input values.

In terms of node composition, the material graph first introduces the Time node as the animation driving source, multiplying it by a scaling coefficient before feeding it into the Noise node to generate a continuously evolving noise field over time. After this noise field undergoes numerical remapping (Remap) processing, it governs the displacement offset of the material surface, thereby producing a wave-like undulation deformation effect across the mesh surface. With respect to color output, the material maps noise intensity values onto a custom gradient curve, forming a smooth color transition from deep blue-violet to cyan-green, consistent with the overall visual style of the Cardia system.

The material exposes several Scalar Parameter interfaces, corresponding respectively to waveform height, intensity, and time multiplier, in order to enable per-particle and per-frame parameter writing within the Niagara system through the Dynamic Material Parameters channel. The material's blend mode is set to Translucent, with two-sided rendering enabled, ensuring that the waveform ring presents a complete and visually coherent appearance from any viewing angle within the VR scene. In addition, the material employs the World Position Offset node to drive vertex displacement, causing the waveform undulations to act directly upon the geometric form rather than remaining confined to the surface texture layer, thereby further enhancing the three-dimensional sense of depth and the perceptual authenticity of the rhythmic motion.

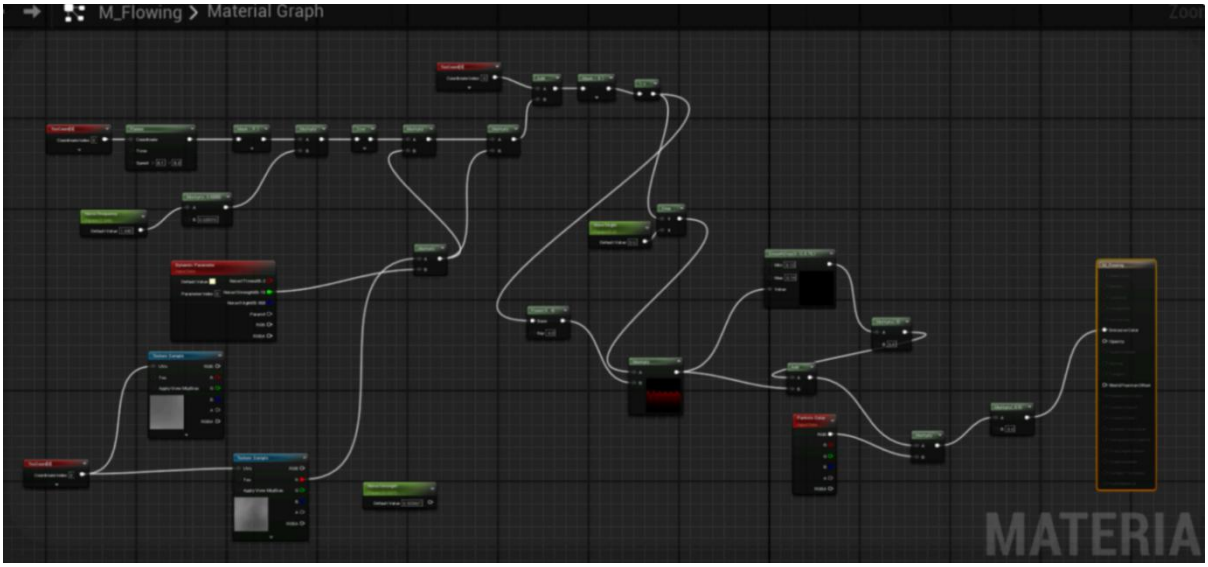


Figure 4.5. Material Node Graph of M_Flowing

*Material Integration within the Niagara System*

The Niagara particle system NS_HeartBeat serves as the rendering container for the material, taking responsibility for instantiating the material, layering multiple instances, and interfacing with the biometric data parameter channels. The system internally contains two primary emitters — Noise1 and Noise1001 — which share an identical module stack structure and correspond respectively to two groups of waveform rings with slightly differentiated visual characteristics within the scene. Through minor parameter offsets between the two emitters, a layered parallax effect is achieved, enhancing the sense of spatial depth in the visual output.

In the render module configuration, both emitters employ the Mesh Renderer rather than the default Sprite Renderer. The target render mesh is specified as a custom waveform ring mesh (Wave mesh), and the M_Flowing material instance is bound within the Materials slot. Once Enable Material Override is activated, the Niagara renderer replaces the mesh's default material with the M_Flowing material instance, simultaneously granting the system dynamic write access to that material's parameters, allowing the values set by the subsequent Dynamic Material Parameters module to be updated into the corresponding Scalar Parameter slots within the material on a per-frame basis.

The critical parameter passing mechanism is implemented through the Dynamic Material Parameters module. This module operates during Niagara's Particle Update stage, writing the current values of user parameters into the material's Dynamic Parameter nodes channel by channel and frame by frame, achieving real-time driving from external numerical values in the particle system to internal visual parameters within the material without requiring the

Blueprint to manually invoke material parameter update interfaces on every frame. During the Particle Update stage, the module writes Niagara User Parameter values into the corresponding material parameter slots by index: **Index 0** passes into Param0 (corresponding to the time multiplier control within the material). **Index 1** passes into Param1 (corresponding to noise intensity). And **Index 2** passes into Param2 (corresponding to waveform height). This mechanism ensures that external Blueprint writes to Niagara user parameters ultimately take effect on the material's visual expression, completing the full mapping chain from physiological values to visual form.

*Exposed Parameter Descriptions*

Within the User Parameters panel of the Niagara system, three core user parameters are exposed, each controlling a distinct visual dimension of the waveform visualization:

**User_Noise1Hight0_5** (default value: 0.5) controls the overall height of the waveform that is, the displacement amplitude of the World Position Offset within the material. The higher this parameter value, the greater the undulation amplitude of the waveform ring, producing a more pronounced rhythmic visual sensation. As the value decreases, the waveform flattens, conveying a quieter and more stable physiological state. Within the experimental logic, this parameter is mapped with a positive correlation to the user's real-time heart rate value — when heart rate is elevated, the parameter value rises and waveform undulation increases. As the user gradually relaxes, waveform height decreases, forming an intuitive visual feedback loop.

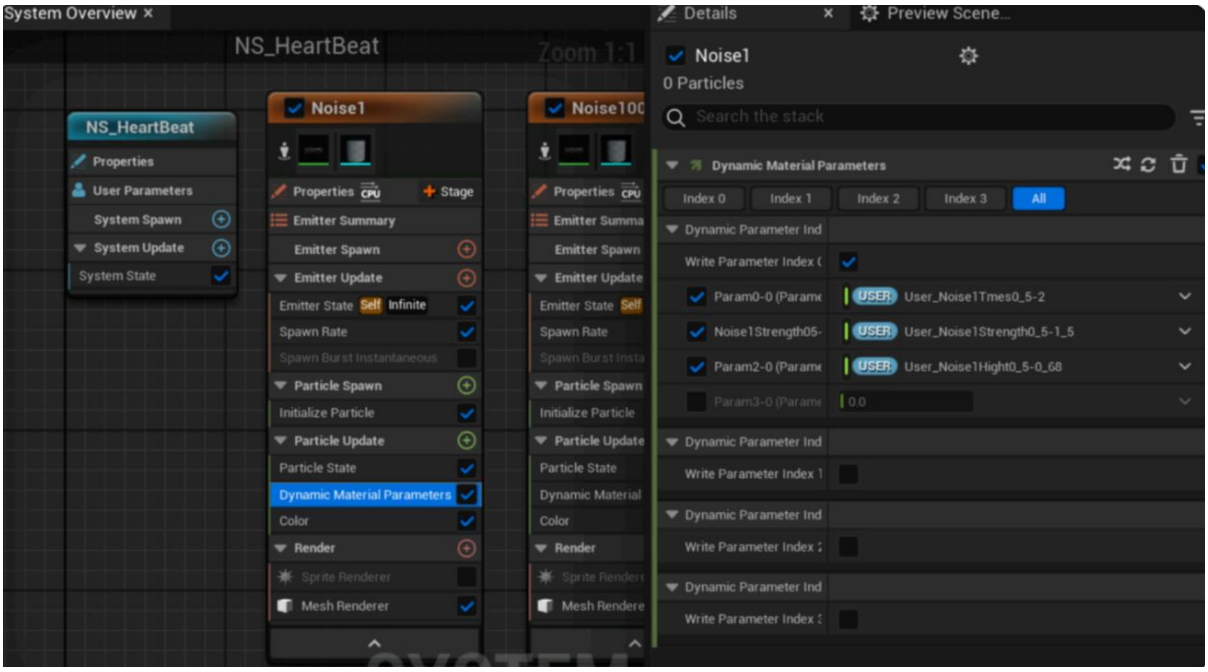

Figure 4.6. Niagara System Configuration: Emitter Stack and Dynamic Material Parameter Settings

**User_Noise1Strength0_5** (default value: 0.8) controls the intensity of the noise field that is, the richness of surface detail and the amplitude of random perturbation across the waveform surface. When intensity is high, the waveform edges exhibit a greater degree of local deformation and textural detail, appearing more visually active. When intensity decreases, the waveform contour becomes smoother and the overall appearance more tranquil. Within the system, this parameter functions as an auxiliary mapping channel for HRV state when the RMSSD value is low (indicating higher stress), noise intensity is modestly elevated to reinforce the visual sense of tension. As RMSSD improves, intensity correspondingly subsides.

**User_Noise1Tmes0_5**(default value: 0.55) controls the temporal evolution rate of the noise field — that is, the speed at which the waveform flows. This parameter directly influences the user's perceived sense of "cardiac rhythm pace": when the value is high, the waveform flows more rapidly, evoking visual associations with an accelerated heart rate. When the value decreases, the waveform flow slows, resonating with the rhythm of calm breathing. The three parameters work in concert to collectively construct a dynamic material system capable of expressing changes in the user's physiological state in a continuous and nuanced manner at the visual level.

*Blueprint Data Input and Normalization Processing*

The complete control logic governing biometric data from OSC reception through to final writing into Niagara user parameters is encapsulated within the Event Graph of the BP_HeartRateReceiver Blueprint. Since the raw heart rate values (BPM) output by the BITalino PPG sensor typically fall within a range of 40 to 120, while Niagara user parameters expect normalized floating-point values within the interval of 0 to 1, the Blueprint incorporates a comprehensive numerical mapping and normalization processing pipeline.

When the OnDataReceived event is triggered, the Blueprint first extracts the heart rate float value from the message body at **Index 0** via the Get OSC Message Float At Index node, obtaining the raw BPM value. This value then enters the normalization computation chain: the Blueprint applies the MapRangeClamped node to linearly map the input range (40, 120) onto the output range (0.0, 1.0), with a Clamp applied to ensure boundary safety and prevent parameter overflow caused by anomalous peak values. The normalized result is assigned to the NormalizedHR variable, serving as the baseline data source for all subsequent parameter writes.

During the parameter distribution stage, the Blueprint retrieves the instance reference of the NS_HeartBeat particle system within the scene via the Get Niagara Component node, and then invokes the Set Niagara Variable Float node to write the NormalizedHR value into each of the three user parameters individually. Specifically, **User_Noise1Hight0_5** directly receives the normalized heart rate value. User_Noise1Tmes0_5 is written as the normalized value multiplied by a baseline coefficient of 0.55, ensuring that the range of flow speed variation remains within a perceptually comfortable range. And **User_Noise1Strength0_5** draws upon the normalized RMSSD value (extracted from OSC message **Index 1**, with a mapping range set to 10–80 ms), enabling HRV state to independently govern waveform surface detail intensity.

This dual-channel parameter mapping design heart rate driving waveform height and speed, HRV driving waveform detail intensity, enables the visual feedback to respond simultaneously to the user's instantaneous heart rate changes and to the deeper state of autonomic nervous system regulation, maintaining visual intuitiveness while encoding a richer dimensionality of physiological information. The entire Blueprint logic operates in an event-driven manner, with each arriving OSC message triggering a complete parameter update cycle, achieving near real-time visual responsiveness at a data sampling rate of 100 Hz.

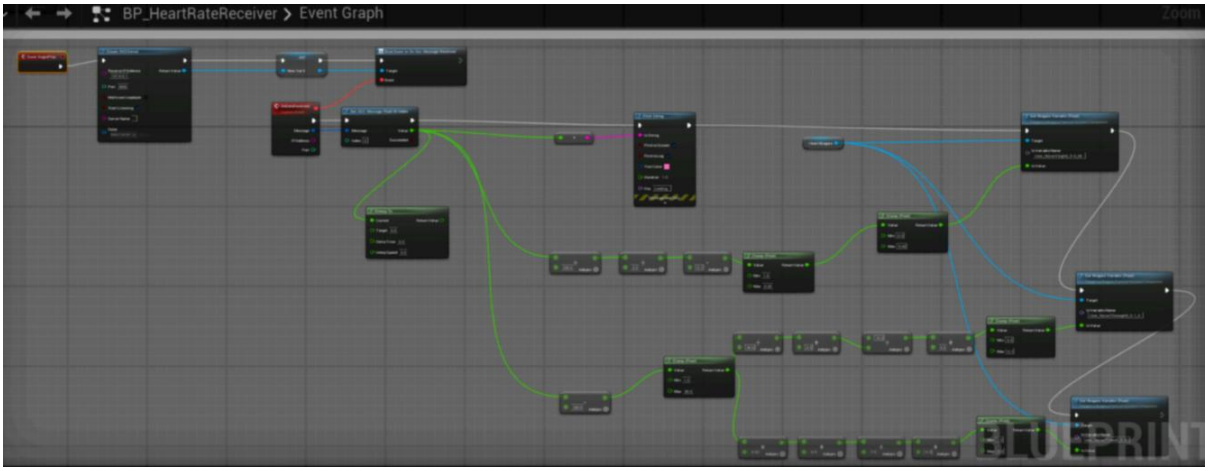


Figure 4.7. BP_HeartRateReceiver Blueprint Event Graph: Full Parameter Normalization and Niagara Write Logic

### *4.1.4. Olfactory Device*

The olfactory stimulation hardware is built around an Arduino Uno R3 microcontroller housed in a custom 3D-printed bracket that clips directly onto the front of the VR headset, positioning the scent-delivery outlet close to the user's nose (Figure 4.8). Mounting the module on the headset itself, rather than as a separate room-based diffuser, keeps the scent delivery distance consistent regardless of the user's head movement during the experience.



Figure 4.8. Olfactory Device Mounted on the VR Headset

The device supports up to three independently addressable scent channels (only one, the lavender channel, is populated in the current build), each driven through a software-implemented pulse-width modulation scheme: an intensity level (10 – 100%, set via serial command) determines how many of ten 50-millisecond steps within each 500-millisecond cycle the output pin is held high, allowing scent intensity to be adjusted in ten discrete steps. Once triggered, a channel remains active for a fixed 4-second period before switching off for an 8-second rest period, repeating this on/off cycle automatically — a design intended to reduce olfactory adaptation (habituation) and keep the lavender cue perceptually noticeable throughout the session, rather than releasing a single, fading burst of scent. The assembled unit, showing the Arduino board and driver wiring within the 3D-printed housing, is shown in Figure 4.9.

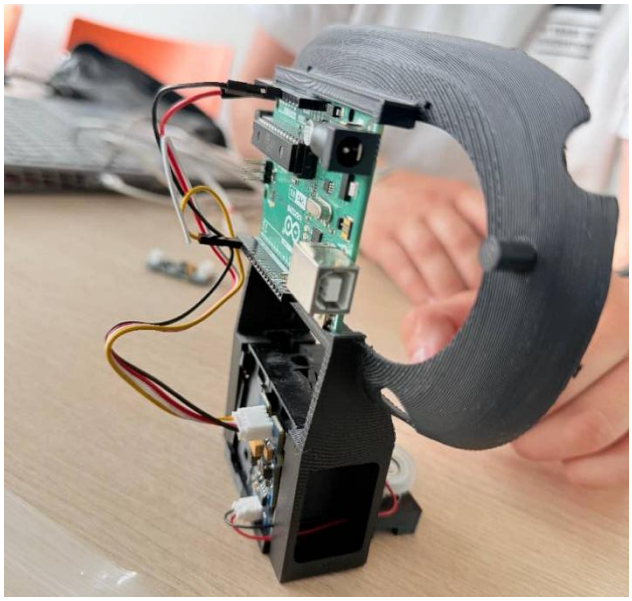

Figure 4.9. Assembled Olfactory Device: Arduino Uno and 3D-Printed Housing

## 4.2. Experimental Design

### *4.2.1. Conditions and Procedure*

Following the completion of the olfactory module integration and technical construction of the Cardia system, this study proceeds to conduct a user testing experiment in order to systematically verify whether the introduction of the olfactory channel can produce measurable relaxation benefits at both the physiological and subjective perceptual levels. This section provides a detailed account of the overall experimental design framework, participant selection criteria, data collection protocols, and analytical methods, thereby establishing the methodological foundation for subsequent experimental execution and results reporting.

The experiment is organized around two core research questions. First, can the multi-sensory VR meditation experience provided by the Cardia system produce measurable relaxation effects across both objective physiological indicators and subjective emotional self-assessment dimensions? Second, does the addition of an olfactory channel on top of the foundational audio-visual version further enhance the sense of immersion and emotional regulation outcomes, and is this incremental effect statistically verifiable?

*Experimental Design Structure*

To address the above questions, this experiment adopts a between-subject design. Condition A constitutes the baseline version (Cardia v1), encompassing visual, auditory, and

heart rate-driven scene transition channels, without the inclusion of any olfactory module, and serving as the experimental control baseline. Condition B constitutes the upgraded version (Cardia v2), which builds upon Condition A by additionally introducing a scent diffusion module, using lavender fragrance as the olfactory intervention stimulus, thereby forming the experimental condition. The two conditions are fully identical in terms of audio-visual content, experience duration, and physiological data collection protocols. The sole variable is the presence or absence of the olfactory channel, ensuring the internal validity of between-condition comparisons. Participants are randomly assigned to either Group A or Group B, with each participant completing only one condition.

| Version | Sensory Channels | Description |
|---|---|---|
| Condition A: Baseline Version (Cardia v1) | Visual + Auditory + Heart Rate-Driven Scene Transition | No olfactory module, serves as the control baseline |
| Condition B: Upgraded Version (Cardia v2) | Visual + Auditory + Heart Rate-Driven Scene Transition + Olfactory | Incorporates scent diffusion module, serves as the experimental condition |

Table 4.1. Experimental Condition Comparison: Sensory Channels and Design Rationale for Cardia v1 and Cardia v2

### *Participant Selection Criteria*

The experiment planned to recruit between 20 and 30 participants, with the target population defined by reference to the core user profile identified in the preceding questionnaire survey. The specific inclusion criteria are as follows: currently enrolled students or recent graduates aged between 18 and 30. The presence of a certain degree of everyday stress or anxiety within the preceding two weeks, operationalized as a PSS-4 score of no less than 6 as the enrollment threshold, ensuring that the sample exhibits observable relaxation needs. No contraindications to VR use, including a history of motion sickness, epilepsy, or severe visual impairment. No olfactory disorders, including hyposmia or anosmia. And no current psychiatric diagnosis, with no psychotropic medication being taken during the experimental period. In addition, to control for the confounding influence of physiological state on heart rate data, all participants were required to wait at least 45 minutes after eating before taking part in the experiment, in order to avoid heart rate baseline shifts induced by the digestive process. A total of 27 valid samples were ultimately collected, with 14 participants in Condition A and 13 in Condition B, a sample size consistent with the established conventions of quasi-experimental research within the field of design research.

*Experimental Procedure*

The first stage comprised pre-testing and informed consent. The researcher introduced the experimental objectives and procedures to each participant, after which participants independently completed the pre-test questionnaire, covering the PANAS mood scale and the SUDS subjective distress self-assessment, in order to record baseline emotional and stress states prior to the experience.

The second stage comprised physiological baseline acquisition. Participants sat in a standardized seated posture, with specific requirements as follows: both feet flat on the floor, hands resting naturally on the thighs, back in contact with the chair back, maintaining a relaxed but alert physical state throughout, with talking and the use of electronic devices prohibited for the entire duration. The BITalino PPG sensor continuously acquired pulse waveform data during this stage for a total collection duration of 5 minutes and 30 seconds, of which the first 30 seconds were excluded due to potential instability in the initial physiological state. The resulting valid baseline data duration of 5 minutes matches the valid data duration of the subsequent VR experience stage, ensuring temporal equivalence for the pre-post comparison.

The third stage comprised the VR experience phase. Prior to entering VR, the researcher provided participants with guidance on the breathing protocol: the system would guide a Box Breathing rhythm of "4 seconds inhale, 6 seconds exhale" through both visual and auditory channels, while explicitly emphasizing that participants were not required to follow this rhythm strictly, and should instead engage with the experience at whatever pace and physical state felt most comfortable to them individually, in order to reduce inter-individual interference caused by differences in prior breathing practice experience. Participants assigned to Condition A wore only the VR headset and experienced the Cardia breathing meditation scene, with the scent diffusion device remaining inactive. Participants assigned to Condition B wore the VR headset while simultaneously activating the Arduino-controlled scent diffusion device, with lavender fragrance released continuously throughout the entire duration of the experience. The VR experience lasted approximately 5 minutes, aligned with the valid duration of the baseline acquisition phase.

The fourth stage comprised the post-test questionnaire. Immediately following the conclusion of the experience, participants completed the post-test questionnaire, which included the PANAS mood scale, the SUDS subjective distress self-assessment, and the RSQ relaxation state scale, used to evaluate pre-to-post changes in emotional valence, subjective stress perception, and degree of physical relaxation following the VR experience.

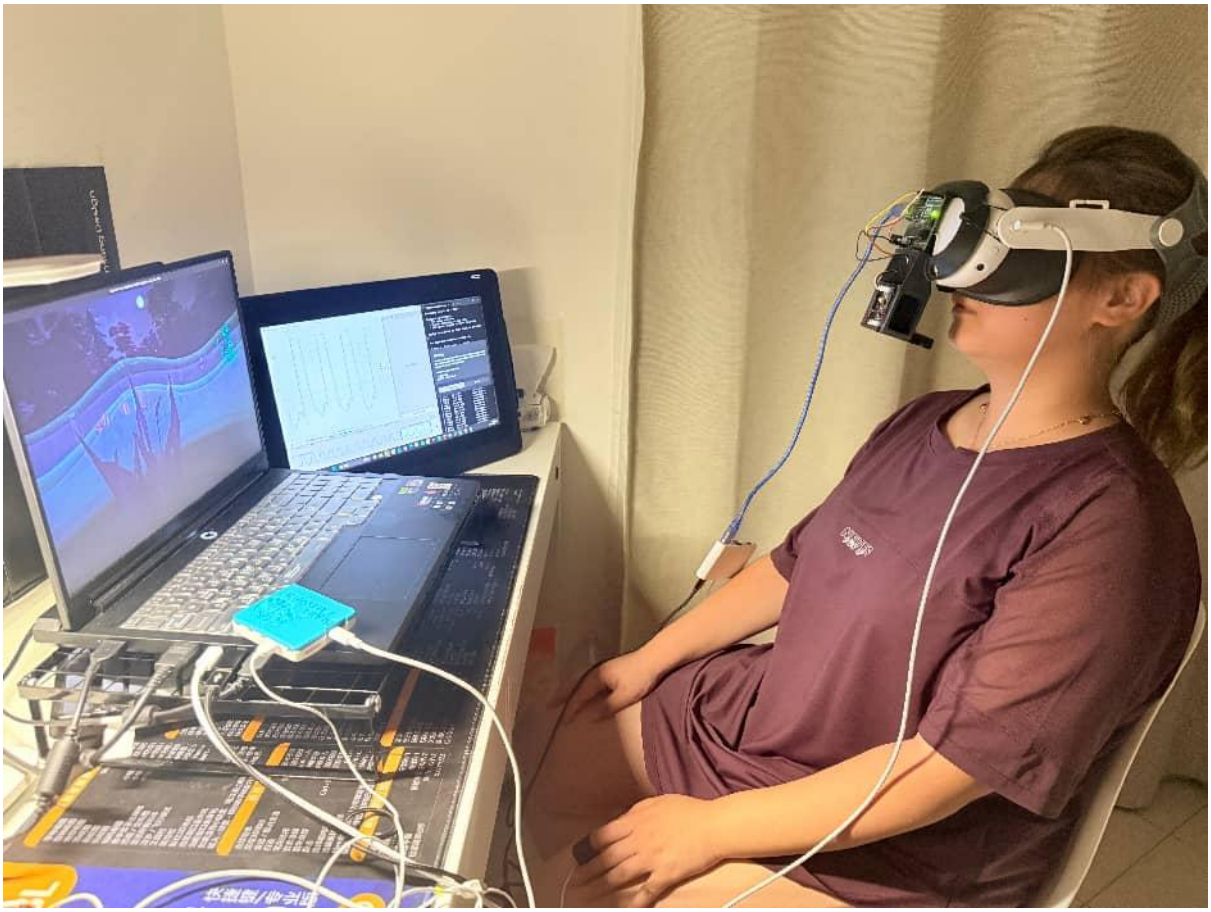

Figure 4.10. Participant Undergoing a Test Session with the VR Headset, Olfactory Module, and Real-Time Physiological Monitoring

### *4.2.2. Signal Acquisition and HRV Computation*

This experiment employs the BITalino biosensing platform for real-time physiological data acquisition, using Heart Rate Variability (HRV) as the core objective assessment metric to quantify changes in autonomic nervous system state induced by the Cardia system under multi-sensory intervention conditions. Compared to single-value mean heart rate, HRV is more sensitive in reflecting the dynamic activation level of the parasympathetic nervous system and has been widely applied in experimental research on stress and emotional regulation.

*Acquisition Metrics and Sensor Configuration*

This experiment acquires cardiac signals via the Photoplethysmography (PPG) sensor of the BITalino platform, with the sensor clipped to the tip of the participant's index finger. The OpenSignals software sampling rate was set to 100 Hz to ensure sufficient temporal resolution for HRV analysis. The acquired signal is transmitted in real time to the Python processing endpoint via the Lab Streaming Layer (LSL) protocol.

The core analytical metrics extracted from the raw PPG waveform are as follows:

| Metric | Full Name | Computation Method | Physiological Significance |
|---|---|---|---|
| HR | Heart Rate | Reciprocal of inter-peak interval × 60 | Reflects instantaneous cardiac frequency, used as an auxiliary reference metric |
| PPI | Peak-to-Peak Interval | Time difference between adjacent PPG peaks (ms) | PPG approximation of the RR interval, foundational unit for HRV computation |
| RMSSD | Root Mean Square of Successive Differences | sqrt( mean( diff(PPI)² ) ) | Core HRV metric for parasympathetic activation, higher values indicate greater relaxation |
| Stress Index | Normalized Stress Index | clip( (RMSSD_baseline − RMSSD_current) / RMSSD_baseline, 0, 1 ) | Real-time stress level relative to individual baseline, 0 = relaxed, 1 = stressed |

Table 4.2. Physiological Metrics: Definitions, Computation Methods, and Physiological Significance

Among these, RMSSD is the metric of highest priority in this experiment. Compared to frequency-domain HRV metrics such as the LF/HF ratio, RMSSD exhibits a higher signal-to-noise ratio for short-term recordings of less than 5 minutes, making it well-suited to the ultra-short-term measurement framework of this experiment.It has also been demonstrated to maintain strong correlations with standard short-term RMSSD across windows ranging from 30 seconds to 5 minutes.

*Peak Detection and PPI Extraction*

Raw PPG signals are processed for cardiac peak detection using the scipy.Signal.Find_peaks function in Python. The detection parameters are configured as follows:

| Parameter | Configured Value | Rationale |
|---|---|---|
| Sampling Rate | 100 Hz | Actual OpenSignals configuration |

| | | |
|---|---|---|
| Minimum Inter-Peak Distance | 50 frames (0.5 seconds) | Corresponds to an upper heart rate limit of approximately 120 BPM |
| Prominence Threshold | 30 (raw signal units) | Based on the actual PPG waveform amplitude range (peak ≈ 250, trough ≈ 60) |
| Valid PPI Range | 300–2,000 ms | Corresponds to a physiologically plausible heart rate range of 30–200 BPM, used to eliminate noise-induced spurious peaks |

Table 4.3. Peak Detection Parameter Configuration and Empirical Justification

To prevent the same cardiac peak from being accepted multiple times across successive frames within the detection loop — a "duplicate counting" problem — the program assigns a unique label to each accepted peak via a global frame counter. A detected peak is added to the peak timestamp list only when its associated global frame index is greater than that of the most recently recorded peak. This deduplication mechanism ensures the temporal accuracy of the PPI sequence.

### *Baseline Protocol*

Baseline data constitutes a prerequisite for within-subject comparative analysis in this experiment. Given that absolute RMSSD values vary substantially across participants — with individual resting RMSSD typically ranging between 15 and 80 ms — cross-participant comparison of absolute RMSSD values would be meaningless without the acquisition of a personal baseline. This experiment follows the baseline acquisition procedure described below:

| Stage | Duration | Procedure | Data Handling |
|---|---|---|---|
| Adaptation Period | 0–30 seconds | Participant sits quietly. BITalino begins acquisition, no instructions given | Data collected but discarded, excluded from baseline computation |
| Stable Acquisition Period | 30–330 seconds | Participant continues to sit and relax, researcher confirms | All PPI data included in baseline RMSSD computation |

| | | | |
|---|---|---|---|
| | | signal quality is stable | |
| Baseline Computation | At 330 seconds | Program automatically extracts valid peak timestamp list | Outputs baseline_rmssd, written to CSV file header comment row |

Table 4.4. Baseline Acquisition Stage Design: Timing, Procedure, and Data Handling

The rationale for discarding the adaptation period is as follows: upon entering the laboratory, participants require time to acclimate to the environment, and heart rate and electrodermal activity during the initial 30 seconds are typically still in a state of dynamic adjustment. Including this period in the computation would inflate baseline noise. Following the discard of the first30 seconds, the subsequent 300-second PPI sequence is used to compute RMSSD, which primarily reflects the activity level of the parasympathetic (vagal) nervous system:

- **Relaxed state:**parasympathetic dominance. Inter-beat interval fluctuation is larg, RMSSD is high (typically > 40 ms)
- **Stressed state:** sympathetic dominance. Heartbeat tends toward metronomic regularity, RMSSD is low

This value serves as the individual physiological baseline for the participant in the current experimental session.

The baseline RMSSD is subsequently used to compute the normalized Stress Index during the VR experience phase, according to the following formula:

| Formula | Interpretation |
|---|---|
| S = clip( (RMSSD_baseline − RMSSD_current) / RMSSD_baseline, 0, 1 ) | S = 0: current RMSSD ≥ baseline (equivalent to or more relaxed than baseline). S = 1: current RMSSD substantially below baseline (significantly more stressed), returns 0.5 (neutral value) when data is insufficient |

Table 4.5. Normalized Stress Index Computation Formula and Interpretation

This normalization procedure eliminates the influence of inter-individual differences in absolute RMSSD values, rendering the Stress Index comparable across participants. It simultaneously serves as the real-time control signal driving the scene transition logic within the UE5 environment.

### *Real-Time HRV Computation*

Upon entering the VR experience phase, the program computes real-time RMSSD using a 30-second sliding window applied to the peak timestamp sequence: at each frame, only the valid peaks falling within the 30-second window preceding the current time point are extracted, and the RMSSD value computed over this window serves as the real-time HRV estimate at that moment.

The selection of a 30-second window is based on the following considerations. First, a 30-second window has been demonstrated to yield statistically stable RMSSD estimates that maintain significant correlations with standard 5-minute short-term RMSSD[58]. Second, windows shorter than 10 seconds produce excessive RMSSD variability that undermines interpretability. Third, a 30-second window aligns with the experiential pacing of this experiment — with a breathing guidance cycle of approximately 8 to 12 seconds — enabling the capture of HRV changes induced by breathing regulation. This window falls within the category of ultra-short-term HRV. The relevant methodological justifications are provided in Shaffer & Ginsberg (2017)[59] and Munoz et al. (2015)[60].

### *Data Logging*

All physiological data are written in real time to a CSV (Comma-Separated Values) file during acquisition, with filenames incorporating the experimental date and timestamp for archival purposes. The CSV recording structure is as follows:

| Field Name | Type | Description |
|---|---|---|
| time_s | Float | Absolute time elapsed since experiment start (seconds), precise to 0.001 s |
| phase | String | "baseline" (baseline phase) or "session" (VR experience phase) |
| raw_ppg | Float | Raw PPG sample value (OpenSignals native units) |
| ppi_ms | Float | Current PPI interval (milliseconds), empty when no new peak is detected |
| bpm | Float | Instantaneous heart rate estimate (BPM) |
| rmssd_window | Float | 30-second sliding window RMSSD (milliseconds), empty when data is insufficient |

| | | |
|---|---|---|
| stress_index | Float | Normalized stress index (0.0–1.0), transmitted synchronously to UE5 |

Table 4.6. CSV Data Recording Fields: Names, Types, and Descriptions

Upon conclusion of the baseline phase, the program automatically writes the following summary information to the CSV as a comment row, facilitating rapid identification during post-hoc analysis: participant baseline mean HR (BPM), baseline RMSSD (ms), number of valid peaks, sliding window duration (seconds), and actual sampling rate (Hz).

*Post-Session Analysis*

Following the conclusion of the experiment, the following statistical analyses will be conducted on the CSV data to evaluate differences in physiological intervention effects between Cardia v1 (Condition A) and Cardia v2 with olfactory module (Condition B):

| **Comparison Objective** | **Method** | **Key Metric** | **Significance Criterion** |
|---|---|---|---|
| Condition A post-experience vs. Condition A baseline | Paired t-test / Wilcoxon signed-rank test | RMSSD change (ms) | $p < 0.05$ |
| Condition B post-experience vs. Condition B baseline | Paired t-test / Wilcoxon signed-rank test | RMSSD change (ms) | $p < 0.05$ |
| Condition B change magnitude vs. Condition A change magnitude | Paired t-test (difference-in-differences) | delta_RMSSD_B − delta_RMSSD_A | $p < 0.05$ |
| Normality test | Shapiro-Wilk test | Determines use of parametric or non-parametric method | $p > 0.05$ indicates normality |

Table 4.7. Physiological Data Statistical Testing Plan: Comparison Objectives, Methods, and Significance Criteria

Given the relatively limited sample size of this experiment (N = 20–30), all statistical results will additionally report effect size (Cohen's *d*) to supplement statistical power information. The time-series trajectory of RMSSD change will be presented in line graph form, with breathing guidance nodes and scene transition time points annotated, in order to

visually reveal the correspondence between physiological changes and designed intervention points.

### *4.2.3. Subjective Assessment Instruments*

The subjective questionnaire battery consists of four standardized, psychometrically validated instruments PSS-10, PANAS, SUDS, and RSQ designed to systematically capture participants' subjective psychological changes before and after the VR experience across four dimensions: stress perception, emotional valence, immediate distress level, and somatic relaxation state. All scales were administered twice once prior to baseline acquisition (pre-test) and once immediately following the conclusion of the VR experience (post-test) presented in identical content and sequence on both occasions, ensuring direct comparability between pre- and post-test data. This two-timepoint measurement design not only enables examination of the Cardia experience's regulatory effects on immediate emotional state, but also provides a standardized subjective data foundation for cross-condition comparison between Condition A and Condition B.

*PSS-10 Perceived Stress Scale*

The Perceived Stress Scale-10 (PSS-10), developed by Cohen, Kamarck, and Mermelstein in 1983, is one of the most widely used subjective stress self-assessment instruments in psychology and health sciences [61]. The scale comprises 10 items, asking participants to recall the frequency with which specific feelings or thoughts occurred over the past month, rated on a 5-point Likert scale (0 = never, 4 = very often), with total scores ranging from 0 to 40, where higher scores indicate greater perceived stress. In terms of scoring rules, items **4, 5, 7, and 8**are reverse-scored. This cross-design of forward and reverse items effectively suppresses response bias and enhances internal consistency of measurement. The PSS-10 typically yields a Cronbach's α coefficient between 0.84 and 0.86, demonstrating strong internal consistency reliability. Its construct validity has been confirmed across multiple cross-cultural studies, with significant correlations observed with objective indicators including cortisol levels, immune function, and health behaviors [62]. In the present study, the PSS-10 serves primarily as a participant screening tool (pre-test score ≥ 6 as the enrollment threshold) and as a quantitative reference for baseline stress level, rather than as a core dependent variable for between-condition comparison, ensuring sufficient homogeneity within the sample on the stress perception dimension.

*PANAS Positive and Negative Affect Schedule*

The Positive and Negative Affect Schedule (PANAS), developed by Watson, Clark, and Tellegen in 1988, is the most widely applied two-dimensional emotional self-assessment instrument in the field of emotion research[63]. The scale comprises 20 emotional adjectives, with participants rating the degree to which they experience each emotional state within a specified time frame — in this study, set to "right now" — on a 5-point Likert scale (1 = very slightly or not at all, 5 = extremely). The 20 items are divided into two independent subscales: the Positive Affect subscale (PA, 10 items, encompassing positively activated states such as "excited," "determined," and "attentive") and the Negative Affect subscale (NA, 10 items, encompassing negatively activated states such as "distressed," "nervous," and "afraid"), with both subscales scoring on a range of 10 to 50. The theoretical foundation of the PANAS rests on the two-factor independent structural model of emotion, which holds that positive and negative affect are not opposing poles of a single dimension but rather relatively independent affective systems capable of varying simultaneously, thereby conferring greater measurement sensitivity in intervention research. The scale demonstrates excellent psychometric properties, with internal consistency coefficients for both the PA and NA subscales falling between 0.86 and 0.90, and has been extensively validated across diverse cultural contexts and clinical settings. In the present study, pre-to-post changes in PANAS PA scores are used to assess the extent to which the VR experience elevates positive emotional activation, while changes in NA scores serve as a sensitive indicator of negative emotion reduction. Together, these constitute the core dimensions of subjective emotional assessment.

*SUDS Subjective Units of Distress Scale*

The Subjective Units of Distress Scale (SUDS) is a single-item immediate distress assessment instrument presented on a continuous scale from 0 to 100, where 0 represents complete calm and 100 represents the most intense distress imaginable (Wolpe, 1969). Its visually intuitive "distress thermometer" format enables participants to rapidly self-locate their current state without verbal guidance, and the scale has been widely used in clinical exposure therapy and stress intervention research to monitor real-time fluctuations in emotional state. The unidimensional design of SUDS renders it highly sensitive to immediate emotional change, making it well-suited for rapid repeated administration before and after an experience, serving as an immediate-state auxiliary indicator complementing other multidimensional scales.

*RSQ Relaxation State Questionnaire*

The Relaxation State Questionnaire (RSQ) is a multidimensional psychometric instrument specifically designed to assess the immediate effects of relaxation interventions [64]. The scale consists of 10 items rated on a 5-point Likert scale, with items **1, 2, 3, and 10** reverse-scored prior to computation, ensuring that higher total scores consistently correspond to greater levels of relaxation. The RSQ encompasses four functional subscales: the Cardiovascular subscale (items **1** and **2**, measuring subjectively perceived changes in heart rate and respiration). The Muscular subscale (items **3, 4, and 5**, assessing the degree of tension and release in somatic musculature). The General Relaxation subscale (items **6** and **7**, capturing subjective sense of calm and overall mind-body relaxation). And the Drowsiness subscale (items **8, 9, and 10**, distinguishing genuine relaxation from mere fatigue or sleepiness, functioning as an internal validity control for the measurement). The overall internal consistency of the scale is $\alpha = 0.86$, and criterion validity has been established through negative correlations with perceived tension and positive correlations with positive affect. In the present study, the RSQ is administered exclusively at the post-test stage, serving as a dedicated instrument for evaluating the immediate somatic relaxation effects of the VR experience. The score changes across its four subscales provide multidimensional supplementary evidence at the physiological perception level, complementing the emotional measurement outcomes of the PANAS and SUDS.

## 4.3. Experimental Results and Analysis

### *4.3.1. Physiological Data Results*

This study collected BITalino PPG physiological data from a total of 27 participants, with 14 participants assigned to Condition A (pure VR experience) and 13 to Condition B (VR + lavender olfactory stimulation). The core quantitative metric is ΔRMSSD — defined as the percentage change in mean session-phase RMSSD relative to each participant's individual baseline — where positive values indicate an increase in parasympathetic nervous system activity (a relaxation tendency) and negative values indicate sympathetic nervous system dominance (a stress or arousal tendency).

*Condition A (Pure VR): Overall Results*

Among the 14 participants in Condition A, 11 exhibited positive RMSSD growth, while 3 showed decreases of varying magnitudes (**A_06**: −4.8%. **A_13**: −4.3%. **A_14**: −21.2%). The mean ΔRMSSD for Condition A was approximately **+11.2%**, with a median of approximately

+7.5%. The most pronounced responses were observed in **A_07** (+44.1%) and **A_08** (+27%), indicating that certain participants were already capable of generating a substantial parasympathetic activation effect within a purely audio-visual VR meditation environment. Examining the group-level trend, the RMSSD time-series profiles of the majority of Condition A participants during the experience phase display a characteristic pattern of brief initial fluctuation followed by gradual stabilization at or slightly above the baseline level — a pattern that aligns closely with participants' subjective descriptions of progressively establishing a relaxed state under the guidance of the VR immersive environment.

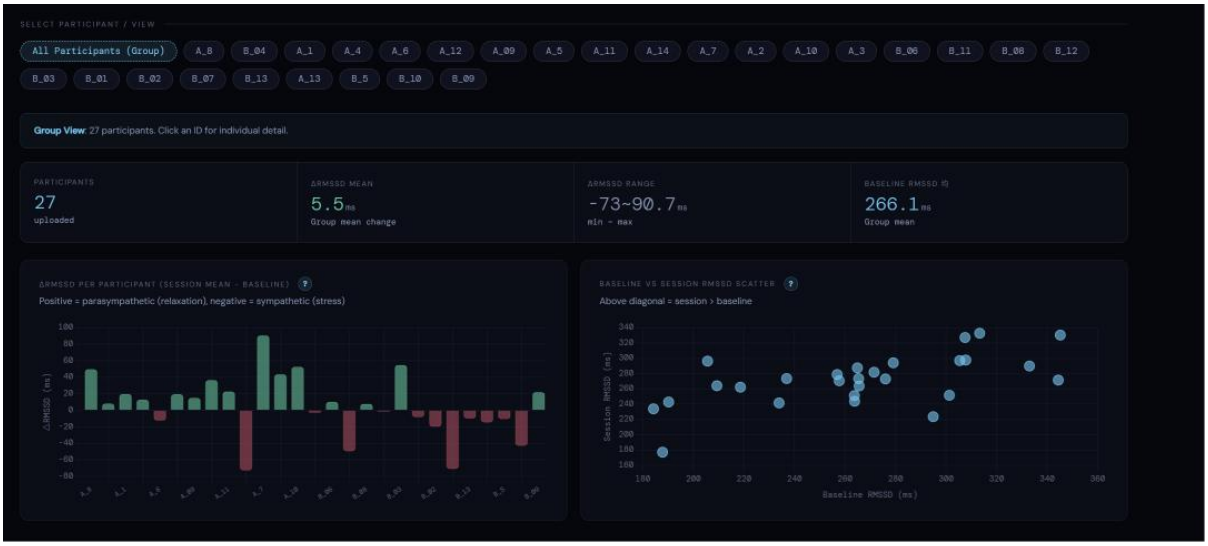


Figure 4.11. Group-Level ΔRMSSD Overview Dashboard: Per-Participant Change and Baseline-to-Session Scatter Plot

*Condition B (VR + Lavender): Overall Results*

Among the 13 participants in Condition B, only 5 exhibited positive RMSSD changes (**B_03**: +26.2%. **B_04**: +3%. **B_06**: +3.8%. **B_08**: +3.3%. **B_09**: +8.6%), while the remaining 8 showed decreases of varying magnitudes, with **B_07** recording the largest decline (−24.1%), and **B_10** (−12.9%) and **B_11** (−16.5%) also exhibiting notably negative changes. The mean ΔRMSSD for Condition B was approximately **−1.7%**, with an overall directional tendency opposite to that of Condition A.

*Between-Condition Comparison and Interpretation*

At the numerical level, the mean ΔRMSSD of Condition A (+11.2%) substantially exceeds that of Condition B (−1.7%), representing a difference of approximately 12.9 percentage points. On initial observation, this result appears to contradict the research hypothesis that "the addition of olfactory stimulation would enhance relaxation outcomes."

However, interpreting this finding requires careful consideration of several important contextual factors pertaining to data quality and experimental conditions.

*Anomalous Data and Discussion of Result Limitations*

First, the baseline data of certain participants in Condition B exhibit marked anomalies. Taking **B_01** as an example, this participant's baseline heart rate was recorded at 188.7 bpm — far exceeding the normal resting heart rate range — suggesting that during the baseline acquisition phase, confounding factors such as physical activity, emotional tension, or inadequate sensor contact may have been present, resulting in an artificially elevated baseline RMSSD and consequently driving the session-phase ΔRMSSD into negative territory. A comparable situation was observed in **B_05** (baseline BPM: 202.5, markedly anomalous). Such baseline quality issues directly undermine the reliability of the group mean and constitute one of the primary confounding sources behind Condition B's overall suppressed ΔRMSSD.

Second, the RMSSD metric is highly sensitive to breathing patterns. In Condition B, the introduction of the olfactory stimulus may have briefly diverted certain participants' attentional resources away from the breathing guidance during the initial phase of the experience, resulting in transient irregularities in breathing rhythm and a corresponding temporary decline in RMSSD. This phenomenon is consistent with the early-phase fluctuation patterns visible in the time-series profiles of several Condition B participants.

Furthermore, the subjective experience of lavender scent is not uniformly calming across individuals with differing olfactory preferences. A small number of participants may have produced a mild stress response to the specific scent, an effect that cannot be fully eliminated through screening in small-sample experiments characterized by substantial individual variability.

It is nonetheless noteworthy that, despite the negative mean RMSSD direction observed for Condition B as a whole, participants such as **B_03** (+26.2%) demonstrated strong relaxation responses under the VR + lavender condition, indicating that the olfactory intervention did produce the anticipated effect for certain individuals. Taken in conjunction with the subsequent subjective questionnaire data, the between-condition differences in self-reported relaxation levels and emotional valence will provide a more complete explanatory framework.

Considered in its entirety, the physiological data outcomes of this study both affirm the feasibility of VR meditation experience in objectively inducing parasympathetic activation

and expose the considerable challenges associated with baseline quality control and individual difference management in multi-sensory biofeedback research conducted within small-sample, real-world laboratory environments. These limitations will be elaborated upon further in Section 4.3.3: Integrated Discussion.

### *4.3.2. Subjective Questionnaire Results*

This section presents a systematic analysis of the subjective experience outcomes for Condition A (pure VR) and Condition B (VR + lavender) across four dimensions — stress perception, emotional valence, immediate calmness, and somatic relaxation state — based on the self-report scale data completed by 27 participants before and after the VR experience.

*Overall Between-Condition Comparison*

Examining the group mean data aggregated by the analysis platform, both conditions exhibited a clearly positive directional trend across multiple dimensions, though the between-condition differences across individual metrics showed varying degrees of divergence. With respect to the PSS-10, the post-test mean for Condition A was 18.3 points and for Condition B was 18.1 points — a negligible difference, with both groups falling within the moderate stress range. This result is attributable to the inherent measurement time frame of the PSS-10 instrument, which asks participants to reflect on their stress experience over the past month. A single 5-minute VR intervention is unlikely to statistically alter participants' retrospective evaluation of their overall stress state across the preceding month. Accordingly, the PSS-10 functions primarily in this study as a participant screening and stress background characterization tool, rather than as a core indicator of immediate intervention effects.




Figure 4.12. Group-Level Post-Test Scale Means Comparison and PSS-10 Post-Test Score Distribution by Condition

| Scale | All (N=27) | Condition A (N=14) | Condition B (N=13) | A vs B Direction |
|---|---|---|---|---|
| PSS-10 | 18.3 / 40 | 18.3 | 18.1 | B better ↓ |
| PANAS-PA | 28 / 50 | 24.8 | 29.5 | B better ↑ |
| PANAS-NA | 21.3 / 50 | 21.3 | 21.8 | A better |
| SUDS | 35.5 / 100 | 41.1 | 29.2 | B better ↓ |
| RSQ | 3.8 / 5 | 3.9 | 3.8 | A better |

Table 4.8. Scale Means Summary: All Participants, Condition A, and Condition B Post-Test Scores Across PSS-10, PANAS-PA, PANAS-NA, SUDS, and RSQ

### *PANAS Emotional Valence: Condition B Demonstrates Substantially Higher Positive Affect Than Condition A*

On the emotional valence dimension, Condition B's mean Positive Affect (PA) score on the PANAS was 29.5 points, markedly higher than Condition A's mean of 24.8 points — a gap of 4.7 points. Furthermore, Condition B's mean PA score exceeds the normative mean of 33.3 points reported by Watson et al. (1988), approaching the level of normative positive affect, while Condition A remains below the normative mean. This divergence suggests that the introduction of lavender olfactory stimulation may have produced an additional augmentation effect at the level of positive emotional activation, leading participants to report higher levels of energy, pleasure, and engagement following the experience. With respect to

Negative Affect (NA), Condition A recorded a mean of 21.3 points and Condition B 21.8 points. Both groups were marginally above the normative level (17.4), with a negligible between-condition difference. This indicates that the immediate effects of the two conditions on negative emotion reduction were comparable, with neither condition achieving a substantial reduction in baseline negative affect within the 5-minute intervention window — an outcome consistent with the ceiling effects commonly observed in short-duration meditation intervention research.

*SUDS Immediate Calmness: Condition B Exhibits Markedly Lower Subjective Distress*

The SUDS proved to be the most sensitive measure of immediate emotional state in this study, and its results revealed the most pronounced between-condition difference of all instruments. The post-test mean SUDS score for Condition B was 29.2 points, compared to 41.1 points for Condition A — a difference of nearly 12 points in Condition B's favor. Since lower SUDS scores correspond to lower subjective distress and greater calmness, this result indicates that Condition B participants reported an overall lower level of subjective tension and discomfort at the conclusion of the VR experience. In terms of pre-to-post change magnitude, Condition A showed an average SUDS reduction of approximately 27 points (estimated from individual change means), while Condition B showed an average reduction of approximately 21 points. However, Condition B's lower post-test absolute score indicates that participants in that condition reached a more tranquil subjective state by the end of the experience, suggesting that the introduction of the olfactory channel produced an observable positive effect on the dimension of immediate calmness.

Individual-level data further corroborate this trend. Multiple Condition B participants recorded extremely low post-test SUDS scores — for example, **B_02** (4/100), **B_01** (19/100), and **B_09**(24/100) all reported subjective states approaching complete calm, and received "Consistent" ratings in the physiology–subjective convergence assessment, indicating that their subjective experience aligned directionally with their physiological data.

*RSQ Relaxation State: Groups Broadly Comparable, with Pronounced Individual Variability*

On the RSQ total score, Condition A recorded a mean of 3.9 and Condition B 3.8 — a minimal difference, with the platform system designating Condition A as marginally superior. However, deeper analysis of the individual subscales reveals that Condition B participants generally scored higher on the Cardiovascular subscale (perceived heart rate and breathing steadiness), with multiple participants — including **B_02**, **B_03**, **B_11**, and **B_13**— achieving maximum scores of 5/5 on this subscale, reflecting a possible auxiliary effect of lavender

scent on the perceptual dimension of breathing. Condition A exhibited more stable score patterns on the Drowsiness subscale, suggesting that the pure audio-visual VR experience produced a degree of cognitive arousal (VR cognitive arousal) in certain participants — an observation consistent with the elevated SUDS scores and platform annotations of "Moderate discomfort — note VR cognitive arousal" associated with participants such as**A_03**, **A_08**, and **A_11**.

*Physiology–Subjective Convergence*

He "Physio × Subjective Convergence" module of the analysis platform conducted item-level convergence assessments between physiological and subjective data for each participant. At the group level, SUDS post-test low scores and RSQ General subscale scores received "Consistent" ratings in the majority of participants, indicating a high degree of internal consistency between self-reported calmness and overall relaxation state assessment. In contrast, the correspondence between PANAS PA–NA difference values and the directional change in RMSSD varied considerably across individuals. Certain participants — including**A_08**, **A_11**, **B_06**, and **B_13,** exhibited divergence between subjective emotional valence and objective physiological indicators, suggesting that the relationship between subjective emotional experience and autonomic nervous system regulation is not linearly straightforward. Individual differences in degree of cognitive engagement with the VR environment, novelty responses to VR, and stress reactivity may all function as mediating variables influencing the correspondence between the two.

Considered as a whole, the subjective questionnaire data provide preliminary support for the research hypothesis that "the introduction of an olfactory channel produces additional positive effects," as evidenced across the dimensions of emotional valence (PANAS-PA) and immediate calmness (SUDS). The RSQ and PSS-10 results, by contrast, more closely reflect the limitations of a short-duration intervention on the dimensions of trait-level relaxation and chronic stress. These findings will be integrated with the physiological data and interpreted holistically in Section 4.3.3: Integrated Discussion.

### *4.3.3. Integrated Discussion*

This study evaluated the relaxation intervention effects of the Cardia system through two parallel pathways — physiological data and subjective self-report — and systematically examined the convergent validity between these two data streams using the platform's integrated analysis module. The following integrated discussion is organized around three

core questions: whether the Cardia system as a whole demonstrates measurable relaxation effects, whether the introduction of the olfactory channel produced observable additive benefits, and what degree of convergence exists between physiological and subjective data.

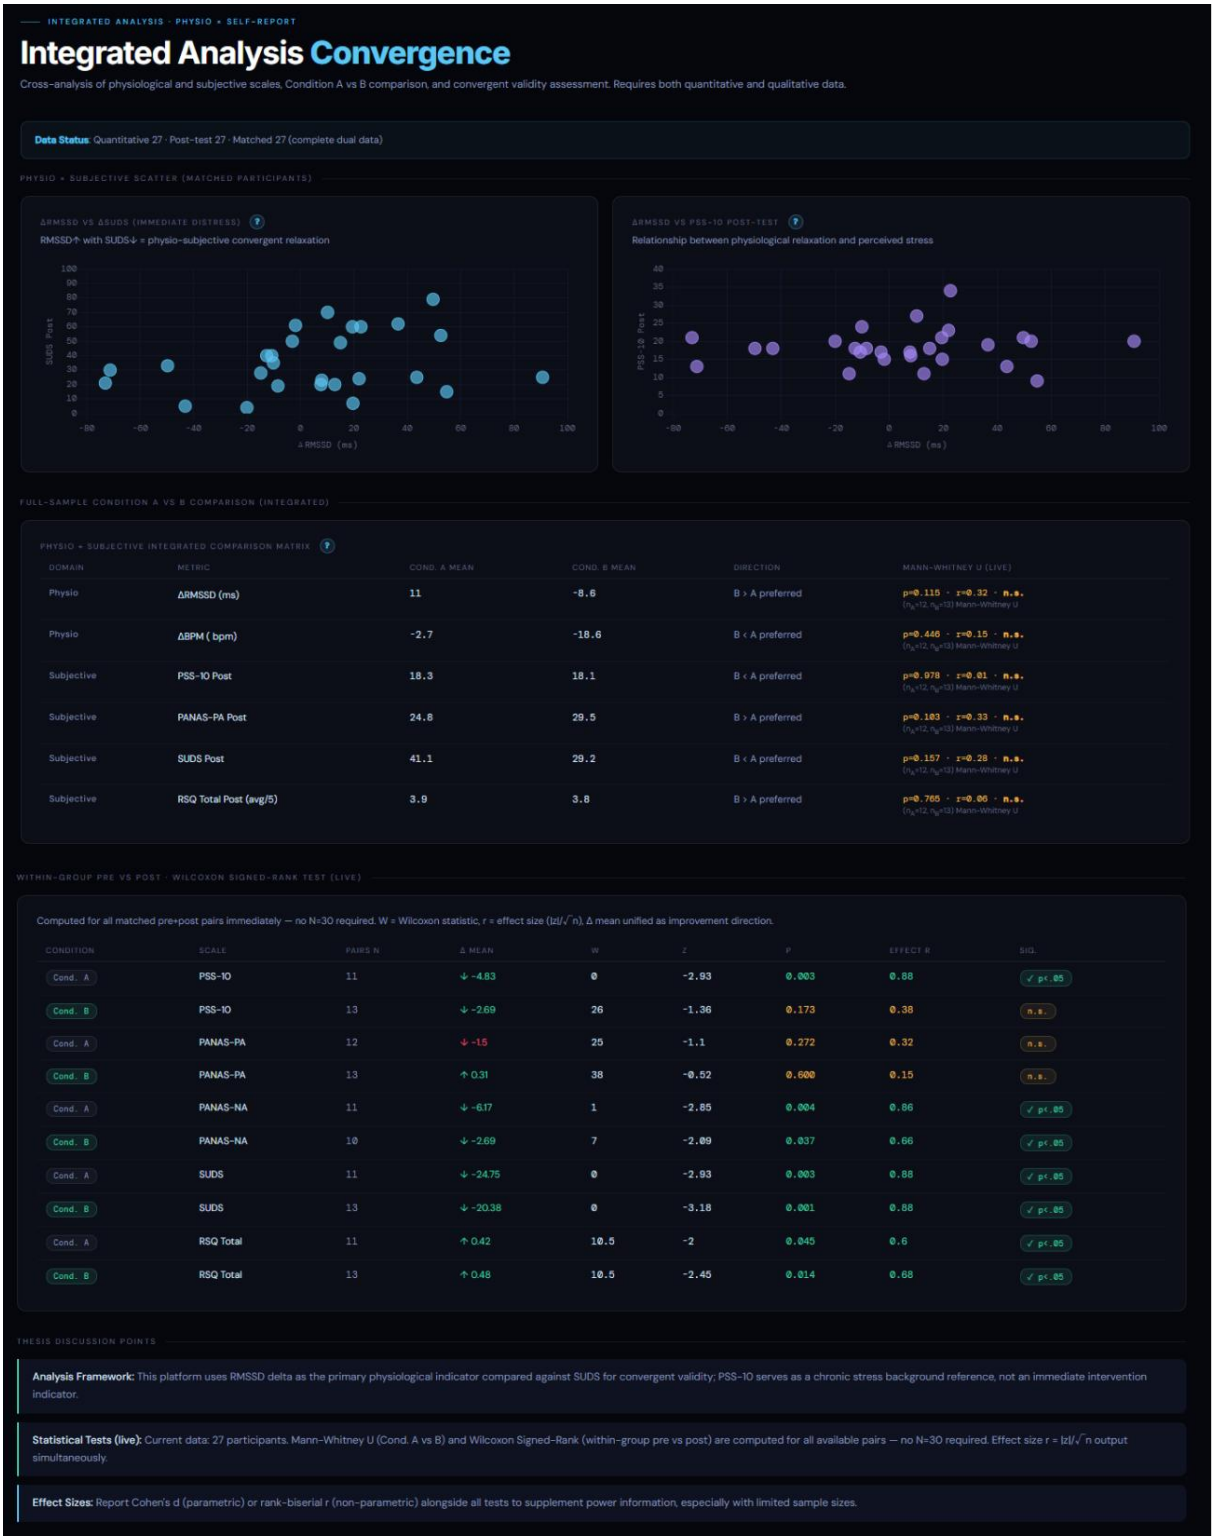

### *Within-Group Pre-to-Post Effects: Both Conditions Produce Significant Subjective Improvement*

The within-group pre-to-post comparison results from the Wilcoxon signed-rank test provide the clearest and most robust core evidence in this study. On the SUDS immediate distress dimension, both Condition A (Δ mean = −24.75, W = 0, Z = −2.93, p = 0.003, effect size r = 0.88) and Condition B (Δ mean = −20.38, W = 0, Z = −3.18, p = 0.001, effect size r = 0.88) achieved statistical significance, with both yielding effect sizes of 0.88 — indicating an extremely large effect. This means that regardless of whether participants experienced the pure audio-visual VR environment or the multi-sensory version augmented with olfactory stimulation, the intervention was capable of reducing immediate subjective distress by approximately 20 to 25 points within 5 minutes, a magnitude of effect that carries substantive significance within the context of design intervention research.

On the PANAS Negative Affect dimension, both Condition A (Δ mean = −6.17, p = 0.004, r = 0.86) and Condition B (Δ mean = −2.69, p = 0.037, r = 0.66) reached significance, confirming that both conditions effectively attenuated negative emotional activation following the experience, with Condition A yielding a marginally larger effect size. On the RSQ Total Relaxation Score, both Condition A (Δ mean = +0.42, p = 0.045, r = 0.60) and Condition B (Δ mean = +0.48, p = 0.014, r = 0.68) were likewise significant, notably, Condition B's effect size (0.68) slightly exceeded that of Condition A (0.60), suggesting that the addition of lavender olfactory stimulation was associated with a modest improvement in participants' self-reported somatic relaxation state at the conclusion of the experience.

With respect to the PSS-10, Condition A demonstrated a significant pre-to-post decline (Δ mean = −4.83, p = 0.003, r = 0.88), whereas Condition B did not reach significance (Δ mean = −2.69, p = 0.173, r = 0.38). As noted previously, the PSS-10 measures chronic perceived stress over the preceding month, the significant change observed in Condition A more likely reflects a retrospective re-evaluation of participants' "past month" stress state at the time of completing the post-test questionnaire, rather than a direct intervention effect of the Cardia system on chronic stress levels, and should therefore be interpreted with caution. PANAS Positive Affect failed to reach significance in either condition (Condition A: p = 0.272. Condition B: p = 0.600), suggesting that 5 minutes of meditation experience produces limited gains in active positive affect activation, and that longer intervention durations or repeated exposure may be required to yield statistically meaningful positive emotion enhancement.

Taken together, the within-group analysis confirms that the Cardia multi-sensory VR meditation experience produces statistically significant and large-magnitude immediate effects on subjective stress reduction (SUDS), negative emotion attenuation (PANAS-NA), and somatic relaxation enhancement (RSQ), thereby addressing the first core research question of this study.

*Between-Condition Comparison: Condition B Shows a Favorable Directional Trend on Subjective Dimensions, Without Reaching Statistical Significance*

The Mann-Whitney U between-condition comparison tests revealed that Condition A and Condition B did not differ significantly on any metric (p values ranging from 0.103 to 0.978). Specifically, on ΔRMSSD, Condition A recorded a mean of +11 ms versus Condition B's −8.6 ms ($p = 0.115$, $r = 0.32$, n.s.), on ΔBPM, Condition A showed a reduction of 2.7 bpm versus Condition B's 18.6 bpm — a larger directional reduction in Condition B — but this likewise did not reach significance ($p = 0.446$, $r = 0.15$), on PANAS-PA, Condition B (29.5) exceeded Condition A (24.8) by 4.7 points ($p = 0.103$, $r = 0.33$, n.s.), on SUDS post-test, Condition B (29.2) was nearly 12 points lower than Condition A (41.1) ($p = 0.157$, $r = 0.28$, n.s.), and RSQ total scores were closely matched between groups ($p = 0.765$, n.s.).

These results indicate that although Condition B's mean scores on PANAS-PA and SUDS were directionally superior to Condition A, and the absolute difference on SUDS (12 points) carries a degree of clinical meaningfulness, none of these differences surpassed the threshold of statistical significance under the current sample size ($n_A = 14$, $n_B = 13$). This is closely related to insufficient statistical power in small-sample testing: given effect sizes in the range of $r = 0.28$ to 0.33, achieving 80% statistical power at $\alpha = 0.05$ would typically require approximately 40 to 50 participants per group. The sample size of this study inherently constrains the inferential capacity of the statistical tests, nonetheless, the consistent directionality of the observed trends provides preliminary exploratory support for the research hypothesis.

*Convergent Validity Analysis of Physiological and Subjective Data*

The platform's ΔRMSSD vs. ΔSUDS scatter plot reveals an important research finding: physiological relaxation (RMSSD increase) and immediate subjective distress (SUDS decrease) do not exhibit a clear linear convergent relationship. Certain participants showed both a positive RMSSD increase and a substantial SUDS reduction (e.g., A_07, A_04, B_03), reflecting an ideal pattern of dual physiological–subjective convergent relaxation, while other participants displayed negative RMSSD change alongside an equally pronounced SUDS

reduction (e.g., B_07, B_11), indicating that the establishment of subjective calmness does not necessarily depend on concurrent improvement in objective autonomic nervous system status.

This dissociation is not uncommon in VR immersive experience research. The cognitive flow generated by VR immersion may enable users to experience relaxation and calmness at the subjective level in advance of the autonomic nervous system's actual regulatory response, which requires a longer time window to reach stable adjustment. Additionally, the mild physiological arousal produced in certain participants by the novelty of the VR environment — such as a transient heart rate increase or RMSSD fluctuation — can disrupt physiological relaxation signals without necessarily impairing the quality of subjective experience. The ΔRMSSD vs. PSS-10 post-test scatter plot similarly displays a broadly dispersed point cloud, further confirming the weak correlation between chronic stress self-report and immediate physiological state, and corroborating the interpretive limitations of the PSS-10 discussed earlier in this chapter.

*Research Limitations and Design Reflections*

Synthesizing the analytical results from both data streams, this study presents several methodological limitations worthy of critical reflection. First, the small sample size constrained the statistical power of between-group comparisons, certain effect differences with genuine practical significance could not be translated into significant p-values at the current sample scale, and future research should expand to a minimum of 40 participants per group. Second, the 5-minute intervention duration is relatively brief for meditation-based interventions that require deep autonomic nervous system regulation, extending the experience to 10–15 minutes may produce more stable physiological responses. Third, the intensity and concentration of the olfactory stimulus were not standardized at the individual level during the experiment, individual differences in olfactory detection thresholds for the same lavender concentration may have resulted in certain participants receiving an effective olfactory exposure below the intended level. Fourth, because Condition A and Condition B were administered to independent groups of participants rather than the same individuals, the between-subject design cannot fully rule out pre-existing differences between the two groups — such as baseline stress reactivity, prior VR experience, or individual differences in olfactory sensitivity — despite random assignment; a larger sample or a supplementary within-subject validation would help address this limitation in future work.

Notwithstanding the limitations outlined above, the integrated analytical results of this study yield valuable exploratory evidence. The Cardia multi-sensory VR meditation system demonstrates statistically significant effects on immediate subjective stress relief and promotion of somatic relaxation, with both conditions producing large-effect-size SUDS improvements ($r = 0.88$) and medium-to-large-effect-size RSQ gains within 5 minutes. While the introduction of the olfactory channel did not produce statistically significant between-group advantages at the current sample size, it consistently demonstrated a favorable directional advantage on PANAS Positive Affect activation and SUDS absolute scores. Grounded in the established neurophysiological literature on lavender's mechanisms of action, these converging directional signals provide a reasonable exploratory basis for future research with larger samples and more refined multi-sensory VR meditation design.

# 5. Conclusion

This study takes the mental health of young adults as its starting point, focusing on the real-world challenge of stress accumulation and emotional dysregulation in contemporary society. In response, a multi-sensory immersive VR meditation experience system — Cardia — was designed and developed, integrating four sensory dimensions: visually immersive scenes, breathing guidance rhythms, real-time physiological biofeedback, and olfactory stimulation. Using virtual reality technology as its core medium and evidence-based psychological theory as its design foundation, the system constructs an active emotional relaxation intervention tailored to young adults aged 18 to 34. This chapter restates the objectives of the research, reviews the work that has been completed, summarises the study's achievements, and reflects on its limitations and directions for future development.

## 5.1. Research Objectives

This study set out to determine whether a multi-sensory, immersive VR meditation system could provide young adults with a meaningful, accessible, and physiologically grounded approach to emotional regulation. Building on this central aim, two more specific objectives guided the research: first, to design and develop a VR meditation system — Cardia — that integrates visual immersion, breathing guidance, real-time physiological biofeedback, and olfactory stimulation into a single closed-loop experience, and second, to empirically test whether the addition of an olfactory channel produces a measurable incremental benefit over an audio-visual-only VR meditation experience, assessed through both physiological (ΔRMSSD) and subjective (PSS-10, PANAS, SUDS, RSQ) measures.

## 5.2. Completed Work

During the research phase, baseline stress levels were assessed among 18 young adult participants using the PSS-10 and PANAS scales. Results indicated that the sample generally exhibited moderate-to-high levels of perceived stress, with primary sources concentrated in academic pressure, workplace competition, and social adaptation, and with only limited existing access to active, non-clinical coping strategies. This finding provided the empirical justification for the intervention proposed in this study: the highly immersive environment created by VR technology can effectively screen out external distractions and redirect attentional focus, lowering the cognitive threshold for entering a meditative state in a way that

conventional meditation practices — constrained by requirements for dedicated spaces, prior training, and sustained habit-formation — often cannot.

Building on this empirical foundation, the Cardia system was designed and developed around three interacting layers: environment, physiology, and subjective experience. On the visual layer, natural landscape scenes rendered in Unreal Engine 5, combined with empirically validated optimal viewing distances (140 to 160 metres) and a low-saturation, cool-toned colour palette, were designed to activate the parasympathetic nervous system's calming response. On the breathing guidance layer, a clinically grounded breathing protocol was paired with particle system animations that visually synchronise with the breath rhythm, guiding participants toward conscious diaphragmatic breathing through non-verbal means. On the physiological feedback layer, real-time PPG signals captured by the BITalino sensor are processed through a Python signal pipeline to compute the RMSSD heart rate variability metric, which is then mapped via OSC protocol to visual parameters in the Niagara particle system, forming a real-time physiological–environmental feedback loop. On the olfactory layer, an Arduino-controlled lavender diffusion module introduces olfactory stimulation at designated points in the experience, extending the system beyond audio-visual feedback into multi-sensory integration.

To test this system, a between-subject controlled experiment was conducted with 27 eligible young adult participants, comparing a pure audio-visual VR condition (Condition A) against a multi-sensory condition augmented with lavender olfactory stimulation (Condition B). A dual-pathway evaluation approach was adopted, combining physiological measures (ΔRMSSD) with multi-dimensional subjective scales (PSS-10, PANAS, SUDS, RSQ), allowing the relaxation outcomes of the two conditions to be compared along both objective and subjective dimensions.

## 5.3. Achievements

Wilcoxon signed-rank test results indicated statistically significant within-group pre-to-post improvements across both conditions on three dimensions — SUDS immediate stress perception, RSQ relaxation state, and PANAS positive affect — with effect sizes reaching large-effect levels ($r \geq 0.85$). These findings support the study's central hypothesis: a VR multi-sensory meditation experience grounded in visual immersion, breathing guidance, and physiological biofeedback is capable of producing meaningful subjective relaxation outcomes

within a single session of approximately five minutes, demonstrating potential as a daily emotional regulation tool.

Although the between-group Mann-Whitney U test did not reach statistical significance at the current sample size, Condition B consistently showed more favourable trends than Condition A on SUDS mean scores (29.2 versus 41.1) and PANAS positive affect (29.5 versus 24.8), suggesting that the addition of an olfactory channel may carry an additive enhancing effect on the overall relaxation experience. Taken together, these results confirm the core feasibility of VR-based multi-sensory meditation as a short-form emotional regulation intervention, while identifying olfactory augmentation as a promising, though not yet conclusively proven, direction for enhancing that intervention further.

## 5.4. Future Work

Several limitations of this study should be acknowledged, and each points towards a concrete direction for future work. First, the effective sample sizes (Condition A: 14 participants. Condition B: 13 participants) were relatively small, resulting in insufficient statistical power for between-group comparisons, future studies should aim for a larger sample ($N \geq 60$) to more conclusively test the incremental benefit of the olfactory condition. Second, baseline PPG signal anomalies in some participants introduced a degree of interference into the reliability of ΔRMSSD calculations, more stable physiological signal acquisition approaches, such as wearable ECG or electrodermal activity sensors, alongside repeated-measures designs, would strengthen future data quality. Third, individual differences in olfactory sensitivity, novelty effects among first-time VR users, and varying degrees of prior familiarity with meditation practice constituted confounding variables that could not be fully controlled within a small-sample framework, future work should systematically control for participants' prior VR experience and meditation background to improve the internal and external validity of experimental conclusions.

Looking further ahead, as VR hardware costs continue to decline and head-mounted displays become lighter and more widely accessible, the large-scale deployment of immersive VR meditation interventions in school mental health services, corporate wellness programmes, and community psychological support settings is becoming increasingly feasible. The multi-sensory interaction design paradigm demonstrated by Cardia — organically integrating visual, auditory, olfactory, and potentially haptic channels into a single psychological intervention — also opens a promising direction for interaction design research within public health more

broadly. The Cardia project suggests that when design practice and psychological science engage in genuine, deep collaboration, they can produce tools that combine experiential quality with clinical intervention potential, offering an initial methodological and practical foundation for future research in multi-sensory VR health intervention.

# Bibliography

1. Ochnik, D., Rogowska, A. M., Kuśnierz, C., Jakubiak, M., Schütz, A., Held, M. J., ... & Cuero-Acosta, Y. A. (2021). Mental health prevalence and predictors among university students in nine countries during the COVID-19 pandemic: a cross-national study. Scientific reports, 11(1), 18644.

2. Cuppen, J., Muja, A., & Geurts, R. (2024). Well-being and mental health among students in European higher education. EUROSTUDENT8 module topic report. https://www. eurostudent. eu/download_files/documents/TM_wellbeing_men talhealth. pdf.

3. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

4. Cuppen, J., Muja, A., & Geurts, R. (2024). Well-being and mental health among students in European higher education. EUROSTUDENT8 module topic report. https://www. eurostudent. eu/download_files/documents/TM_wellbeing_men talhealth. pdf.

5. Cuppen, J., Muja, A., & Geurts, R. (2024). Well-being and mental health among students in European higher education. EUROSTUDENT8 module topic report. https://www. eurostudent. eu/download_files/documents/TM_wellbeing_men talhealth. pdf.

6. Ochnik, D., Rogowska, A. M., Kuśnierz, C., Jakubiak, M., Schütz, A., Held, M. J., ... & Cuero-Acosta, Y. A. (2021). Mental health prevalence and predictors among university students in nine countries during the COVID-19 pandemic: a cross-national study. Scientific reports, 11(1), 18644.

7. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

8. Ochnik, D., Rogowska, A. M., Kuśnierz, C., Jakubiak, M., Schütz, A., Held, M. J., ... & Cuero-Acosta, Y. A. (2021). Mental health prevalence and predictors among university students in nine countries during the COVID-19 pandemic: a cross-national study. Scientific reports, 11(1), 18644.

9. de Lijster, J. M., Dieleman, G. C., Utens, E. M., Dierckx, B., Wierenga, M., Verhulst, F. C., & Legerstee, J. S. (2018). Social and academic functioning in adolescents with anxiety disorders: A systematic review. Journal of affective disorders, 230, 108-117.

10. Cuppen, J., Muja, A., & Geurts, R. (2024). Well-being and mental health among students in European higher education. EUROSTUDENT8 module topic report. https://www. eurostudent. eu/download_files/documents/TM_wellbeing_men talhealth. pdf.

11. de Lijster, J. M., Dieleman, G. C., Utens, E. M., Dierckx, B., Wierenga, M., Verhulst, F. C., & Legerstee, J. S. (2018). Social and academic functioning in adolescents with anxiety disorders: A systematic review. Journal of affective disorders, 230, 108-117.

12. Fergusson, D. M., & Woodward, L. J. (2002). Mental health, educational, and social role outcomes of adolescents with depression. Archives of general psychiatry, 59(3), 225-231.

13. Ochnik, D., Rogowska, A. M., Kuśnierz, C., Jakubiak, M., Schütz, A., Held, M. J., ... & Cuero-Acosta, Y. A. (2021). Mental health prevalence and predictors among university students in nine countries during the COVID-19 pandemic: a cross-national study. Scientific reports, 11(1), 18644.

14. Salari, N., Hosseinian-Far, A., Jalali, R., Vaisi-Raygani, A., Rasoulpoor, S., Mohammadi, M., ... & Khaledi-Paveh, B. (2020). Prevalence of stress, anxiety, depression among the general population during the COVID-19 pandemic: a systematic review and meta-analysis. Globalization and health, 16(1), 57.

15. Cuppen, J., Muja, A., & Geurts, R. (2024). Well-being and mental health among students in European higher education. EUROSTUDENT8 module topic report. https://www. eurostudent. eu/download_files/documents/TM_wellbeing_men talhealth. pdf.

16. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

17. Salari, N., Hosseinian-Far, A., Jalali, R., Vaisi-Raygani, A., Rasoulpoor, S., Mohammadi, M., ... & Khaledi-Paveh, B. (2020). Prevalence of stress, anxiety, depression among the general population during the COVID-19 pandemic: a systematic review and meta-analysis. Globalization and health, 16(1), 57.

18. Ochnik, D., Rogowska, A. M., Kuśnierz, C., Jakubiak, M., Schütz, A., Held, M. J., ... & Cuero-Acosta, Y. A. (2021). Mental health prevalence and predictors among university students in nine countries during the COVID-19 pandemic: a cross-national study. Scientific reports, 11(1), 18644.

19. de Lijster, J. M., Dieleman, G. C., Utens, E. M., Dierckx, B., Wierenga, M., Verhulst, F. C., & Legerstee, J. S. (2018). Social and academic functioning in adolescents with anxiety disorders: A systematic review. Journal of affective disorders, 230, 108-117.

20. Agnafors, S., Barmark, M., & Sydsjö, G. (2021). Mental health and academic performance: a study on selection and causation effects from childhood to early adulthood. Social psychiatry and psychiatric epidemiology, 56(5), 857-866.

21. de Lijster, J. M., Dieleman, G. C., Utens, E. M., Dierckx, B., Wierenga, M., Verhulst, F. C., & Legerstee, J. S. (2018). Social and academic functioning in adolescents with anxiety disorders: A systematic review. Journal of affective disorders, 230, 108-117.

22. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

23. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

24. Varma, P., Junge, M., Meaklim, H., & Jackson, M. L. (2021). Younger people are more vulnerable to stress, anxiety and depression during COVID-19 pandemic: A global cross-sectional survey. Progress in Neuro-Psychopharmacology and Biological Psychiatry, 109, 110236.

25. Agnafors, S., Barmark, M., & Sydsjö, G. (2021). Mental health and academic performance: a study on selection and causation effects from childhood to early adulthood. Social psychiatry and psychiatric epidemiology, 56(5), 857-866.

26. Cipresso, P., Giglioli, I. A. C., Raya, M. A., & Riva, G. (2018). The past, present, and future of virtual and augmented reality research: a network and cluster analysis of the literature.*Frontiers in psychology*, *9*, 2086.

27. R. Masters, V. Interrante, M. Watts, and F. Ortega, ‘‘Virtual nature: Investigating the effect of biomass on immersive virtual reality forest bathing applications for stress reduction,’’ in Proc. ACM Symp. Appl. Perception, New York, NY, USA, Sep. 2022, pp. 1–10, doi: 10.1145/3548814.3551459.

28. Saredakis, D., Szpak, A., Birckhead, B., Keage, H. A. D., Rizzo, A., & Loetscher, T. (2020). Factors associated with virtual reality sickness in head-mounted displays: A systematic review and meta-analysis. Frontiers in Human Neuroscience, 14, 96. https://doi.org/10.3389/fnhum.2020.00096

29. Cipresso, P., Giglioli, I. A. C., Raya, M. A., & Riva, G. (2018). The past, present, and future of virtual and augmented reality research: A network and cluster analysis of the literature. Frontiers in Psychology, 9, 2086. https://doi.org/10.3389/fpsyg.2018.02086

30. Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. JMIR Mental Health, 4(2), e7785. https://doi.org/10.2196/mental.7785

31. Nichol, B., Sunner, A., Wilson, N., & Hewitt, J. (2024). Exploring the impact of socially assistive robots on health and wellbeing across the lifespan: An umbrella review and meta-analysis. International Journal of Nursing Studies, 154, 104746. https://doi.org/10.1016/j.ijnurstu.2024.104746

32. Papadopoulos, C., Castro, N., Nigath, A., Davidson, R., Faulkes, N., Menicatti, R., Khaliq, A. A., Recchiuto, C., Battistuzzi, L., Randhawa, G., & Sgorbissa, A. (2022). The CARESSES randomised controlled trial: Exploring the health-related impact of culturally competent artificial intelligence embedded into socially assistive robots and tested in older adult care homes. International Journal of Social Robotics, 14, 245–256. https://doi.org/10.1007/s12369-021-00781-x

33. Sung, H., Kim, M., Park, J., Shin, N., & Han, Y. (2024). Effectiveness of virtual reality in healthcare education: Systematic review and meta-analysis. Sustainability, 16(19), 8520.

34. Angra, S., Jangra, S., Gulzar, Y., Sharma, B., Singh, G., & Onn, C. W. (2025). Twenty-two years of advancements in augmented and virtual reality: A bibliometric and systematic review. Frontiers in Computer Science, 7, 1470038.

35. Hendric, S. W. H. L., & Rosmansyah, Y. (2025). Exploring immersive virtual reality in higher education: Research gap and future direction—A scoping review. IEEE Access, 13, 76308-76321.

36. Hendric, S. W. H. L., & Rosmansyah, Y. (2025). Exploring immersive virtual reality in higher education: Research gap and future direction—A scoping review. IEEE Access, 13, 76308-76321.

37. Rubio-Tamayo, J. L., Gertrudix Barrio, M., & García García, F. (2017). Immersive environments and virtual reality: Systematic review and advances in communication, interaction and simulation. Multimodal technologies and interaction, 1(4), 21.

38. Angra, S., Jangra, S., Gulzar, Y., Sharma, B., Singh, G., & Onn, C. W. (2025). Twenty-two years of advancements in augmented and virtual reality: A bibliometric and systematic review. Frontiers in Computer Science, 7, 1470038.

39. Anthes, C., García-Hernández, R. J., Wiedemann, M., & Kranzlmüller, D. (2016, March). State of the art of virtual reality technology. In 2016 IEEE aerospace conference (pp. 1-19). IEEE.

40. Anthes, C., García-Hernández, R. J., Wiedemann, M., & Kranzlmüller, D. (2016, March). State of the art of virtual reality technology. In 2016 IEEE aerospace conference (pp. 1-19). IEEE.

41. Cipresso, P., Giglioli, I. A. C., Raya, M. A., & Riva, G. (2018). The past, present, and future of virtual and augmented reality research: A network and cluster analysis of the literature. Frontiers in Psychology, 9, 2086. https://doi.org/10.3389/fpsyg.2018.02086

42. de Freitas, F. V., Gomes, M. V. M., & Winkler, I. (2022). Benefits and challenges of virtual-reality-based industrial usability testing and design reviews: A patents landscape and literature review. Applied Sciences, 12(3), 1755.

43. Hendric, S. W. H. L., & Rosmansyah, Y. (2025). Exploring immersive virtual reality in higher education: Research gap and future direction—A scoping review. IEEE Access, 13, 76308-76321.

44. Shang, W., & Alena, K. (2025, May). Advancements in virtual reality technology: A systematic review of user experience and application trends. In Artificial Intelligence and Applications (Vol. 3, No. 4, pp. 341-358).

45. Berrezueta-Guzman, S., & Wagner, S. (2025). Immersive multiplayer vr: Unreal engine's strengths, limitations, and future prospects. IEEE Access.

46. Zhang, Z., et al. (2025, August 18). Choosing the right engine in the virtual reality landscape. arXiv. https://arxiv.org/abs/2508.13116

47. Game-Ace. (2025, November 21). How to make a multiplayer game in Unreal Engine. https://game-ace.com/blog/multiplayer-game-in-unreal/

48. Chen, Z., et al. (2025). Blue Book of Mental Health: Report on the Development of Mental Health in China (2023–2024). Institute of Psychology, Chinese Academy of Sciences.

49. Survey Report on Mental Health of Chinese Society (2025). Retrieved from https://health.baidu.com

50. International Healing Industry Association. (2025). In-depth analysis of the mental health and healing industry landscape. Retrieved from https://www.ihiac.org.cn

51. Shu, Y., Wu, C., & Zhai, Y. (2022). Impacts of landscape type, viewing distance, and permeability on anxiety, depression, and stress. International Journal of Environmental Research and Public Health, 19(16), 9867. https://doi.org/10.3390/ijerph19169867

52. Vaquero-Blasco, M. A., Perez-Valero, E., Lopez-Gordo, M. A., & Morillas, C. (2020). Virtual reality as a portable alternative to chromotherapy rooms for stress relief: A preliminary study. Sensors, 20(21), 6211. https://doi.org/10.3390/s20216211

53. Wang, X., Shi, Y., Zhang, B., & Chiang, Y. (2019). The influence of forest resting environments on stress using virtual reality. International Journal of Environmental Research and Public Health, 16(18), 3263. https://doi.org/10.3390/ijerph16183263

54. Vaquero-Blasco, M. A., Perez-Valero, E., Morillas, C., & Lopez-Gordo, M. A. (2021). Virtual reality customized 360-degree experiences for stress relief. Sensors, 21(6), 2219. https://doi.org/10.3390/s21062219

55. Vaquero-Blasco, M. A., Perez-Valero, E., Morillas, C., & Lopez-Gordo, M. A. (2021). Virtual reality customized 360-degree experiences for stress relief. Sensors, 21(6), 2219. https://doi.org/10.3390/s21062219

56. Shu, Y., Wu, C., & Zhai, Y. (2022). Impacts of landscape type, viewing distance, and permeability on anxiety, depression, and stress. International Journal of Environmental Research and Public Health, 19(16), 9867. https://doi.org/10.3390/ijerph19169867

57. Ma, J., Li, H., Zhao, M., & Williams, J. (2025). Mind nature: A virtual reality nature therapy for reducing anxiety and negative mood in university students. International Journal of Clinical and Health Psychology, 25(4), 100636. https://doi.org/10.1016/j.ijchp.2025.100636

58. Munoz, M. L., van Roon, A., Riese, H., Thio, C., Oostenbroek, E., Westrik, I., de Geus, E. J. C., Gansevoort, R., Lefrandt, J., Nolte, I. M., & Snieder, H. (2015). Validity of (ultra-)short recordings for heart rate variability measurements. PLOS ONE, 10(9), e0138921. https://doi.org/10.1371/journal.pone.0138921

59. Shaffer, F., & Ginsberg, J. P. (2017). An overview of heart rate variability metrics and norms. Frontiers in Public Health, 5, 258. https://doi.org/10.3389/fpubh.2017.00258

60. Munoz, M. L., van Roon, A., Riese, H., Thio, C., Oostenbroek, E., Westrik, I., de Geus, E. J. C., Gansevoort, R., Lefrandt, J., Nolte, I. M., & Snieder, H. (2015). Validity of (ultra-)short recordings for heart rate variability measurements. PLOS ONE, 10(9), e0138921. https://doi.org/10.1371/journal.pone.0138921

61. Cohen, S., Kamarck, T., & Mermelstein, R. (1983). A global measure of perceived stress. Journal of Health and Social Behavior, 24(4), 385–396. https://doi.org/10.2307/2136404

62. Cohen, S., & Williamson, G. M. (1988). Perceived stress in a probability sample of the United States. In S. Spacapan & S. Oskamp (Eds.), The social psychology of health (pp. 31–67). Sage Publications.

63. Cohen, S., & Williamson, G. M. (1988). Perceived stress in a probability sample of the United States. In S. Spacapan & S. Oskamp (Eds.), The social psychology of health (pp. 31–67). Sage Publications.

64. Krampen, G. (1987). Konzept und Methodik der Entspannungsforschung [Concept and methodology of relaxation research]. Zeitschrift für Klinische Psychologie, 16(1), 1–21.