

INTELLIGENT AMBIANCE

Digitally mediated workspace atmosphere, augmenting experiences and supporting wellbeing


ASAPH AZARIA

B.Sc. in Mathematics and Physics, Hebrew University of Jerusalem, 2005

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of Master of Science in Media Arts and Sciences at the Massachusetts Institute of Technology.

SEPTEMBER 2016

© Massachusetts Institute of Technology 2016. All rights reserved.



Author:

Signature redacted

Asaph Azaria
Program in Media Arts and Sciences
August 19, 2016

Certified by:

Signature redacted

Joseph A. Paradiso
Professor in Media Arts and Sciences
Thesis Supervisor

Accepted by:

Signature redacted

Pattie Maes
Academic Head
Program in Media Arts and Sciences

The following people served as readers for this thesis:

Thesis Reader:
Pattie Maes
Professor in Media Arts & Sciences
MIT Media Lab

Signature redacted

Thesis Reader:
Rosalind W. Picard
Professor in Media Arts & Sciences
MIT Media Lab

Signature redacted

INTELLIGENT AMBIANCE

Digitally mediated workspace atmosphere, augmenting experiences and supporting wellbeing

ASAPH AZARIA

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, on August 19, 2016 in partial fulfillment of the requirements for the degree of Master of Science in Media Arts and Sciences at the Massachusetts Institute of Technology.

ABSTRACT

Cues from the physical environment are constantly sensed and interpreted – unconsciously finding their way into our cognitive schemas and influencing our perceptions and experiences. Manipulating them has been shown to be powerful, affecting cognitive performance, mood and even physiology. Inspired by this, we propose a workspace capable of dynamically transforming its ambiance.

This work presents the Digital Cubicle – a cubicle workspace which uses lighting, video projection and sound to manipulate its physical characteristics. A set of software tools is developed to create digital compositions in such a workspace, basing the compositions on objective physical measures from existing popular workspaces. Through a set of user studies, we evaluate how these elicit occupant perceptual, cognitive and physiological responses. Detailed in this work, are physiological responses indicating stress development and restoration, with interesting implications for health and wellbeing applications.

We develop a sensor data collection infrastructure to complement the workspace's controllable ambiance, with rich, real-time information about the occupant context and state. Leveraging Reinforcement Learning techniques, we present a framework to devise adaptive control agents. These utilise sensor data to recommend and transform the workspace ambiance in a closed-loop fashion. A prototype intelligent agent is implemented, optimising for occupants' heart rate recovery, yet counterbalancing for occupants preferences and requests. Through evaluative simulations, anchored with real occupant data, we demonstrate and discuss the effectiveness of our proposed approach.

ACKNOWLEDGMENTS

First and foremost, I wish to thank my advisor, Professor Joseph A. Paradiso. His invaluable advice glued together fragments of ideas into one cohesive concept and a clear path for investigation. I am grateful for his trust taking me into the Responsive Environments family, and value the moments I had learning from his never-ending knowledge and experience. Likewise, I thank my readers Professor Pattie Mass and Professor Rosalind Picard for their constructive critique and mentorship in this work.

A special acknowledgment goes to Nan Zhao. Without her guidance and support much of this work would not have been possible. Her perspective and creativity have been a source of influence and inspiration. I could not have wished for a better partner in collaboration.

I additionally thank the Isabella Stewart Garden Museum, who kindly opened up their gates and generously allowed me to record and explore their inspiring spaces; Gershon Dublon, Spencer Russell and Brian Mayton, who provided feedback and insights and were always ready to lend their capable shoulders; The remainder of the Responsive Environments group for challenging me to think differently and their constant encouragement; Judith Amores, for our mutual work exploring memory and perception and experimenting with her intriguing work in olfactory interfaces; Amna Carreiro, Linda Peterson and Keira Horowitz for always being there for everyone, and smoothly making things happen.

Last but not least, I thank my own parents and family, who raised and supported me, pushing me to become who I want to be.

CONTENTS

1	INTRODUCTION	17
1.1	Overview	18
1.2	Contributions	19
1.3	Supporting Work	20
2	RELATED WORK	21
2.1	Affective and Cognitive effects of the physical environment	21
2.2	The impact of nature experience on stress and attention restoration	22
2.3	Ambient Display of Information	23
2.4	Intelligent Adaptive Control Agents	24
3	SYSTEM DESCRIPTION	27
3.1	Design Considerations	27
3.1.1	Digital control	27
3.1.2	Real-time response	28
3.1.3	Extensibility	28
3.1.4	Applicability	29
3.1.5	Composition	29
3.2	Capturing ambiance of real working environments	30
3.2.1	Describing lighting conditions	30
3.2.2	Video recording	31
3.2.3	Audio recording	32
3.3	Physical Characteristics Manipulation	32
3.3.1	Lighting	34
3.3.2	Video projection	37
3.3.3	Sound	41
3.3.4	Occupant interface	42
3.4	Real-time Sensor Data Collection	42
4	PHYSIOLOGICAL EFFECTS	47
4.1	Experimental Design	47
4.1.1	Study Protocol	47
4.1.2	Scenes Selection	50
4.1.3	Collected Data	51
4.2	Data Processing and Characterisation	53
4.2.1	Signal Preprocessing	55
4.2.2	Law of Initial Values (LIV)	57
4.2.3	Building Probabilistic Models	58
4.2.4	Stress Development and Restoration Metrics	60
4.3	Results and Discussion	61
4.3.1	Were stress development and restoration elicited?	61
4.3.2	Individual Physiological Responses	64
4.3.3	Physiological Responses Generalised per Scene	72

4.3.4	Focusing capabilities	76	
4.3.5	Subjective Perception of the Scene	78	
4.4	Summary	81	
5	INTELLIGENT CONTROL AGENT	83	
5.1	Problem Statement	83	
5.2	Design Principles	84	
5.3	Applying a Reinforcement Learning Framework	86	
5.4	Modelling Actions and States	87	
5.5	Designing the Rewards Function	90	
5.5.1	Occupant Interactions	91	
5.5.2	Physiology-based Rewards	92	
5.6	Solving with Q-Learning	93	
5.6.1	The Learning Step	93	
5.6.2	Choosing an Action	94	
5.7	Implementation	94	
5.8	Evaluation Method	96	
5.8.1	Implementation	96	
5.9	Results and Discussion	100	
5.9.1	Heuristic 1: Learning from occupant preferences	101	
5.9.2	Heuristic 2: Learning an optimal heart rate recovery sequence	102	
5.9.3	Heuristic 3: Combining focus and stress restoration	104	
5.10	The challenge of convergence time	105	
5.10.1	Reducing state-action space and parameter tuning	106	
5.10.2	Boost learning with simulation	107	
6	CONCLUSIONS AND NEXT STEPS	109	
A	APPENDIX: PHYSIOLOGICAL STUDY MATERIALS	113	
A.1	Cognitive Task	113	
A.1.1	Text (presented in the left of the screen)	113	
A.1.2	Questions (presented in the right of the screen)	113	
A.2	Survey	115	
	BIBLIOGRAPHY	117	

LIST OF FIGURES

- Figure 1 The recording setup, including a camera and a tetrahedral microphone, in the Isabella Stewart Gardner museum 31
- Figure 2 Top view of the measurement tools used to describe the lighting conditions. From the left: Colour reference card, rod for shadowing pattern and a LuxMeter. 32
- Figure 3 Manipulating physical characteristics in the Digital Cubicle. 33
- Figure 4 Physical layout of the Digital Cubicle. 34
- Figure 5 Top-level software design of the lighting control 35
- Figure 6 User interface of the lighting mapping controller. The yellow circles are virtual light sources within the room's two dimensional representation. Lightness, colour temperature and decay parameters for each are set on the right. 37
- Figure 7 The display with and without active projection 38
- Figure 8 Max Msp Controller developed for video composition 40
- Figure 9 Max Msp Controller developed for audio speaker configuration 42
- Figure 10 Top-level design of the *Sensor Collection Service* 43
- Figure 11 Occupant wearing the set of supported sensors in the Digital Cubicle 45
- Figure 12 Facial feature points, head and pupil orientation extracted by the Intraface code library. 46
- Figure 13 Schematic illustration of the study's protocol. Each participant was invited to a single sitting consisting of six identical sessions. Sessions were structured as illustrated above. 47
- Figure 14 Necker Cube Pattern Control Test. Participants were presented with the leftmost image. The following images illustrate the two possible perceived orientations. 49
- Figure 15 Images of the six scenes used in the physiological user study. From left to right and top to bottom: Neutral, Forest, Rotch Library, Kites, Shibuja and Sunset. 51

- Figure 16 Sample heart rate, respiration rate and HRV signals, collected from Subject 6 when he was exposed to the Shibuja scene. 52
- Figure 17 Sample EDA signal, collected from Subject 7 when he was exposed to the Neutral scene. 53
- Figure 18 Signal preprocessing example illustrating Subject 4's heart rate exposed to Shibuja. From the top: Original signal, Moving average filtering, Standardisation and Delta scores. 56
- Figure 19 Application of the Law of Initial Values to Subject 1's heart rate changes. The delta scores (Δz_t) are plotted against their preceding levels ($z_{t-\kappa}$). The scores are corrected to become $\Delta z'_t$, which is indicated by the dashed line 57
- Figure 20 Heart rate change pdfs estimated for Subject 1 in the Neutral and Forest scenes. pdfs on the left and right based on test and break segment data respectively. The markers below indicate $\Delta z'$ samples from which the pdfs were estimated. 60
- Figure 21 Histograms comparing stress development and restoration between the test and break conditions. The top row illustrates Z-score levels for each of the physiological signals. The bottom row illustrates corrected delta scores ($\Delta z'$), representing physiological changes. 63
- Figure 22 MLE of the intensity of heart rate change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break) 65
- Figure 23 MLE of the intensity of respiration rate change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break) 66
- Figure 24 MLE of the intensity of HRV change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break) 67
- Figure 25 Probability of heart rate change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense heart rate change (>0.5 std). Asterisks indicate statical significance. 68

Figure 26	Probability of respiration rate change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense respiration rate change (>0.5 std). Asterisks indicate statistical significance.	69
Figure 27	Probability of HRV change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense HRV change (>0.5 std). Asterisks indicate statistical significance.	70
Figure 28	Probability density function of heart rate change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.	73
Figure 29	Probability density function of respiration rate change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.	74
Figure 30	Probability density function of HRV change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.	75
Figure 31	Necker Z-score per participant, for each of the rendered scenes. Rightmost bars are the average values from all the participants.	77
Figure 32	Self-reported restoration and focusing ratings for the rendered scenes	79
Figure 33	States and actions model for the intelligent agent	89
Figure 34	Implicit reward function based on occupant interactions	91
Figure 35	Top-level design of the Intelligent control agent implementation	95
Figure 36	Top-level design of the simulation implementation	97
Figure 37	Sequence of contexts for a simulated session	98
Figure 38	Rewards progression throughout simulations of Heuristic-1	101
Figure 39	Rewards progression throughout simulations of Heuristic-2 for each of the study participants	104

Figure 40	Rewards progression throughout simulations of Heuristic-3 for each of the study participants	106
-----------	--	-----

LIST OF TABLES

Table 1	Significant correlations ($p < 0.05$) found between perceived stress restoration and physiological recovery during breaks.	80
Table 2	Simulation results for Heuristic 2. For each participant the table summarises the optimal policy the agent converged to, i.e. the optimal scene for each working context. λ is the ϵ -greedy decay parameter of our exploration strategy.	103
Table 3	Simulation results for Heuristic 3. For each participant the table summarises the optimal policy the agent converged to, i.e. the optimal scene for each working context. λ is the ϵ -greedy decay parameter of our exploration strategy.	105

INTRODUCTION

Our senses constantly perceive and interpret information about our physical environment. Many times unconsciously, textures, colours, shapes, sounds and lighting conditions find their way into our cognitive schemas and serve as cues to retrieve skills, knowledge, feelings and behaviours. This phenomenon has sparked researchers' interest particularly for its implications on cognitive performance, health and wellbeing. Indeed, manipulating the physical environment was demonstrated to be powerful in supporting memory, fostering creativity, enhancing sensitivity to details, and balancing cognitive load [13]. It was additionally shown to manage important aspects of psychological affect. Art and architecture are prominent such examples for moulding mood, attitude, narratives, experiences via physical manipulation [53] [5] [46].

In a complex interplay, the human physiology also responds to its physical environment. Light and temperature modulate pineal melatonin production, which maintains the body's circadian rhythm and in turn affects its resilience to various disease [25]. Noise commonly triggers excitation, which for prolonged subjection is highly associated with chronic-stress [37]. Nature experiences, on the other hand, exhibit healing qualities, where even minimal exposure to plants was demonstrated to facilitate stress regulation and physical recovery [27].

Concurrently, the progress towards truly ubiquitous computing has made information readily accessible everywhere. Employees of the 21st century are increasingly interacting, engaging and completing their assignments in dynamic locations [6]. Likewise, student instruction is provided not only in classrooms but also through mobile contexts, and even virtual worlds [70]. Advances in sensing technologies and affect recognition have been recently introduced, improving related computer interactions even further. Tools to capture, process and retrieve information are literally at the tips of one's fingers. Too very often, however, we find ourselves in pursuit not of such tools, but of the appropriate atmosphere in which to use them.

If the physical environment is indeed so salient, why is it then that our workspaces cannot dynamically change their physical characteristics to support our activities? In an age that strides towards personalisation, why can't individuals capture and access working ambience, the same way they do with information? Shouldn't the way we experience the environment be an integral part of our work?

We envision an intelligent workspace that is capable of dynamically transforming its ambience to manipulate one's experiences and

perception. A workspace that, when asked, can replicate the tension of a control room or the restorative qualities of a loved childhood place. A workspace that can instantly trade the engaging diligence of a library with the liberating sensation from a wilderness stroll. We ask that it would be attentive, leveraging sensing and pattern recognition technologies; continuously monitoring occupants' physiology and behaviour to infer context and affect. Through data we foresee such workspace gaining intelligence, optimising our environment to either enhance performance or support our health and wellbeing. A workspace that suggests the atmosphere of a local coffee shop when it is time to spark creativity, and the atmosphere of a study when it is time to maintain focus. A workspace that integrates nature escapes into our workday, balancing our schedule and physiological needs to create a sustainable, healthier routine.

1.1 OVERVIEW

Our work is structured as follows: We begin by summarising related academic work, on which we either base our assumptions or from which we draw inspiration in our designs (Chapter 2). Next, we develop the Digital Cubicle – a prototype workspace that using digital media, manipulates many of its physical characteristics. A set of tools are developed to control the cubicle's smart illumination, video projection and sound. We experiment utilising those to artificially mimic some perceptual qualities of exiting, real, working environments, basing our composition on objective measurements of physical quantities. We additionally equip the Digital Cubicle with a modular real-time data collection infrastructure. Integrating a set of commercial sensors, we demonstrate the rich information set that can be obtained, and the infrastructure's flexibility (Chapter 3).

Through three user studies, we set out to assess our prototype's effects on perception, cognitive performance and physiology. In this essay, we concentrate our review a user study, in which we examine occupants' physiological responses to various compositions rendered in the Digital Cubicle. Measuring heart rate, respiration rate and heart rate variability, we develop a signal processing approach and corresponding metrics to assess stress development and restoration. These responses, though individual, are found significantly different, during both mentally demanding and break conditions. These findings are presented and discussed in detail (Chapter 4).

Finally, we propose and evaluate a Reinforcement Learning approach to build intelligent control agents on top of the Digital Cubicle. We implement a prototype agent which leverages collected sensor data, autonomously rendering compositions in the Digital Cubicle in such way. The agent optimises for occupant physiological stress recovery, yet counterbalances for their preferences and requests.

We present evaluative simulations, anchored with real occupant data, suggesting significant possible improvements in occupants' heart rate recovery. We conclude by discussing the promise and peril of our proposed approach, and outline steps for future investigation. (Chapter 5)

1.2 CONTRIBUTIONS

This work takes the first steps towards the vision outlined above. We summarise our contributions as follows:

- To the best of our knowledge, we are the first to present a workspace prototype capable of digitally mediating its ambiance. We build our workspace as an extensible, multimodal platform, supporting future research and exploration. A set of tools are developed to compose lighting, video projection and sound, experimenting with utilising objective physical measurements from existing workspaces as the basis to compose artificial ambiance. Likewise, a flexible sensor data collection infrastructure is developed. We demonstrate its strength – fusing data streams in real-time – using a set of different sensors from multiple vendors.
- We test the hypothesis that such a workspace does have significant effects on occupant physiology through a user study. We provide a framework to process occupant physiological signals and quantify stress development and restoration in multiple contexts. We show that physiological effects are highly personal and partially correlate with occupants' self-reported subjective perception. Likewise, we present the workspace's effects on occupants' ability to direct attention, through both subjective and objective measures.
- Finally, we present a framework to design a closed-loop, intelligent control agents for our platform – bridging sensor data and physical manipulation. We formulate the problem as a Markov Decision Process and borrow tools from Reinforcement Learning literature to build an adaptive control scheme. We implement a working agent, learning from occupant interactions and real-time heart rate recovery measures, to find an optimal balance between preference and physiological benefits. Through simulations, we demonstrate the potential effectiveness of our approach in a new practical domain.

1.3 SUPPORTING WORK

For brevity, our user studies examining the perceptual and cognitive effects of the Digital Cubicle are not fully described in the scope of this essay; they will be appropriately detailed in future publications. We briefly describe them here to provide the reader with relevant context.

The first user study collected and analysed participants' subjective perception of different compositions rendered in the Digital Cubicle. We compared those to differences we found in subjective perception collected from the real working environments on which we base the artificial compositions. Covering broad perceptual categories, including affordance, affect, spatial impressions, immersion, and presence, we established two main statistically significant perceptual differences among the artificial compositions. We generalise those to be the composition's suitability to direct focus (affording conditions for detailed work, not demanding or distracting, cohesive, etc.) and its restorative qualities (as defined by Attention and Stress Restoration Theories - See Section 2.2).

In the second user study, we followed the Encoding Specificity Principle (See Section 2.1), and compared cognitive performance in a memorisation task between an intervention and control groups. Both groups were asked to memorise a set of unknown Hindi words in a real working environment and recall them the next day. Whereas the control group recalled the words in a regular cubicle office, the intervention group recalled them in the Digital Cubicle. During recall, the Digital Cubical was mimicking physical features from the original environment in which the words were learned. On average, we observed free recall improvements in the intervention group. However, due to major individual differences and the relatively small sample size of our study, we were unable to determine statistical significance.

RELATED WORK

2.1 AFFECTIVE AND COGNITIVE EFFECTS OF THE PHYSICAL ENVIRONMENT

Evidence, accumulated for over half a century, indicates the significance of the physical environment in influencing one's cognitive performance and affective state. Research spans areas such as environmental psychology, urban planning, educational instruction design, cognitive functioning and mental health. A recent revision of the famous theoretical framework of cognitive load has even reconceptualised the physical environment as one of its main three causal factors [13]. It argues that the physical characteristics of the space inherently interact with a learner's task and state, to determine the occupancy of the working memory. In turn, it affects not only current processing and performance, but also construction and retrieval of long-term cognitive schemas, governing the development of personality traits, dispositions, knowledge and skills.

The Encoding Specificity Principle, articulated by memory researchers Thomson and Tulving [83], is a famous example of such an interaction. The principle suggests that memory recall is most effective when the physical conditions at the time of encoding match the physical conditions at the time of retrieval. Indeed, superior memory performance has been consistently observed when the physical learning and test environments are similar [78]. Analogous results have been observed for motor skills development [51]. A well-known empirical example is the home-field advantage, in which athletes usually have better chances winning a game when it is played on the team's own practice field [74]. Paas et al. extended this idea, and additionally exemplified that if the learning context cannot mimic the test one, learning best take place in a variety of contexts. This way transfer to new unfamiliar contexts is facilitated [63].

Researchers have also studied the effects of specific physical characteristics on learning. Air quality and thermal conditions were observed to affect individual's learning performance; presumably through oxygen-related physiological mechanisms [50]. McCoy and Evans identified several environmental characteristics that facilitate creative performance – complexity of visual details, the use of natural materials, and a limited use of cool colours [55]. Levels of noise, height of the ceiling and lighting colours were shown to act differently on different types of cognitive capabilities [58]. Zhu and colleagues, for example, demonstrated that the colour blue enhances cognitive performance

on creative tasks whereas red improves detail-oriented tasks performance [56]. Mehta et al. found that moderate levels of noise (e.g. the ambient sound of a coffee shop) facilitates abstract processing [57].

Particularly interesting is the interplay between the physical environment and the learner's affect. Nizam and Marnie [43] found particular colour lighting conditions in office buildings to affect occupants' aspects such as calmness and comfort. Lighting intensity was interestingly correlated with the intensity of induced emotions [95]. Hall discusses some affective influences of the space design [22], showing how open-space designs produce social cues, which can be interpreted as an infringement on personal space or feelings of crowdedness. Emotions, mood and motivation additionally mediate the relationship between the physical environment and cognitive performances [42]. Thus, a consistent relationship was identified between the preferred learning environment of a learner and his or her achievements in it [31] [93]. In a similar vein, aesthetically appealing multimedia learning-materials were shown to promote positive emotions and reduce the perceived difficulty of the learning task [88]. Evans and Stecker observed that noise can lead to diminished motivation, feelings of helplessness, and consequently result in lower learning outcomes [23].

2.2 THE IMPACT OF NATURE EXPERIENCE ON STRESS AND ATTENTION RESTORATION

A large body of literature is dedicated to the role of nature in providing feelings of wellbeing and supporting physical health and mental functioning [10]. Attention restoration theory (ART) [45] and Stress reduction theory (SRT) [87] are the two major explanatory theories for the restorative qualities of nature. Both draw heavily on human evolution. SRT posits a healing power to natural places – watersides and visible horizons – in which our species had greater rates of survival [84] [85]. These moderate stress and negative valence through autonomic psychophysiological pathways, most noticeably in individuals who have been stressed beforehand [86] [30]. Affective responses to such natural settings are theorised to be preconscious, affecting even individuals who think they have habituated themselves to nature deprived settings [85]. Ulrich et al. compared nature and urban settings, monitoring subjects heart rate, skin conduction muscle tension and systolic blood pressure. All measures have shown significantly higher recovery speed viewing natural rather than urban scenes [87].

Kaplan and Kaplan [44] formulated ART, centring on nature's power to replenish mechanisms of attention. It claims that urban life taxes attentional capacities more consistently than natural environments, as the latter were more common in our species' collective past. Tenessen and Cimprich showed increased capacity to direct attention in stu-

dents who had access to natural views through their dormitories windows [80]. Berman et al. compared break walks in urban settings to an arboretum [7]. They established participants increased positive affect and attentional performance in the latter. A famous study by Taylor et al. investigated children's views from home in a housing complex in Chicago [79]. On average, children with better access to natural views, exhibited higher levels of concentration, impulse inhibition and delay of gratification; All shown to act as mediators for lower levels of aggression and higher scholastic and career success.

The Kaplans postulated that there are five essential components to an environment's restorative effects. Accordingly, the Perceived Restorativeness Scale (PRS) [29] – a frequently reported psychometric scale assessing these components – was developed. *Being away* is the feeling of an escape from the habitual activities and concerns for daily life. *Fascination* is the environment's capacity to effortlessly capture one's attention. The *Coherence* and *Scope* perceived in an environment, also affect its restorative qualities and include the possibility of feeling immersed in it. Finally, *Compatibility* describes the match between an individual's intentions, inclinations or purposes and the environment.

We note some previous work indicating that the aforementioned physiological effects diminish when the environment is digitally mediated. Kahn et al. compared heart rate recovery from low-level stress, exposing participants to one of three conditions – a glass window, an artificial plasma "window" (both affording similar natural views) and a blank wall [41]. As expected, the glass window was significantly more restorative than the blank wall and participants' heart rate tended to decrease more rapidly the more time spent looking at it. This was not the case with the plasma window, which was found no more restorative than the blank wall. Kort et al. suggested that the effects of surrogate nature, mediated through technology, depend on immersion [18]. They showed an interaction between screen size, self reported presence and stress recovery measured through heart rate and skin conductance.

2.3 AMBIENT DISPLAY OF INFORMATION

Ambiance or atmospherics design has been extensively researched in both architecture and psychology. The idea of digitally manipulating it, however, was introduced by HCI (Human Computer Interface) literature. Skog et al. discuss design principles for ambient information visualisation in public spaces – attempting to strike a balance between aesthetic appeal, usefulness and relevance [77]. Ishii et al. envisioned that the physical architectural space would be a new form of interface between humans and digital information [35]. In their work "ambientRoom", they exemplified how information can be dis-

played and processed in the background of attention [36]. Rhodes and Maes also emphasised the need for an accessible yet non-intrusive display of information [69]. They introduced information agents, which proactively retrieve and display information based on a person's local context, as it is inferred by a set of sensors. A similar approach, was presented in "Ambient Agoras", where user interfaces and ambient displays were embedded in the architectural envelope to foster information flow through an organisation (e.g. company news and announcements) [67].

2.4 INTELLIGENT ADAPTIVE CONTROL AGENTS

The area of ubiquitous computing envisions people surrounded by intelligent intuitive interfaces embedded in various objects. This has given rise to the idea of Ambient Intelligent (or Smart Environments), as articulated by the IST Advisory group – "environments that are capable of recognising and responding to the presence of different individuals in a seamless, unobtrusive and often invisible way" [20]. In this technological paradigm, environments are equipped with complementary technologies for intelligent control, namely sensing, reasoning and acting [16] [3]. Sensing technologies continuously gain knowledge about the preferences, intentions, needs and habits of their occupants. Reasoning technologies leverage the sensed data, deciding how to act upon the environment to achieve an intended goal. Finally, acting technologies carry out these decisions (e.g. activating actuators or suggesting interfaces), only to provide feedback again to the sensing technologies.

As occupant interactions, preferences, intentions and needs may change over time, adapting is key for such control processes. Planning and Machine Learning techniques have been utilised to recognise and learn patterns in occupant behaviour. Mozer et al. introduced an intelligent lighting system, based on Reinforcement Learning techniques, to minimise energy consumption [59]. The system would learn from penalties, whenever inhabitants revoked its decisions, to derive a control policy balancing energy and comfort. Project iDorm [33] demonstrated learning from multiple sensing and acting technologies, using an unsupervised fuzzy technique. It monitors seven sensor inputs in a room (chair and bed pressure, light levels and temperature) to adaptively control ten actuators including lamps, window blinds, a heater and the media playlist. Aztiria et al. survey and suggest a framework to compare Machine Learning techniques in the context of Ambient Intelligence [3].

Throughout our work, we additionally notice an interesting analogy between our adaptive control scheme and affective music playlist generation. Music is intrinsically intertwined in our every day life, setting an atmospheric tone which affects our experiences. It was

demonstrated to evoke emotions and mood and to stimulate a variety of physiological responses [39]. The large amount of digital music, hence, opens up opportunities for music players with affective intelligence. Picard was the first to describe it in her seminal work on affective computing [65]. Challengingly, music listening experiences are highly personal, contextual and sequential, providing fertile ground for academic exploration. Bonnin and Jonnach survey algorithms and learning techniques applied for automated generation of music playlists [9]. Of particular interest is the work of Oliver et al. who incorporated contextual sensor information to enhance exercise performance through playlist generation [61]. Likewise, Van der Zwaag et al. demonstrated an affective music player, capable of directing listeners' Electrodermal Activity and Skin Temperature [38] [99].

SYSTEM DESCRIPTION

To explore our vision, we set out to embody it in a working prototype. We approach this embodiment in a three-phase process. First, we construct a workspace capable of manipulating many of its physical characteristics. Then, we build a sensor infrastructure, collecting rich real-time data about the occupant. Finally, we explore designing an intelligent control agent, bridging collected sensor data with the workspace’s capabilities to manipulate its appearance.

When we embarked on our journey, designing an intelligent agent was an abstract problem, open to much speculation. We therefore decided to regard the first two phases as building a platform – flexible enough to add or remove as necessary either sensing or appearance manipulation components. Such a platform will allow iterative experimentation, anchoring the agent’s design with invaluable real occupant data. The design and implementation of this platform is described throughout this chapter. We call our platform the Digital Cubicle (DC). Chapter 5 will describe the implementation of the intelligent control agent.

This chapter is structured as follows: Section 3.1 outlines the design considerations, which guide our prototype implementation. In Section 3.2 we describe how we recorded real working environments to serve both as content to display in the Digital Cubicle and as reference for evaluation. Sections 3.3 and 3.4 detail the technical implementation of our platform.

3.1 DESIGN CONSIDERATIONS

In this section we describe the design considerations, which guided the Digital Cubicle implementation. These were formed mainly by our aforementioned vision, but also by occupant feedback we obtained through multiple short design iterations. They are of course only one possible interpretation of our vision. We describe them here, however, to shed light on tradeoffs and choices we had made in the process of building our prototype.

3.1.1 *Digital control*

We require that all elements, capable of manipulating the workspace’s physical characteristics, be fully controlled digitally. Our prototype, therefore, primarily uses digital media (images, videos and sound), which can be fully controlled with software, to manipulate appear-

ance. Other elements, which include mechanical, haptic or even olfactory manipulation [62] [81] [2], are slowly becoming available. These are an extensive area of research in the field of HCI (Human Computer Interfaces). We hope to include such elements in future iterations of our prototype.

The requirement, that manipulation is controlled digitally, is key to enable the development of future smart applications on top of our design. It is crucial to establish a closed-feedback loop, i.e. allowing an intelligent agent to carry out its recommendations and gauge them to adapt future recommendations accordingly.

3.1.2 *Real-time response*

We envision the smart applications built on top of our platform to be data and computation driven. To predict the optimal configuration of the Digital Cubicle, they must rely on current and accurate data from the cubicle's occupant. This data should reflect aspects of the occupant state, context and even intentions – all of which change in real time. This requirement is relevant for appearance manipulation but also, and primarily, for data collection. The Digital Cubicle needs to handle multiple streams of information, from various sensors, and process them in real time. It must also integrate them into a single computation context, allowing data fusion from multiple streams to make recommendations and predictions.

3.1.3 *Extensibility*

As mentioned, we view the Digital Cubicle as a platform for research and exploration. Our prototype takes the first steps towards an implementation of digitally mediated ambiance in the workspace. As such it should provide opportunity to manipulate as many physical characteristics as possible in the room. We attempt to create a comprehensive set of these, including different stimulus modalities and employing multiple techniques for ambient display. We put particular emphasis on the platform's extensibility. Our design lays the groundwork for combining with software multiple elements that can result in physical changes in the room. In fact, we many times favoured covering a breadth of possible elements, rather than refining particular ones. As a result many of the elements we pick can be further perfected and improved. We note possible immediate improvements, where relevant, throughout the following sections.

We take the same approach when designing our infrastructure for real-time sensor data collection. We focus on building an engine that many different sensors can be plugged into. We demonstrate its robustness by integrating a sample set of sensors from different vendors and for different modalities.

3.1.4 *Applicability*

We build the Digital Cubicle with how economically viable, scalable and easily deployed it would be, in mind. As a result we limit our design to include relatively accessible mid-range commercial elements. This characterises the video projection and display materials, audio speakers and physical construction. The lighting fixtures in our prototype are less accessible at this point, but we expect them to be more affordable in the near future. We focus on designing a single employee cubicle, which is widespread in contemporary office spaces and can be integrated into existing built environments. We draw inspiration from the original Herman Miller cubicle design [89], which is configured from modular elements depending on the occupant's needs.

3.1.5 *Composition*

Assume one could completely manipulate the physical interior of a workspace, how would they stage it? Psychological research into physical manipulation of workspaces has primarily been concerned with isolating a single environmental feature and measuring its effects. Experience and perception, however, are higher psychological constructs, shaped via multiple features, such as colours, textures, movement, artefacts, etc. Architects in the built environment and Stage designers in theatres have long adopted a more holistic approach to transform and define spaces. They name the art of putting together a self-contained harmonious environment, conveying designated ambiance, composition.

Our Digital Cubicle is a blank canvas for composition. It challengingly provides countless opportunities for creative combination of sounds, images colours, etc. It even allows those to evolve in time and space. To scope our work, we turn to existing working environments, already designed to foster different experiences. Instead of composing environments on our own, we experiment with recording these environments and artificially recreating them in the Digital Cubicle. This novel experimentation is also valuable as an evaluative tool. It allows us to compare the perception and experiences of the Digital Cubicle with real working environments experiences. In the following section (3.2), we describe how we visited and recorded some real working environments. We used these as an evaluative tool in the first user study mentioned in the introduction (See 1.1).

From hereafter, we distinguish between two types of users of the Digital Cubicle: composers and occupants. The former use the cubicle's controllable elements as a canvas, creating compositions to induce experiences and perception. The latter use the cubicle as a

workspace, choosing between presets and adjusting a minimal set of parameters for convenience.

3.2 CAPTURING AMBIANCE OF REAL WORKING ENVIRONMENTS

As mentioned in the previous section, we turned to real working environments to facilitate the composition of different working experiences in the Digital Cubicle. We visited and recorded about 10 working spaces around the MIT campus. The selected spaces range from popular spaces for focusing to inspiring and thought evoking open art displays. We intentionally included a mixture of nature and outdoor settings as well as urban and indoor ones. The recordings were used as content from which we rendered an artificial composition in the Digital cubicle, mimicking some of the physical characteristics of the real environment. The process of mapping the recordings to the Digital Cubicle was done manually, leveraging a set of control tools we specifically developed for this purpose. We intentionally kept it, however, simple in nature to allow automation of the process in the future.

We erred on the side of caution to explicitly obtain consent from space owners for each video recording. A sign was placed near the recording setup notifying passer-bys that video is recorded, and allowing them to request that the videos are erased. Figure 1 displays an image of our recording setup in action, recording the inspiring indoor gardens at the Isabella Stewart Gardner museum. Note the simplicity and compactness of our setup.

3.2.1 *Describing lighting conditions*

To express lighting conditions in real working environments we employed a simple model with three descriptors: light sources configuration, colour temperature and overall luminosity. This model is far, of course, from fully describing an environment's lighting composition. It neglects for example movement, contrast, reflections and shadowing, all of which arguably are equally important to induce experiences and affect perception. We choose, however, this simplified subset in light of the limitations of the artificial lighting we can create in the room.

To measure light sources configuration, we photograph the shadowing pattern created by a simple vertical rod. The number of shadows, their length, and how diffused they are, indicate the number of light sources in the environment, their directionality and brightness. We use a colour reference card (ColorChecker) to measure colour temperature, and a LuxMeter to measure the luminosity on the working surface in the environment. Figure 2 displays these measurement devices.




Figure 1: The recording setup, including a camera and a tetrahedral microphone, in the Isabella Stewart Gardner museum

3.2.2 *Video recording*

We used a Canon 55D camera to record videos from real working environments. We set the camera in a stationary angle looking horizontally at the environment from a viewer's height. The videos are recorded in a 1920 X 1080 resolution and a 60 frame per seconds rate. According to the sensor specifications, the camera's focal length and zoom are set so that recorded objects on the focal plane will maintain their realistic size, when viewed on the projection display from the working desk. The resulting focal length is relatively short (about 1m), naturally blurring distant objects and supporting an illusion that the objects are in the visual periphery. Each recorded video is about 30 minutes long. For future recordings (pending on budget), we recommend to consider 360° video recording technologies. These will allow composers higher flexibility, setting viewer perspective retrospectively through video processing.




Figure 2: Top view of the measurement tools used to describe the lighting conditions. From the left: Colour reference card, rod for shadowing pattern and a LuxMeter.

3.2.3 Audio recording

We used a tetrahedral microphone to record audio from the real working environments. The microphone is placed next to the video recording camera in identical orientation. The microphone was developed by Mayton, and previously used as part of the Tidmarsh Living Observatory Restoration Project [54]. It uses four cardioid capsules in a tetrahedron configuration, calibrated all to have equal gain. Using the microphone we record four audio signals from the microphone capsules in a format called A-format. With a mathematical matrix operation we convert the recordings to B-format, which also contain four audio channels. Three channels correspond to the signal that would have been picked up by figure eight capsules oriented along the X, Y and Z axes respectively. The fourth channel contains the signal that would have been picked up by an omnidirectional microphone, recording a perfect three-dimensional sphere. We use Tetraproc [1] for the aforementioned format conversion. Once the signal is in B-format, It can be further processed, to artificially mimic the spatial qualities of the original sound.

3.3 PHYSICAL CHARACTERISTICS MANIPULATION

In the first phase of our implementation, we set out to design and build a workspace that can manipulate as many of its physical characteristics as possible. We do so by integrating three elements: lighting, video projection and sound. Images of the final result are presented in Figure 3. In the following subsections, we describe each element in detail – the physical setup we developed and a set of software tools tailored to allow composers to control each of the elements. We then bundle all these tools to a single service we call *Scene Control Service*. The service provides occupants (or smart agents on their behalf)



Figure 3: Manipulating physical characteristics in the Digital Cubicle.




Figure 4: Physical layout of the Digital Cubicle.

primitives to choose from a set of presets and adjust some of their parameters. The control tools and the *Scene Control Service* all run on a single iMac computer connected to a local area network and located at the back of the video projection display.

3.3.1 Lighting

The first controllable visual manipulation we develop is lighting. We base its physical design on the work of Zhao et al. [98], who built an office room with highly controllable lighting fixtures. The office room is a windowless rectangular room (4.2m X 2.8m with ceiling height of 2.6m), with 20 individually controlled luminaire groups installed in the ceiling. Six wall-washing fixtures (*Colour Kinetics Skyribbon Wall Washing Powercore*) are installed along the long edges of the room and light the walls to the occupant's left and right. Two ceiling-recessed fixtures (*Colour Kinetics Skyribbon Linear Direct Powercore*) are installed in the centre, directly on top of the occupant's workstation. The physical layout of the room is sketched in Figure 4.

A central server, running *Color Kinetics Data Enabler Pro*, controls all fixtures. It processes and executes control commands in an OSC format, received via its local network. Each luminaire group exports 5 channels – Red, Green, Blue, Warm White (2700K) and Cold White (4000K). The wall-washer groups can be set with an 8-bit resolution to any value between 0 and 1451, 2315 or 2867 lumens for the RGB, Warm White and Cold White channels, respectively. The ceiling-recessed groups can be similarly set with an 8-bit resolution between 0 and 1344, 2176 or 2888 lumens for the RGB, Warm White and Cold White channels, respectively.

Though extremely flexible, the aforementioned lighting configuration raises a challenging control problem. There are 100 parameters that can be varied (20 groups with 5 channels each), with 255 setting options for each parameter. Every parameter configuration results in

distinguishable lighting changes in the room, and from a human-centric perspective, sets a completely different tone for the room's atmosphere. Zhao et al. [98] discuss this problem in further detail, and propose a control scheme, which maps intended working context into parameter configurations. Their scheme is based on human perceptual rating of a set of compositions, whose parameters are then cleverly extrapolated to create contextual control axes.

Imitating the lighting conditions of existing workspaces requires a different mapping approach. Such mapping will have to translate lighting conditions we observe in the space into a parameter configuration for the room. We refer back to the descriptors we used to measure the lighting conditions in the real working environments, and design our mapping accordingly. Figure 5 illustrates the top-level software design of the implementation of our scheme. It translates a user interface, which uses light sources, colour temperature and luminosity terminology, into RGB parameter configurations for each of the fixtures in the room.




Figure 5: Top-level software design of the lighting control

1. Fixtures

Fixtures are convenience software objects that represent each of the physical fixtures in the room. They introduce additional colour representations, other than RGB, that one can use to set the fixture lighting with. Hue Saturation Value (HSV) is a cylindrical coordinate representation, which spans the same colour space as RGB. It is a necessary representation for some colour blending models, and particularly convenient in our context to aggregate lightness values from multiple light sources. Correlated Colour Temperature (CCT) is another colour representation, which only represents a subset of the colour space – black

body radiation. It is useful to describe natural light, indicating the whiteness of the light source. It ranges from low temperatures for warm white light to high temperatures for cold white light. Python's *coloursys* library is used to implement RGB to HSV conversions and vice versa. We implement simple lookup table conversion functionality for CCT, based on factory calibration data for our fixtures.

2. Virtual Light Sources

A "virtual light source" is the basic primitive of our lighting mapping-scheme. It can be thought of as an illuminating point in a two dimensional representation of the room (Width X Length). It has a composer-set lightness, which decays with distance like a Gaussian function. The Gaussian's standard deviation is an additional composer-set parameter. Our user interface for lighting control allows adding as many virtual light sources as necessary to the two dimensional representation, and dragging them freely around. Figure 6 illustrates the controller's user interface and its virtual light sources. We built the user interface as a web page, accessible via the room's local area network.

The virtual light sources are mapped to HSV parameters for each of the room's fixtures. To map a single virtual light source, we calculate its Euclidian distance to each of the fixtures in the two dimensional representation of the room. According to the distances, we compute the source's lightness contributions to each of the fixtures (Value in an HSV representation), decaying the source's lightness like a Gaussian. By aggregating contributions from all virtual light sources, we determine the total lightness value for a particular fixture. Colour temperature is set globally to all virtual light sources. Our implementation however can easily be extended to support colour blending of virtual light sources with different colour temperatures.

This controlling scheme allows us to imitate some of the lighting characteristics of the real working environments we observed. To set the artificial lighting, we arrange a source configuration that conforms as best as we can to the shadowing pattern recorded in the real working environment. Moving the sources around, we can set the shadowing directionality and each shadow's length. The room's colour temperature is set according to the colour we have recorded. We additionally measure the luminosity levels on the table and adjust the light sources' lightness, to achieve luminosity levels similar to the ones measured in the real environment.

3. **Lighting Fader** The *Lighting Fader* is responsible for smoothly transitioning between lighting compositions. Once a lighting composition is artificially created, its parameter configuration




Figure 6: User interface of the lighting mapping controller. The yellow circles are virtual light sources within the room’s two dimensional representation. Lightness, colour temperature and decay parameters for each are set on the right.

can be saved for later use. Commonly it is matched with audio and video recordings from the same real environment. Either the occupant or a smart agent can request loading existing lighting compositions and rendering them in the room. When a composition is loaded, the *Lighting Fader* generates a set of intermediate configurations to smoothly transition between the current lighting state and the loaded composition. It logarithmically extrapolates between the current and destination lightness and colour temperature of each fixture in the room. The extrapolation curve is then finely sampled to result in intermediate configuration steps, that match the logarithmic visual perception of the human eye (Weber–Fechner law [32]).

3.3.2 Video projection

The second visual manipulation we developed uses video projection. Unlike lighting, video projection allows content display and move-




Figure 7: The display with and without active projection

ment – two powerful tools widely used to generate experiences and alter perception in artistic contexts. We design a video display using 183cm by 76cm light defusing acrylic material. Without projection, the display is white and opaque, mimicking the aesthetics of a modern cubicle wall-divider. The display is placed in front of a working desk and is levelled to the height of a sitting person. We carpentered a wooden base to hold the display frameless, and give it the appearance of an integral part of the working desk. Figure 7 shows on the bottom the display when no content is projected onto it. The top images show the display with active projection.

This design is the result of numerous iterations, with potential occupants and composers, in which we tested a wide range of possible displays. The current design attempts to reconcile the main feedback points our users surfaced. For example, the display was intentionally cropped not to follow familiar display ratios. We identified that familiar ratios (e.g. 16:9 or 4:3), automatically reduced occupant immersion and thus result in lower effectiveness changing occupant perception. For similar reasons, we designed the display to be frameless and float in mid air.

We also intentionally placed the display in the front, covering most of the occupant's peripheral vision. Horizontal and vertical viewing angles are 100° and 53° , respectively. Side displays were found less effective. Our experimentation also included back displays with a mirror placed in the front, in which the occupants could see themselves with the display as a backdrop. We notice a delicate tradeoff between immersion and ambiance in the display design. On the one hand, the display should be immersive enough to affect occupant's experiences

and perception. On the other, and unlike virtual reality systems, occupants are expected to engage in productive work in its presence. The display, hence, must maintain a balance, remaining ambient but still in the outskirts of the occupant's attention.

We project onto the display from behind, using a high luminous projector (NP-PA571W 5700 Lumen) with short throw lens (NENP30ZL). Rear projection technology was selected to allow enhanced image contrast and brightness in the presence of ambient light. This is of particular importance in our case, as ambient light is one of the physical features we manipulate. Throughout our iterations, in fact, we note that the combination of rear projection and ambient light can be leveraged to effectively create an artificial impression of depth. Placing light sources with matching colours behind the display creates an illusion of extending it and changes one's perception of how deep the displayed image is. Rear projection also assures there are no occlusions between the projector and the display. This prevents shadowing from tainting the desired experience.

We would like to finally mention two possible areas of improvement for our current design. These we believe can upgrade the setup's immersive experience and applicability. First, the projector's noise should be attenuated. The projector's cooling noise has frequently been complained to be a source of disturbance, interfering with occupants' experiences. Second, the distance between the projector and the display should be shortened to make the setup more applicable. This can be achieved by a combination of using an ultra short throw projector and modifying the display material. Using an ultra short throw projector decreases the projection distance to about 40cm. The acute throw angle, however, changes the angular visible range, such that the optimal viewing angle is above the display rather than in front of it, substantially reducing the image quality. This can be corrected by manufacturing a directed acrylic material that can refract the light beam, compensating for the loss in brightness. The price point for this alternative however will unfortunately be significantly higher.

To control video projection we program a video composition controller using the data flow framework Max MSP 7. An image of the controller interface is presented in Figure 8. The controller is responsible for two main tasks. First, it handles various image transformations needed to accurately aim and calibrate the projected images onto a display. Second, it allows applying video processing filters to the projected videos and change their parameters in real time. The former is important to enable flexible experimentation with display configurations in our Digital Cubicles; Similar to the ones described above. The latter is important to provide both occupants and composers with the capability to fine tune the visual stimuli.




Figure 8: Max Msp Controller developed for video composition

The controller creates an OpenGL context for display, in which it renders loaded video frames. We use OpenGL so that we can take advantage of hardware-accelerated processing and display features, keeping the controller as computationally efficient as possible. In fact, in previous iterations we built an OpenCV video controller, based on CPU processing. This attempt was unsuccessful, and could not computationally support the video processing performance capabilities we require. The controller, is capable of loading any type of video files onto the OpenGL context using objects from the *jit.gl* library.

After selecting the files to load, the controller allows a set of image transformations to aim and calibrate the projected video onto a display. These include, cropping the video and the display window to a configurable size (e.g. 183cm X 76cm in our final setting), changing image position and perspective and scaling its resolution. We again use *jit.gl* objects to implement these transformations and program a simple sliders interface to control them.

Finally our controller allows application of image processing filters to tweak the images in real time. We implement this by chaining *jit.gl.slab* objects, which provide a streamlined interface to preform general-purpose GPU-based grid evaluations. Those manage compiling binding and submitting GPU-shaders to the controller's OpenGL context. The GPU-shaders are computer programs that run on graphics hardware with high degree of flexibility. They efficiently render visual effects on top of the original video frames. Our controller allows programming any type of image processing using GPU-shaders and exporting its parameters to the user interface.

The use of GPU-shaders allows composers and occupants to control the visual stimuli in real time. This is of particular importance in cases where, for instance, the visual stimulus is too demanding and the occupant wishes to blur, attenuate or distort it in different ways.

Other GPU-shaders can be programmed to moderate movement, alter contrast, modify warmth, and even artificially create various image qualities (like the popular Instagram filter). Though programming them takes some effort, they provide substantial flexibility in video projection composition. In our specific implementation we demonstrate the development of two such GPU-shaders – controlling in real time blur and opacity.

3.3.3 *Sound*

To play audio in the Digital Cubicle we use a quadrophonic speaker configuration. The speakers are positioned in a square, equally distant from the occupant who sits in the centre, and are levelled to ear height. In this configuration, we cannot reproduce the height information from the original recorded environments. Other configurations, however, (e.g. a cube configuration of eight speakers) can easily be introduced to preserve the sound’s height quality. In the scope of this work, unlike the video display, we did not experiment with additional speaker configurations.

The speakers are controlled with a MOTU 4Pre audio interface connected to the main system computer. In our current prototype they are placed on simple speaker stands in the workspace. For improved aesthetics, they can be easily integrated into the cubicle furniture, as exemplified by the Steelcase’s SonetQt and QtPro sound products [34].

To control the played audio, we program an audio configuration controller using Max MSP 7. An image of the controller interface is presented in Figure 9. This controller’s main purpose is to support different speaker configurations in the Digital Cubicle, while preserving the spatial sound qualities of the recorded from the real working environment.

We use ambisonics [72] to implement this control capability. As mentioned, we record the original sound using a tetrahedral microphone and process the recorded signal to a representation called B-format. Unlike multichannel surround sound formats, this is a speaker-free representation of a sound field. We can decode it using a first-order ambisonics decoder to a composer-defined speaker array. The speaker configuration is provided using a simple graphical interface, determining the orientation, distance and height of each speaker from the listener. The result, played by the speakers, is a first order approximation of the sound field at the listener position. In practice, a sound, for example, of a page flipping in the library, will appear to originate from the same direction and proximity in the Digital Cubicle.




Figure 9: Max Msp Controller developed for audio speaker configuration

3.3.4 Occupant interface

To provide occupants with an easy control over the Digital Cubicle, we bundled the aforementioned capabilities into a single service called *Scene Control Service*. The service saves lighting, video and sound compositions into what we hereafter call "scenes". We implement a simple web-based interface allowing occupants to select between previously set scenes, changing the atmosphere in the Digital Cubicle. Furthermore, we add three simple sliders enabling occupants to modify the scene volume, to reduce the projection brightness (opacity) and to blur the projected video to their liking.

The *Scene Control Service* additionally implements smooth transitioning between the various scenes. Lighting is transitioned by the *Lighting Fader*, previously described. Sound is cross-faded between the corresponding scenes' audio. Video projection is transitioned by logarithmically reducing the current scene brightness to zero, switching scenes and then logarithmically increasing brightness for the new one. The resulting outcome is a smooth and pleasant scene transition, which takes 8 seconds.

3.4 REAL-TIME SENSOR DATA COLLECTION

The second main part of our platform is a sensor data collection infrastructure. Figure 10 illustrates its top-level design. The infrastructure is a flexible code library which lays out the ground work to integrate a wide range of sensors into a single service we call *Sensor Collection Service*. When the service is executing, it initiates connections with a configurable list of sensors and begins harvesting data




Figure 10: Top-level design of the *Sensor Collection Service*

samples from all of them. External applications can register to the service, obtaining a processed version of the data samples from specific or all data streams. The sensor collection service executes as a single process on the Digital Cubicle central iMac computer, located at the back of the projection screen.

We built the infrastructure as an asynchronous code library implemented in Python. The library interfaces with a variety of sensors and can concurrently handle multiple real-time sensor data streams. We implement it based on the Twisted library [24], following a Reactor design pattern [73] (marked in red in the diagram). This design pattern implements an event driven scheme, in which multiple data streams are multiplexed into a single computation context. Whenever sufficient data is available on any of the streams, an appropriate event handler is called to process it. Processing can be done in layers to improve concurrency, breaking it to smaller sequential event handlers, which trigger each other also through events. A load balancing mech-

anism prevents starvation, assuring that all streams receive adequate processing time.

We select this design pattern to allow processing of different real-time sensor data streams in the same context. This property is particularly important to allow fusing information from different streams to make predictions or inference. For example, suppose we need to infer one's concentration levels using their pupil orientation, and can improve our prediction based on their arousal levels from skin conductance. Our programming scheme allows us to access both information streams from the same context, and as soon as they become available, avoiding cumbersome intermediate mediation through files or inter-process communication. It also allows us to make real time computation involving the different streams to design and assess relatively complex classification features.

With an underlying event-driven design, we implemented software support for three types of sensors: Empatica E4 wristband [26], Zephyr Bioharness 3 [96] and a simple Logitech web camera. Figure 11 displays a photograph of an occupant, instrumented with these sensors in the Digital Cubicle. The camera, in this case, is used to extract facial feature information from the occupant, utilising the Intraface [94] code library. For each of the sensors we implemented a stack of handlers, which process their data online and parse their unique protocol (marked in green in the diagram). Once the data are processed it is queued, in a convenient tuple format, at the application layer. High-level applications can register to our service, specifying the information streams they are interested in. From then on, parsed stream data that was queued in the application layer is passed on to them whenever it is available. The high-level applications can use the data freely, e.g. to make predictions about the occupant's concentration. The smart agent we build in Chapter 5 is an example of such high-level application.

Zephyr's Bioharness 3

Zephyr Bioharness 3 is a physiological monitoring module. It is attached to a strap that is wrapped around the occupant's chest. We primarily use it to collect occupants' Electrocardiography (ECG) and respiration signals. It is equipped with on-board algorithms processing the raw signals and extracting measures, such as heart rate, respiration rate and heart rate variability (HRV). A 3-axis IMU in the module additionally collect acceleration signals and estimates the wearer's posture and activity.

The module has a Bluetooth interface through which it can connect to a Zephyr proprietary software server (OmniSense 4), logging data and visualising it. Since the server does not support the real-time capabilities we require, we contacted the company who kindly released the device raw Bluetooth protocol specifications. Based on




Figure 11: Occupant wearing the set of supported sensors in the Digital Cubicle

these we build our event-driven handler stack, communicating with the module from our central server and parsing incoming information in real-time.

The first protocol layer aggregates bytes received over Bluetooth into frames, in a protocol similar to HDLC. It verifies their integrity and implements a simple continuous handshake protocol to maintaining the link active. The second layer is responsible for parsing the various messages encoded into the frames. Internally, the device allows registering to each of its sensors separately. That is, we can request to listen separately to the ECG signal, the respiration signal, the acceleration signal, etc. An additional *Summary* stream is also available, it contains processed data signals, computed on-board, such as heart and respiration rates. The final layer transforms incoming data to a convenient object-based format, e.g. handling sample compression in some of the streams and synchronising timestamps.

Empatica E4 Wristband

The Empatica E4 wristband is also a physiological monitoring module. It senses physiological signals from the wearer's wrist. It is equipped with two electrodes measuring Electro Dermal Activity (EDA), i.e. variations in the electrical characteristics of the skin. These are related to activation of the sympathetic nervous system and are an indication of changes in physiological or psychological arousal [26]. The module also contains a 3-axis accelerometer, an infrared Thermopile measuring skin temperature, and a Photoplethysmography sensor, measuring blood volume pulse (BVP) and extracting heart rate and heart rate variability.

The module connects wirelessly via Bluetooth to an Empatica BLE server, which parses the transmitted data. Our Sensor Collection in-



Figure 12: Facial feature points, head and pupil orientation extracted by the Intraface code library.

frastructure connects to the server and pulls in real-time incoming data in a text format. A set of handlers parses the text messages in a similar manner to the one described above. As a result high-level applications can access in a convenient tuple format streams of samples from all of the sensors mentioned above.

Intraface

We positioned a Logitech web camera on the Digital Cubicle desk, to capture videos of the occupants as they work. Videos are recorded in real-time in a 12fps frame rate. The frames are then processed using the Intraface code library for facial image analysis. The library extracts four types of features from the images (illustrated in Figure 12):

1. Coordinates of 49 facial feature points, describing the eyebrows, eyes, nose and lips.
2. A vector estimating the 3D head orientation of the occupant.
3. Vectors estimating the 3D viewing orientation for each of the occupant's pupils.
4. Intensity predictions of 6 emotions the occupant expresses (neutral, angry, disgust, happy, sad and surprised).

Since the Intraface library is written in C++, we implemented the facial feature extraction process in a separate process. The process executes on our main server, and communicates its results to the *Sensor Collection Service* via an internal connection. Collected data are passed on to the high-level application, just like any other sensor data.

PHYSIOLOGICAL EFFECTS

In this chapter we set out to examine occupants' physiological responses to a set of different scenes artificially rendered in the Digital Cubicle. We focus our examination on physiological responses associated with the autonomic nervous system, which indicate arousal and relaxation. These, if found, can be leveraged by future applications to target chronic stress and related physical and mental illnesses.

We rely on the Digital Cubicle's real-time sensor data collection infrastructure to conduct a user study, and collect a rich dataset of physiological responses. The first section of this chapter delineates the study's protocol, how we selected its scenes, and the dataset we collected. In the second section, we develop a method to process the physiological signals, and define a set of metrics to quantify stress development and restoration. Based on those, we present the study's results and discuss our observations in Section 4.3. We conclude by summarising our learnings and highlighting relevant observations that will guide the development of an intelligent control agent in the next chapter.

4.1 EXPERIMENTAL DESIGN

4.1.1 Study Protocol

Sitting Structure						
Calibration	Tutorial Session	Session 1	Session 2	...	Session 6	

Session Structure:						
Clock:	30 Sec.	1.5 Min.	3 Min.	3 Min.	30 Sec.	Unlimited
Task:	Wait	Test Segment (Reading Comprehension)	Break Segment (Stress restoration)	Necker	Survey	
Scene:	Neutral	Randomly Selected Scene (Neutral, Forest, Rotch, Kites, Shibuja or Sunset)				

Figure 13: Schematic illustration of the study's protocol. Each participant was invited to a single sitting consisting of six identical sessions. Sessions were structured as illustrated above.

A set of 9 participants (4 female) were recruited for a study, following IRB protocol #1601357773. All 9 participants were healthy adults between the ages of 19 and 38 with no history of cardiovascular disease or medication for hypertension. The participants were all right handed and fluent in English.

The study participants were invited to an approximately one hour sitting in the Digital Cubicle. The sitting was divided into six identically structured sequential sessions. In each session, the participants were exposed to a different scene rendered in the cubicle. The order of the rendered scenes was randomised. In one of the six scenes no content was rendered, i.e. the office divider was blank, no sound was played, and lighting was set to common office settings. We name this scene "Neutral" and regard it our control condition. The remaining five rendered scenes were selected as described in Subsection 4.1.2.

Within each session participants were asked to perform a sequence of cognitive tasks sitting in front of a laptop computer. Throughout the session their physiological signals were recorded using a set of wireless sensors, communicating in real time with the Digital Cubicle's backend. We designed the sessions as illustrated in Figure 13. To disconnect the session from any previous ones, participants were first exposed to 30 seconds of the "Neutral" scene. Immediately after, the room transitioned into the scene whose effects are tested. Participants were given a waiting period of a full minute in the rendered scene before they were presented with any tasks. This waiting period was designed to allow acclimation and avoid possible novelty effects, originating from the initial exposure, which may bias the physiological signals during the tasks.

The two main segments of the session were a stress-eliciting task, followed by a restorative break. We used a graduate-level ETS GRE [21] reading comprehension assignment as the stress-eliciting task. We programmed the assignment into a website, and it was presented to the user on a laptop computer. Each assignment contained 2 to 4 paragraphs of an academic text and three questions relating to it. Participants were given three minutes to select a correct answer to each question, after which their answers were automatically submitted and the test was cleared from the screen. A three minutes restorative break followed the test. During the break, participants were requested to remain seated, but were free to relax and explore the surroundings as they please.

To increase participants' stress during the test, we reduced the average time allotted per question to a minute, rather than 1.5 minutes allotted in the original design of the test. Additionally, participants were informed that their test scores will be compared to the test scores of others and that their performance is of critical importance to the success of the study. Finally, a timer and a progress bar at the top of the website indicated the amount of time remaining for the test.

Once the break was over, the website assigned the participant with an additional task – the Necker Cube Pattern Control Test [14] [15]. This test is designed to measure one's capacity to direct mental effort, and is widely administered in attention research. Participants were presented with the cube in Figure 14, which can be perceived in two

different orientations (with the top square closest to you or furthest from you). When viewed for prolonged period of time the cube spontaneously reverses its perceived orientation. The participants were presented with the cube for 30 seconds during which they were asked to sustain, for as long as they can, a particular orientation for the cube. Whenever the orientation reversed, they would click on the cube and the website recorded their interaction. The time period one can sustain the cube in a particular orientation was previously correlated with one's ability to direct attention [80]. This suggests that the more reversals a participant reported during the Necker test, the less he or she was able to direct their attention in the rendered scene.




Figure 14: Necker Cube Pattern Control Test. Participants were presented with the leftmost image. The following images illustrate the two possible perceived orientations.

To conclude the session, participants were asked to answer a short survey about the rendered scene. No time limit was enforced on answering the survey. We focused our survey on the restorative qualities of the scene, hoping to later correlate those with the recorded physiological signals. Two questions asked the participants to imagine themselves in a particular situation and assess how suitable the rendered scene is for it. The participants rated the scene suitability on a Likert scale from -2 ("Not suitable") to 2 ("Very suitable"). The first question described a situation when one is recovering from prolonged mental effort. The second question described a situation when one is full of energy and has to direct effort towards a new task.

The remaining questions in the survey were based on the revision and verification of the PRS Questionnaire done by [64]. Five additional questions were designed to measure the five different facets of a restorative environment: Compatibility, Coherence, Being-away, Fascination and Scope. Coherence was measured inversely, asking the participants how chaotic and confusing the scene is. The Compatibility was phrased to focus on the occupant's personal preference and liking. A copy of the survey with the specific phrasing is affixed in Appendix a.

To reduce the novelty effect of the experimental sessions and the wearable sensors, each sitting began with a tutorial guided by the study personnel. Through an example session, the tutorial walked the participant through the different parts of the session and made

sure they are all well understood. Furthermore, the tutorial included a short calibration phase, in which each participant's viewing orientation was matched with the borders of the laptop screen. This allows us to later on process the facial orientation data and label viewing directions whose trajectories are inside and outside the screen. Once the tutorial was over, the participants were left alone in the workspace, and our website guided them through the sessions and their different parts. The website also controlled the transitions between the scenes in the room (through the Digital Cubicle backend), timing the various parts, and collecting the participants' interactions.

4.1.2 *Scenes Selection*

We hypothesise that scenes perceived differently in their restorative qualities would result in distinguishable related physiological responses. We expect to find these differences during intermissions or following periods of directed mental and cognitive efforts in the Digital Cubicle. Such a relationship was demonstrated before in the context of attention restoration theory. [11], for example, showed a correlation between PRS scores of images and viewers' restoration, measured through BVP, EMG and EEG, during the first 10 viewing seconds. Similarly, [27] demonstrated significant changes in BVP, when taking break periods in natural landscapes rather than urban ones.

Note, however, that environmental stress restoration is context dependant. It is not likely to function in all working situations, particularly, when occupants attempt to direct their attention and focus. Empirical studies have shown that ambient stimuli from the physical environment can impose working-memory load, taking away resources from the occupant's intended cognitive process [19] [71]. In this sense, some of the rendered scenes may in fact act as stressors. This is regardless of perceived restorative qualities. Their visual and auditory features may distract the occupant and adversely result in heightened stress.

We picked five different scenes for our user study, attempting to represent diversity of perceived restorative and focusing qualities. Images of the five scenes are presented in Figure 15. The scenes' audio and lighting were matched to the projected videos. Sounds of a street tumult, for instance, were played in the Shibuja scene and sounds of wind and waves were played in the Kites scene. Each scene was coded with two binary labels corresponding to what we hypothesised would be its restoring and distracting effects. The Library scene, for example, was coded with negative restoration and positive focus. In a similar manner, the Kites scene was coded with positive restoration and negative focus. We verify this coding scheme through self-reported measurements and present the related results in the next




Figure 15: Images of the six scenes used in the physiological user study. From left to right and top to bottom: Neutral, Forest, Rotch Library, Kites, Shibuja and Sunset.

sections. The situation in which no scene is rendered in the room is named "Neutral" and we consider it our control condition.

In this study, the scenes' audio was played through a set of BOSE noise cancelling headphones. This is the result of feedback from our previous user study, in which the projector's sound was reported substantially inconvenient and distracting. We adjusted the ambisonic processing of the Digital Cubicle accordingly. The scenes' sound levels were also tuned, such that all scenes would have similar volume.

4.1.3 Collected Data

We used a single Empatica E4 wristband [26], worn on the non-dominant hand wrist, to record participants' EDA signals. To reduce noise in the recorded signal, the participants were asked to minimise the movement of their left arm throughout the sessions. Zephyr's Bioharness 3 [96] was used to record participants' ECG and Respiration signals. A Logitech Webcam was used to record videos of the participants' face. The video frames were streamed through the Digital Cubicle backend software and were processed to extract facial feature points and facial orientation using the Intraface [94] code library.




Figure 16: Sample heart rate, respiration rate and HRV signals, collected from Subject 6 when he was exposed to the Shibuya scene.

Figure 17 exemplifies heart rate, respiration rate and HRV signals, that were collected through Zephyr’s Bioharness 3. The presented signals were collected from a single session. The vertical lines mark the different segments within the session. Similar signals were collected for each session from each of our participants – a total of 54 sessions per physiological signal. Due to low reliability, indicated by the device, we removed from our dataset the HRV signals recorded during the Sunset sessions of Subjects 1, 3 and 5.

The aforementioned signals are the result of the device’s built-in algorithm which processes the raw ECG signal. The sampling rate of the processed signals is 1Hz. Accordingly, each signal contains about 500 samples, depending on the survey segment duration. We smooth the heart rate and respiration rate signals using a 20 samples moving average filter. The device’s HRV extraction method is proprietary. It uses a relatively large window, which introduces a similar time shift. The raw ECG signal is additionally recorded to allow applying different processing algorithms in the future (e.g. frequency domain methods for HRV analysis). For similar reasons, we also record the participants’ RR signal, which measures the time interval between consecutive heart beats. The sampling rates for the aforementioned signals are 250Hz and 18Hz, respectively. Raw respiration signal is recorded in a 18Hz sampling rate. Finally, participants’ posture is




Figure 17: Sample EDA signal, collected from Subject 7 when he was exposed to the Neutral scene.

recorded by the device’s built-in 3-axis accelerometer, in a 50Hz sampling rate.

Figure 17 exemplifies an EDA signal collected through the Empatica E4 wristband. We collect one such signal from each participant in each of the sessions, resulting in 54 EDA signals. The sampling rate for each EDA signal is 4Hz. Perhaps due to low skin perspiration, we identify erroneous EDA signals collected from participants 1 and 4, and mark those missing in our dataset. The collected videos were recorded and processed in a 10fps rate. Intraface features including viewing directions, head orientation and facial expression feature points, as detailed in Section 3.4, were extracted for each of the frames.

Collected data is available for release by request, according to the restrictions delineated in IRB protocol #1601357773. In the scope of this work, we will only analyse the heart rate, respiration rate and HRV signals processed through the Zephyr’s Bioharness 3 built-in algorithm. The remaining signals are left for future analysis.

4.2 DATA PROCESSING AND CHARACTERISATION

We must first define an appropriate metric from the physiological signals, to assess one’s arousal levels and restoration when experiencing the different artificially rendered scenes. A wide range of such metrics has been previously proposed in the literature. The simplest ones report aggregated statistics (such as median, mean and standard deviation) of the physiological signals and compare those between the different experimental conditions. Such an approach would be very limited in the context of our study. Physiology is responsive to many psychological and physical influences besides the physical settings. Cognitive workload, for instance, is a varying factor throughout our study, which is very likely to affect physiology, even more than the physical settings [92]. Furthermore, physiology and physiological regulation are highly personal and state dependent. Autonomic heart rate regulation, for example, can vary greatly among individuals with

different training routines [82] [17]. Even within individuals, simply a cup of coffee can influence the recorded physiological signature [4].

Other approaches attempt to characterise trends in the signal, circumventing absolute values. In the aforementioned work by [41] heart rate recovery was measured by fitting a regression line to the first 60 seconds of each experimental activity. When applied to our recorded signal, this approach failed to truthfully capture the patterns in our data. Significant differences in the response time among individuals made it challenging to round up arousal and restoration with a fixed time window. While some participants tended to build up stress in anticipation for the test to begin (e.g. Subject 8's respiration rate signal), others tended to build it shortly before the test expired (e.g. Subject 1's respiration rate signal). Likewise, some participants showed multiple trends within a single segment inside a session. Subject 7's heart rate, for instance, tended to mostly recover during the break, but then build up again at some point before the Necker Test. The designed restoration metric must stay robust to such inconsistencies in the data. We wish for it to probabilistically capture those, dealing with real world noise and divergence in response.

Furthermore, the sequential design of the sessions made it difficult to disentangle the dependency between the Test and Break periods. We notice that distracting scenes, which caused high arousal levels during the Test, were likely to decrease the same levels sharply during the break. It would clearly be inaccurate to conclude these scenes were more restorative during the break, neglecting the physiological levels before the Break onset.

This observation coincides well with Wilder's Law of Initial Values (LIV) [91]. This law states that the effect of a stimulus on physiological change depends to a large degree on the pre-stimulus physiological level. Heart rate, for example, tends to innately be regulated regardless of any stimulus: decrease when high and increase when low. In turn, a stimulus that normally increases heart rate might, when heart rate is already high, only keep it at the same high level. Our analysis should therefore control for this effect in an attempt to isolate the restorative influence of the rendered scenes. When doing so, it would also take into account personal differences in natural physiological regulation.

Based on these considerations we propose a three-step data analysis approach to extract measures of the restorative qualities of the rendered scenes. Our analysis is inspired by the work of [38], who built an affective music playlist recommendation system to direct skin temperature and conductance levels. Drawing analogies from their system's training phase, we present a restoration metric that follows three design principles:

1. *Individual*: We attempt to model individual differences in physiology, recovery tendency, and stimulus response times and can-

cel them from the recorded signals. This, we believe, reflects better the effects of the rendered scenes, per participant, through individual experience and preference.

2. *Probabilistic*: We avoid aggregated statistics and utilise probabilistic modelling to represent the scenes' effects. These deal better with divergence in the physiological signals and describe better multiple trends that may arise during the relatively long sessions.
3. *Law of initial values*: We acknowledge pre stimulus levels and use baselining to incorporate the Law of Initial Values. The physiological signals are corrected accordingly and only then compared.

4.2.1 Signal Preprocessing

We begin by standardising the physiological signals, transforming each data point to its Z-score value [60]. For each participant and each signal we calculate the signal's mean μ and standard deviation σ over all the sessions and all their segments. We then transform each data point x_t , using:

$$z_t^i = \frac{x_t^i - \mu^i}{\sigma^i} \quad (1)$$

This transformation is applied to each of the physiological signals – heart rate, respiration rate and heart rate variability – separately. i indexes the study participants. It allows equaling out absolute signal values and representing the signals in units of individual standard deviations.

The signals are then filtered using a rolling mean window, which allows smoothing out high frequency changes. We chose a 20 samples (corresponding to 20 seconds) window for our filter and slide it one sample at a time. The top three graphs in Figure 18 illustrate filtering and standardisation of the heart rate measured from Subject 4 when she was exposed to the Shibuja scene. Note that changes in heart rate are quite visible between the different segments in the resulting signal. This participant's heart rate begins rising gradually as soon as the test begins and recovers slowly during the break. A second shorter increase can also be identified at the end of the break going into the Necker test.

Next, we compute delta scores Δz , indicating changes in the physiological signal. Each data point is subtracted from its K -th predecessor using:

$$\Delta z_t = z_t - z_{t-K} \quad (2)$$




Figure 18: Signal preprocessing example illustrating Subject 4's heart rate exposed to Shibuja. From the top: Original signal, Moving average filtering, Standardisation and Delta scores.

We chose $K = 20$ by visually inspecting the recorded signals, noting that typical changes in the signals' trend are apparent in 20-second intervals. This choice also takes into account the basic time scale of our experimental protocol, changing between segments with time intervals as short as 30 seconds. A 20-second time difference guarantees that physiological changes will only be computed between consecutive segments.

Since the delta scores are computed post filtering, each delta score in the resulting signal effectively represents the difference between averages of a 20-second time window and its preceding 20-second window. The bottom graph in Figure 18 continues the previous example and illustrates the resulting delta signal from Subject's 4 heart rate signal. It is important to note that the delta scores at the first 20 seconds of the break now contain information about the signal levels from the last 20 seconds of the test. Thus, we may now identify indi-




Figure 19: Application of the Law of Initial Values to Subject 1’s heart rate changes. The delta scores (Δz_t) are plotted against their preceding levels ($z_{t-\kappa}$). The scores are corrected to become $\Delta z'_t$, which is indicated by the dashed line

cations of a highly arousing scene during the test, by looking at the first 20 seconds of the break’s delta signal.

4.2.2 Law of Initial Values (LIV)

Figure 19 illustrates how the delta scores Δz_t computed in the previous section relate to their preceding physiological levels $z_{t-\kappa}$. The figure exemplifies the heart rate delta scores collected from all sessions and segments of Subject 1. A point on the graph ($z_{t-\kappa}, \Delta z_t$) indicates that when Subject 1’s heart rate was z , at some point during the break, 20 seconds afterwards her heart rate increased by Δz . The point’s colour indicates the scene Subject 1 was exposed to when it was measured. Wilder’s LIV is clearly apparent through this presentation. When the signal level is high the delta score tends to decrease, whereas, when the signal level is low the delta score tends to increase.

We model this effect as a simple regression line, describing a linear relationship between physiological change and its preceding physiological level:

$$y(z) = w_1 z + w_0 \quad (3)$$

Once the regression parameters are computed, we use them to correct each of the delta scores accordingly, following:

$$\Delta' z_t = \Delta z_t - y(z_{t-\kappa}) \quad (4)$$

The corrected delta scores represent one’s physiological changes throughout the sessions relative to his or her innate tendency to regulate towards a neutral state. Arguably, we can now regard them to be independent of any preceding physiological state and reflect only

the effects of the rendered scene and the experimental task. Corrected for individual tendency and scaled by individual standard deviation and mean, these are now also more suitable for comparison between participants.

We fit a regression model for each of the participants and each of the physiological signals separately. Each model describes a participant's tendency to regulate a specific physiological signal. We estimate the model's parameters on the participant's particular signal recorded in all sessions, throughout all of their segments.

We indeed find that the regulation slopes w_1 differ per participant and physiological signal. Heart rate regulation slope, for example, varies from -0.15 for Subject 7, to -0.87 for Subject 1. Respiration rate and HRV regulation slopes vary from -0.17 to -0.43 (Subject 9 and 8) and -0.1 to -0.03 (Subjects 7 and 4). In line with LIV, all fitted models result in negative slopes indicating regulation, rather than amplification. All regression models were found statistically significant with p values smaller than e^{-56} , e^{-55} and e^{-6} for the heart rate, respiration rate and HRV, respectively.

It is interesting to note that the models' R^2 values are relatively low and differ on average between the physiological signals (0.30 for heart rate, 0.12 for respiration rate, and 0.03 for HRV). This is an encouraging finding, in this context, as it indicates that LIV does not explain all the variance in our data. We assume that the experimental conditions and the rendered scenes will explain the remaining variance. To support this assumption, when we fit the models to data only collected during breaks, the average R^2 values of our participants increase to 0.4, 0.2 and 0.06 for heart rate, respiration rate and HRV respectively.

4.2.3 Building Probabilistic Models

Next, we wish to build a probabilistic model for the physiological changes for each participant. For each participant, we treat the corrected delta scores as random variables, and estimate them conditioned on the experimental task and the rendered scene. That is, we partition each participant's corrected delta scores by the scene in which they were collected. Then we partition them again by the test and break segments and for the scope of this work discard data from the remaining segments. We use each of the sets as observations from which we can estimate a probability density function (pdf): $P(\Delta'z \mid \text{Scene}, \text{Segment})$. Where $\text{Scene} \in \{\text{Neutral}, \text{Forest}, \dots\}$ and $\text{Segment} \in \{\text{Test}, \text{Break}\}$.

If we know what task the participant is conducting and the scene he or she is exposed to, this pdf can be interpreted as an estimation of what would be the probability to observe physiological change $\Delta'z$ in the next 20 seconds. The more similar physiological changes were

previously observed in our data, the higher this estimated probability will be.

We use Kernel Density Estimation (KDE) [60] with a Gaussian Kernel to construct the pdfs from the observations. This well established non-parametric method makes no prior assumptions on the probability distribution of the random variable and takes into account equally all available observations. It constructs Gaussians with bandwidth β around every data point and averages them to generate a smooth pdf, following:

$$P(\Delta'z) = \frac{1}{N} \sum_{n=1}^N K(\Delta'z | \Delta'z_n, \beta) \quad (5)$$

Where:

$$K(\Delta'z | \Delta'z_n, \beta) = \frac{1}{\sqrt{2\pi}\beta} e^{-\frac{\Delta'z - \Delta'z_n}{2\beta^2}} \quad (6)$$

To choose the Gaussian bandwidth, we utilise Silverman's rule-of-thumb [76] and correct it to introduce some robustness to outliers as proposed by [28]. The bandwidth parameter is calculated separately per estimation, following:

$$\beta = 1.06 \cdot \min\left(\sigma, \frac{R}{1.34}\right) \quad (7)$$

Where R is the interquartile range and σ is the standard deviation of the corrected delta scores.

Figure 20 exemplifies four pdfs of heart rate change estimated by the aforementioned procedure. The illustrated pdfs are for Subject 1 estimated from observations when she was exposed to the Neutral and Forest scenes during either the test or the break. Note we obtain different pdfs for the different conditions.

This probabilistic representation of the induced physiological change has some key advantages in the context of our work. The smoothing characteristics of the Gaussian Kernel improve our robustness to outliers. Unlike aggregated statistics, it captures more naturally the uncertainty in our data. If one of the scenes were to have 70% chance of decreasing a physiological signal but 30% chance of increasing it, such information would be preserved. Furthermore, a pdf representation allows easy visualisation, inspection and juxtaposition of the collected data. Visual inspection may serve as a powerful tool, identifying if two conditions have distinguishable distributions, and may even surface underlying dependencies unaccounted for. In fact, through a similar process, we have discovered the need to condition our estimation by the different segments. The different segments manifested at first as two distinguishable peaks in pdfs over full sessions, hinting they should be treated separately. In future research, we hope our proposed approach will facilitate gaining similar insights, where other factors are involved.




Figure 20: Heart rate change pdfs estimated for Subject 1 in the Neutral and Forest scenes. pdfs on the left and right based on test and break segment data respectively. The markers below indicate $\Delta z'$ samples from which the pdfs were estimated.

Our probabilistic representation was also developed in light of smart applications we foresee implemented in our mediated workspace. It is better suited to make unsupervised predictions for physiological changes occupants may experience, based on previous observations. For example, a smart agent can base its decisions on the constructed pdf to make a maximum likelihood estimation of the physiological change a scene will induce. It may even tune its prediction to a specific magnitude of change it is looking for (e.g. increasing heart rate by over one standard deviation); The risk of an adverse reaction can be assessed in a similar way and be taken into account in the agent's optimisation function.

In the following sections, we additionally show how we leverage this probabilistic representation to conduct evaluative simulations. Once the pdf is constructed it can be sampled, simulating one's physiological reactions. We note that constructing the pdfs is computationally efficient and that they can be easily updated as more observations become available to improve the quality of the simulation.

4.2.4 Stress Development and Restoration Metrics

Using the pdfs detailed above, we specify four metrics to evaluate stress development and restoration in the context of our work. We develop our metrics for a narrow notion of stress, based only on objective physiological signals and their well-established correlations with stress as an affective state [75]. In that sense, we define stress development as increasing heart and respiration rates and decreasing HRV. Contrarily, we consider stress restoration as reduction in heart and respiration rates and increase in HRV. The first two metrics we specify capture the intensity of physiological change and the second two capture its likelihood. These metrics will be applied on what we

established as physiological changes throughout this section. That is, changes observed in 20 seconds intervals, on top of (or under) one's innate tendency to regulate his or her physiological signal.

1. **Mean** - The average intensity of physiological change observed in the analysed population.
2. **Maximum Likelihood Estimation (MLE)** - The most likely intensity of physiological change observed in the analysed population. Based on the pdf we have estimated in the previous section, this value would be the $\Delta'z$ that has the maximal probability density. If a pdf is normally distributed the Mean and MLE are equal. As we will later see, this is not the case in many of our observations.
3. **Probability of intended physiological change** - We define the intended physiological change as the change intended by the experimental protocols. That is, stress development for the test segment and stress restoration for the breaks. The probability that an expected physiological change will fall within a certain range, $P(\Delta'z \in [a, b])$, is the area under the graph of the constructed pdf function. Therefore, the probability of intended heart and respiration rate changes during the test is the area under the relevant pdf graph for $\Delta'z \in [0, \infty]$. During the break segments, the probability of intended heart and respiration rate changes is the area under the corresponding graph for $\Delta'z \in [0, -\infty]$. The equivalent metrics for HRV are derived identically by flipping $\Delta'z$'s sign.
4. **Probability of intense, intended physiological change** - We denote an intense physiological change as a change of over 0.5 standard deviations in the physiological signal. We calculate the probability for such a change similarly to the previous metric. During the test segments, an intense physiological change (in the intended direction) is the area under the pdf where $\Delta'z \in [0.5, \infty]$ for heart and respiration rates and $\Delta'z \in [-0.5, -\infty]$ for HRV. Likewise, during the break it is the area under the pdf where $\Delta'z \in [-0.5, -\infty]$ for heart and respiration rates and $\Delta'z \in [0.5, \infty]$ for HRV.

4.3 RESULTS AND DISCUSSION

4.3.1 *Were stress development and restoration elicited?*

First we should assess how well the experimental protocol elicited the intended stress development and restoration. For intuitive analysis, we ask the reader to refer to the physiological samples presented in Figure 17. These heart rate, respiration rate and HRV signals were

collected from Subject 6 when she was exposed to the scene Shibuja. Note that all three signals clearly respond to the experimental segments. We observe heart and respiration rates building up during the test segment and slowly abating during the break; HRV sharply decreases during the test and recovers during the break. We see additional moderate stress development in all signals shortly before the break ends. We postulate this is a result of the occupants' anticipation to the Necker test, which immediately follows the break period.

Though per participant the physiological signals followed distinct patterns, we report general trends common to all of them. That is, increasing heart and respiration rates and decreasing HRV during the test; Inverse trends were observed during the breaks. Interestingly, individual stress development and restoration patterns repeated across sessions. In our previous example following Subject 6, we see resemblance in the signals' patterns, when Subject 6 is exposed to other scenes, such as Forest, Kites and Neutral. This indicates the dominance of the experimental protocol over the scenes in affecting stress development and restoration. It alludes that scene effects should be assessed by inspecting differences in individual responses.

For a more formal analysis, we aggregate session data from all our participants, in all scenes, and partition it by the test and break segments. Figure 21 juxtaposes the histograms of the two partitions for heart rate, respiration rate and HRV Z-score readings. The bottom row similarly presents the histograms for corrected physiological changes $\Delta z'$ - following the signal processing approach we developed in the previous section. For convenience, the histograms are normalised, and the partition means are marked by vertical lines.

From the bottom row, we note the same trends in physiological change in the aggregated data. That is, on average we see increasing heart and respiration rates during the test, and decreasing heart and respiration rates during the break. The opposite trends appear for HRV. Interestingly, the Z-score readings for HRV (reading level rather than change) are on average lower during the break than during the test. This may indicate that the break was insufficient to induce full recovery. It may however also be explained by the inherent signal delay, caused by the Bioharness' built-in HRV extraction algorithm. We use one-way ANOVA to establish the differences in significance between each pair of test and break partitions. The F-test results are attached under each graph, with $n=10,075$ for all pairs. As intended, we found the test and break conditions to be significantly different ($p<0.05$) in all cases.

We additionally investigate if there are possible correlations between the scene ordering and the computed stress development and restoration metrics. This allows us to assess the sequential nature of the experimental protocol, lining the sessions one after the other. Such correlations, possibly stemming from exhaustion or habituation,




Figure 21: Histograms comparing stress development and restoration between the test and break conditions. The top row illustrates Z-score levels for each of the physiological signals. The bottom row illustrates corrected delta scores ($\Delta z'$), representing physiological changes.

should be informed alongside our stress development and restoration results. Differences, which align with such correlations, can be explained by either the scenes or the ordering interchangeably – weakening possible conclusions. We compute 24 pairwise correlations between the scene sequence number and our four stress metrics, for each of the physiological signals, during both test and break segments. We use the Spearman rank-order method [8] to calculate the correlation coefficients and test for statistical significance. This non-parametric method is better suited for datasets, which are not normally distributed. It describes monotonicity relationships between random variables with coefficient values between -1 and 1, like Pearson’s correlation.

We found three statistically significant ($p < 0.05$) correlations in our data set. During test segments we report inverse correlation of -0.31 ($p = 0.001$) between the session number and the probability of observing an increase in respiration rate. Likewise, we report an inverse correlation of -0.2 ($p = 0.03$) between the session number and the probability of observing an intense increase in heart rate (above 0.5 standard deviations). A third inverse correlation of -0.24 ($p = 0.03$) is found during break segments between the session number and the probability of HRV recovery. Intuitively, these correlations indicate that as the sessions progressed the effects of the experimental conditions attenuated. That is, as participants took more sessions, stress developed less in the test segments and we observed less restoration during breaks. We take these correlations into account in the next section, when presenting individual physiological responses to the scenes.

4.3.2 Individual Physiological Responses

Next, we evaluate individual responses for each participant in the presence of the rendered scenes. We follow the data processing approach described in the previous section and present stress development and restoration through four metrics: MLE and mean of physiological changes, probability of restoration or arousal and probability of intense restoration or arousal. For each participant we obtain the metrics separately for the test and break segments.

Figures 22 to 24 illustrate the MLEs and means of physiological changes per participant. The results are presented for each of the physiological signals separately. Each coloured bar shows the MLE of physiological change under the corresponding scene. On each bar we overlaid the mean physiological change with a small black line. The disagreements between the MLE and mean demonstrate the variance we observe in physiological changes and their non-Gaussian distribution.

For example, please refer to Figure 22, which illustrates heart rate changes results. From the top graph, Subject 4 is most likely to exhibit an increase of 0.66 standard deviations in her heart rate when exposed to the scene Shibuja, in a 20 seconds period during a test. From the bottom graph, she is then most likely to exhibit a 0.48 standard deviations reduction in her heart rate if exposed to the same scene during the break. We point out that these values represent the MLE of the LIV corrected physiological change rather than absolute changes. Therefore, increase of 0.66 signifies an even greater increase beyond Subject 4's natural tendency to regulate the test. This should be kept in mind for values, whose sign contradicts what one would expect for a particular segment. Accordingly, when Subject 4 is exposed to Sunset during the break, she is most likely to change her heart rate by positive 0.1, only relative to her natural regulation. Her absolute heart rate might still decrease during the break, but to a lesser extent than what she naturally exhibits.

Figures 25 to 27 illustrate the probability for stress development and restoration during the test and break respectively, per participant. The results are presented for each of the physiological signals separately. As mentioned, this metric does not take into account the magnitude of change, but only its direction. Within each figure, the top set of coloured bars shows the probability of stress development during the test, i.e. a positive change $\Delta'z$ for heart and respiration rates and a negative change $\Delta'z$ for HRV. Similarly, the bottom set of coloured bars shows the probability of stress restoration during the break.

The different scenes are indicated by the different bar colours. We use the colour intensity to overlay the probability of an intense response $P(|\Delta'z| \geq 0.5)$, out of the overall probability for a response in




Figure 22: MLE of the intensity of heart rate change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break)




Figure 23: MLE of the intensity of respiration rate change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break)




Figure 24: MLE of the intensity of HRV change, per participant, in the presence of the different scenes (top: during Test, bottom: during Break)




Figure 25: Probability of heart rate change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense heart rate change (>0.5 std). Asterisks indicate statistical significance.




Figure 26: Probability of respiration rate change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense respiration rate change (>0.5 std). Asterisks indicate statistical significance.




Figure 27: Probability of HRV change in the intended direction, per participant in the presence of the different scenes. Top bars indicate the probability during Test and the bottom bars during Break. Darker colours indicate probability of intense HRV change (>0.5 std). Asterisks indicate statical significance.

the desired direction. For example, in a Neutral scene and during the break, Subject 8 exhibited a heart rate decrease in 63% of the measured 20 seconds intervals. An intense decrease of over 0.5 standard deviations was observed in 33% of the intervals – about half of the times a decrease was observed.

The graphs clearly indicate that the physiological response to the rendered scenes is extremely personal. Personal differences are noticeable in both the intensity of the physiological change and its probability to be in the desired direction. Subjects 4 and 5, for example, displayed similar heart rate recovery probabilities when exposed to either the Forest scene during the break (0.66 and 0.64). However, the intensity in which they responded to the scene was very different. The MLE change in Subject 4's heart rate in the Forest scene was almost 6 times larger than the one Subject 5 exhibited (0.3 vs. -0.05 standard deviations).

Even within participants, physiological responses to the same scene vary between the test and break segments. Subject 4, whose heart rate highly increased in the presence of the Kites scene in the test (+0.98 MLE and 89% increase probability), displayed only moderate heart rate recovery in the break (-0.1 MLE and 69% decrease probability). Across many participants and scenes, we see that high stress development intensity or likelihood during the test does not necessarily result in the inverse response during the break.

Individuals have also reacted differently in their physiological signals. That is, changes in heart rate, respiration rate and HRV did not necessarily match, even under the same rendered scene and task. Subject 1's heart rate for example had high likelihood to increase with the Kites scene during the test and a moderate one with the Forest scene (0.84 vs. 0.59). Her respiration rate on the other hand displayed the inverse, i.e. high likelihood to increase with the Forest scene and moderate one with Kites (0.90 vs. 0.57). This is also the case for a similar comparison of the intensity of change for other participants.

Nonetheless, we report statistically significant differences in the recorded physiological responses. Since some of our data is not normally distributed, we utilised the non-parametric Kruskal-Wallis test [48] to evaluate if the physiological changes under the various conditions are statistically significant. For each participant we conducted pairwise comparisons between the physiological changes recorded in the Neutral scene and the ones recorded in each of the other scenes. Comparisons were conducted separately for the test and break data. Differences were considered statistically significant if $p < 0.05$. We mark the populations that were statistically significant with an asterisk on top of their corresponding bar in all figures (22 to 27). We note that differences that align with correlations previously found between the session sequence number and the restoration metrics

should be regarded with caution. These possibly stem from both the scene and the order in which it was presented.

Note that the vast majority of the scenes exhibited statistically significant differences. This is an important finding in the context of our work. It suggests that the Digital Cubicle can elicit different physiological responses through digital transformation. This is of particular interest for stress restoration. Subjects 1, 2, 3 or 6, for instance, can supposedly use the Digital Cubicle to increase their likelihood of heart rate recovery by almost 30%, rendering the Forest scene instead of the Neutral one during breaks. Likewise, Subjects 5, 7 and 8 can render the Rotch scene to reduce the probability of their HRV decreasing in about 25% during tests. If such effects turn out to be long lasting and resistant to habituation, they can potentially be leveraged through smart applications to target chronic stress and its related illnesses.

4.3.3 *Physiological Responses Generalised per Scene*

Next we describe the general physiological response produced by each scene. Such a description is of course very limited from the data we have obtained. Particularly in light of the relatively small number of participants in our study (N=9), and the personal nature we see in the recorded results. We continue to do so, however, in an attempt to learn about the choice of scenes we have made in the experimental design.

For each scene we aggregate the observed physiological changes from all our participants. Through standardisation and the Law of initial values, these are corrected for absolute values and personal regulation tendencies, and thus are relatively proportionate. We aggregate participants' data for each of the physiological signals separately and, as before, partition it by the test and break segments. We then use KDE to construct the probability density function of physiological change in a 20-second interval, under each scene. Figures 28 to 30 present the resulting probability density functions. For each physiological signal, the top and bottom rows illustrate the test and break segments respectively. The coloured area under the graphs represent the likelihood the scene had to result in a change in the projected direction, i.e. stress development during the test and restoration during the break. We present each scene next to the "Neutral" scene to allow straightforward comparison.

Our probabilistic presentation allows us to gain insight into the agreement in our data about the physiological change under every scene. The more narrow and spiked the pdf is, the more systematic and likely the physiological change was. Look for example at the pdfs for HRV recovery during the break for the Kites and Shibuja scenes. Next to the Neutral scene, they both present higher chances




Figure 28: Probability density function of heart rate change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.




Figure 29: Probability density function of respiration rate change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.




Figure 30: Probability density function of HRV change for each of the scenes (top row: during Test, bottom row: during Break). Coloured area indicates probability of change in the intended direction. Vertical lines mark MLE.

of HRV recovery, signified by the larger coloured areas under the graph. However, their relatively flat and wide shapes indicate high variance in participants' physiological reactions to them.

Contrarily, the Forest and Sunset scenes exhibit relatively high agreement in responses across most of the physiological signals and segments. This makes them good candidates for future research, when there is a need to select scenes with higher "physiological consensus". We additionally report generally higher agreements in responses during the test segment. This may indicate that participants tend to experience the scenes more similarly as stressors than as a source of restoration.

Narrowing our discussion to restoration, we note that in comparison to the Neutral scene, the Forest, Sunset and Kites scenes resulted in higher chances of heart rate recovery during the break. Higher chances of respiration rate recovery were observed during the breaks only in the Rotch scene. Finally, improved HRV recovery was noticed in the Kites and Shibuja scenes. These findings that the physiological signals respond differently, coincide with our previous observations in the per-participant analysis.

4.3.4 *Focusing capabilities*

To get a fuller image of the scenes' effects, we complement our stress development and restoration analysis with an assessment of participants' performance in the presence of the rendered scenes. This is an important dimension of the Digital Cubicle evaluation, since the optimal scene choice does not solely depend on stress development and restoration. Accordingly, when using music as stimuli, [47] have systematically shown, that excessively relaxing stimuli can harm one's task performance; overly arousing stimuli was shown to cause distraction. We expect similar situations with the Digital Cubicle. In fact, in preparation for the experiment we explored attempts to direct attention with a larger set of possible scenes, and have informally experienced severe performance reduction in the presence of overly busy ones.

In the scope of this work we assess participants' performance only through the Necker test. We avoid using the reading comprehension test, due to high variance in participants' individual scores between the different sessions. The limited sample size of our study, does not allow us to cancel differences in the reading comprehension tasks' difficulty.

To use the Necker test as a measure of one's ability to direct attention, we first standardise the results per individual. For each participant we calculate the mean and standard deviation of the cube orientation reversals during the test from all the different sessions. We then transform each sample to its Z-score value, representing its




Figure 31: Necker Z-score per participant, for each of the rendered scenes. Rightmost bars are the average values from all the participants.

distance from the mean in units of standard deviation. This standardisation preserves the meaning of the Necker score, i.e. the higher the score is, the lower one's capacity was to direct his or her attention. Figure 31 illustrates the Necker Z-score results for each participant in the presence each of the scenes. The rightmost column in the graph presents averaged results from all of the participants.

Note that for 5 out of the 9 participants the Shibuja scene was among the most distracting scenes. This coincides well with what we anticipated, selecting it during the experiment design, and attention literature regarding the effects of high movement and tumult. For a different set of 5 participants, the Rotch scene was among the least distracting scenes. The latter is a surprisingly encouraging finding when juxtaposed to the Necker Z-scores in the Neutral scene. It suggests that one's capacity to direct attention can in fact be increased in the presence of stimuli rendered through the Digital Cubicle. This bolsters our design choices in the cubicle's implementation and supports the balance we hoped to achieve between the occupant ability to focus and scene immersion.

Finally, we investigate possible correlations between the Necker scores and the calculated stress development and restoration metrics. We use the Spearman rank-order test and consider $p < 0.05$ for significant correlation. We report a correlation of 0.33 ($p = 0.01$) between the Necker Z-score and the probability of intense HRV recovery during the break segment; an inverse correlation of -0.28 ($p = 0.03$) is found between the same measurements during the test segment. This intuitively suggests high arousal levels measured by HRV in the presence of distracting scenes. Encouragingly, our results show that this correlation is not necessarily a limiting factor for our cubicle. The relatively low correlation coefficient points out there is enough room to find scenes with a balance between restoration and distraction.

4.3.5 *Subjective Perception of the Scene*

In this last section we analyse how the rendered scenes were perceived based on the subjective reports our participants submitted at the end of each session. First we wish to corroborate that the scenes we chose were indeed perceived the way we intended. That is, the scenes were perceived differently in both their restoration and focusing qualities and according to how we labeled them. We then continue by examining possible correlations between the scenes' perception and our previous results for both attention direction performance and stress development and restoration.

Figure 32 illustrates on a two-dimensional plane how the scenes were perceived in their restorative and distracting qualities. The presented results are the averaged scores of the first and second survey questions which ask about the scene's general suitability for taking




Figure 32: Self-reported restoration and focusing ratings for the rendered scenes

breaks and directing attention respectively. The error bars represent the standard deviation in these reported measures. The sample size for these results is relatively small and consists of 9 data points (one for each participant) per scene.

Despite high variance we note that some scenes were indeed perceived distinguishably different in their focusing and restoration qualities. In agreement with attention restoration literature, nature scenes (Forest, Kites and Sunset) were perceived better for taking breaks and restoration than their urban counterparts (Neutral, Shibuja and Rotch). It should be noted though that scene rating for Shibuja and Rotch shows high variance in this respect. Nonetheless, the same pair of scenes was considered discernibly different in its suitability for directing attention. Overall the scenes' ratings satisfy our experimental design intentions. We find the set of selected scenes to represent diversity in the aforementioned dimensions.

We find statistically significant pairwise correlations between the scenes' perception and the Necker scores. To obtain the correlations we used the Spearman's rank-test on pairs of participants' Z-scored Necker results and their reported perception. The data pairs were taken from all the participants and all their sessions ($n=54$). The strongest correlation (-0.35 , $p=0.0001$) was found between the scene's perceived focus and the Necker score obtained in its presence. This correlation indicates that participants were able to correctly gauge their capacity to direct attention in the presence of a scene. Weaker, yet significant, correlations (~ 0.2 , $p < 0.05$) were found between the Necker scores and the scene restorative facets: Being Away, Fascination, Scope and Co-

herence. The higher the perceived restorative facet was the less participants were able to direct their attention.

We similarly calculate pairwise correlations between the scenes' perception and the stress development and restoration metrics we obtained. Participants' reports on Perceived Restoration, Compatibility and Fascination were found correlated with probability for intense HRV stress development during test segments. Calculated correlation coefficients were 0.35 ($p=0.02$), 0.3 ($p=0.03$) and 0.29 ($p=0.03$) respectively. These indicate that increased probability of stress development was observed during the test in scenes which were perceived more restorative, more fascinating or better matched occupant liking.

During the break segments, more significant pairwise correlations are found between facets of perceived stress restoration and both heart rate and HRV recovery. Table 1 summarises these results. For example, we found that the more the scene provided an escape experience (Being away), the higher the probability was for heart rate recovery during break. Likewise, the more occupants liked the scene (Compatibility), the more intense were their heart rate reductions. In a similar manner, scenes that were perceived to provide opportunity for exploration and discovery (Scope), resulted in more intense HRV recovery. It is interesting to note that correlations were found for both recovery probability and intensity (MLE). Due to high disagreement among the participants on the restorative qualities of the Rotch and Shibuja scenes, we conduct an additional calculation for a data subset excluding them. This calculation reveals additional correlation between perceived restoration and respiration rate recovery during the break. Interestingly the correlation that is found is inverse, i.e. respiration rate tended to develop more (and more intensely) during the breaks when the scene was perceived to be more restorative.

Perception	Phys. response		Correlation	
	Signal	Metric	Coef.	p-value
Perc. Restoration	HRV	MLE	0.29	0.03
	Heart rate	MLE	-0.32	0.02
Compatibility	HRV	MLE	0.31	0.02
	Heart rate	MLE	-0.45	0.0009
	Heart rate	Restoration Prob.	0.33	0.01
Being away	Heart rate	Restoration Prob.	0.33	0.02
Fascination	HRV	MLE	0.28	0.04
Scope	HRV	MLE	0.31	0.02

Table 1: Significant correlations ($p<0.05$) found between perceived stress restoration and physiological recovery during breaks.

4.4 SUMMARY

To conclude this chapter we summarise our main findings. These will guide the development of an intelligent control agent in the next chapter.

- We established significant differences in stress development and restoration between the test and break contexts. These manifested through the metrics we have developed and apply to all physiological signals – heart rate, respiration rate and HRV.
- In all metrics, the physiological responses to the rendered scenes are highly personal. Even within participant, in the presence of the same scene, we note that the different physiological signals do not necessarily correlate. That is, different participants may exhibit stress development and restoration in one physiological signal and not in another (e.g. heart rate, rather than respiration rate).
- Nonetheless, we observed statistically significant differences between the scenes in the recorded responses. This suggests that the Digital Cubicle can elicit changes in occupant physiological response through its digital manipulation of the physical characteristics.
- Aggregating our results per scene, we note that some scenes acted better than others as a source of stress restoration. For example, occupants in the Sunset, Forest and Kites scenes exhibited higher likelihood of heart rate recovery than when they were in the Neutral scene. A larger sample size however is required to generalise such results.
- We additionally report differences in occupants' ability to direct attention under the different scenes. These were found correlated with occupants self reported perceptions. Similar correlations were found between the scenes' perceived restorative qualities and their physiological stress development and restoration measures. These indicate occupant ability to gauge scene effects.

We have so far attained only a partial demonstration of the full vision we outlined in the introduction: We presented a workspace prototype that can digitally manipulate many of its physical characteristics. Then, through a set of user studies we exemplified how it can affect aspects of occupants' perception and experience. Focusing on stress development and restoration, we identified how such aspects can be sensed through the occupants' physiological signals. Our vision, however, is not complete unless we close the loop – leverage the captured data to autonomously and intelligently control the physical characteristics of the workspace.

In this chapter, we propose and evaluate a machine learning approach for building an intelligent control agent for our workspace. We begin by refining and scoping down our problem. Next, we point out the four main design principles leading us to approach our problem through a Reinforcement Learning framework. In this framework, we develop a full algorithmic solution in Sections 5.3 through 5.6 and describe a working implementation we built in Section 5.7. To conclude, we evaluate our intelligent agent through a set of simulations, anchored by real human subject data.

5.1 PROBLEM STATEMENT

Finding the optimal approach to control the characteristics of the workspace is a complex and multi-layered problem. It grows particularly difficult due to the many dependencies involved and the high dimensionality that characterises each of the variables. The workspace itself provides ample opportunities: from changing merely the colour temperature, to expressing a whole different narrative through visual and auditory content. The sensor framework, likewise, produces a rich real-time dataset, which is open for broad interpretations. It can be used either directly or in combination, to infer many high-level constructs about the occupant. The occupants are, of course, impractical to fully model. They vary, in time and between each other, in their preferences, intensions, objectives, affective states, capabilities, etc.

We therefore begin by narrowing down our problem according to the following guidelines:

1. **Control objectives**

We set the control objective of the intelligent agent to optimally foster stress restoration and lower stress development. We are

encouraged to do so by our findings in the physiological study presented in Chapter 4. On the one hand, these suggest our workspace can produce significant differences in stress development and restoration. On the other hand, the effects that were found were highly individual, stochastic rather than deterministic, and non-straightforward to apply. These characteristics rule out naive control algorithms and make it an interesting optimisation problem for a smart agent.

2. Stress estimation

For simplicity, we define stress development and restoration, as described in Section 4.2. We look at direct physiological changes measured by our sensors, and regard the sign of physiological change to indicate development or restoration. The metrics we developed in the same section are used to quantify the workspace effects. In Section 5.5 below, we discuss how this simplistic model can be extended to include more accurate models for stress estimation. For further reading on stress estimation techniques, we refer the reader to [75].

3. Content rendering

We scope our problem further, by significantly limiting the content that can be rendered by the agent. We program a fixed set of 6 predefined scenes from which the agent can choose, and restrict the agent from natively controlling the workspace components. How to utilise (and even map) the full potential offered by our digitally controlled physical characteristics is a complex problem that is addressed outside the scope of this work. We refer the reader to [97] for further reading on this subject.

5.2 DESIGN PRINCIPLES

When designing the intelligent agent we wish to take into account learnings from our previous experiments. We shape those into four design principals, which will hereafter guide the majority of our implementation choices.

1. Personalisation

The agent must acknowledge personal differences. Throughout our studies we witnessed such differences both in individual perception, performance and physiological response. Under the limitations of our sample size, we could not identify similarities or response patterns that allow clustering or profiling. We, hence, must assume that the agent's recommendations should be tailored to individual occupants.

This is a particularly relevant design principle for any learning process the agent may apply. It emphasises the agent's need

to focus on learning from observations rather than searching for similarities. An immediate tradeoff of such choice is bootstrapping – prolonging the initial time before the agent learns to make effective recommendations on its own. We believe, however, this is an appropriate tradeoff for our system. As a workspace the Digital Cubicle is intended to be routinely used, for long periods of time and by a limited set of individuals.

2. Time and context

Note that the Digital Cubicle was experienced many times differently due to changes in context. The most obvious is the differences we saw between responses, when taking a test or a break in the physiological study. Perceptually, participants have also described their experiences differently according to what phase of their work they wish to achieve (creative thinking vs. focused execution), and what stage they were at during the day. These effects take place regardless of affective state or mood and depend mainly on occupants' intentions and objectives.

We also identified a dependency between the occupants' experience and the order of scenes to which they were exposed. In the first perceptual study we saw through free subjective reports how an environment experience was many times described in relation to its predecessor. In the physiological study, we saw correlations between the scenes' order and physiological response.

We combine these findings and propose that our smart agent should output a sequence of recommendations rather than discrete ones. Using machine-learning terminology, this translates into an accumulative loss function the algorithm should optimise for in making its predictions. Note that this choice relaxes the agent's need for context awareness but does not necessarily satisfy it. Instead, it allows the agent to assume there is latent context. As such, the context effects will be averaged on the entire sequence and, more importantly, if the context follows a temporal pattern, the agent will have the opportunity to learn it.

Any information that can be gathered about the occupant's context would still benefit the agent's ability to make effective recommendations. Context recognition in workspaces, particularly for wellbeing and productivity, is a live and ongoing field of research. We therefore intentionally deemphasise it in our proposed solution and focus our efforts on the many remaining aspects of our problem. In Section 5.4 we mention how our solution can be adapted to incorporate relevant breakthroughs in context recognition.

3. Occupant preference and interaction

Autonomous as the agent may be, it must still conform to occupant requests and respect their preferences. In fact, our studies suggest that the occupant preferences are informative cues from which the agent can learn. Corroborating this idea, we found significant correlations between occupant perceived restoration and corresponding physiological responses. Occupants were likewise successful in ranking the scenes in which they would be able to direct attention.

4. Probabilistic modelling of physiological response

We note that occupants' physiological responses behave stochastically. This was recorded early in the physiological study and directed our approach to describe physiological changes probabilistically. Accordingly, our agent is required to handle stochastic physiological measures. This requirement is particularly important as physiological measures are the ones on which the agent's control objectives are based.

Machine learning algorithms are natural candidates for handling such a requirement. Their underlying statistical models make them robust to handle data probabilistically, and optimise for expected outcomes.

5.3 APPLYING A REINFORCEMENT LEARNING FRAMEWORK

Based on the design principles above we chose a Reinforcement Learning framework [40] to approach the agent's design.

Reinforcement learning is a sub field in machine learning, which aims to train an agent what actions it should take to maximise its long-term rewards. The learning process happens through sequential steps and usually within episodes, in which the agent interacts with its environment. In each step the agent takes an action that transitions the environment from one state to another and receives a reward. The agent's goal is to learn a policy, which defines what actions to take in every state in order to maximise the accumulative rewards by the end of an episode.

A key component, which makes Reinforcement Learning methods powerful, is their ability to deal with exceptionally large environments. They tend to require no a priori knowledge about the environment, and collect information about the environment through interaction. Unlike supervised machine learning techniques, they do not require pairs of labeled data to begin learning, which are particularly challenging to collect in the context of our work. In that sense, Reinforcement Learning methods are well tuned to online performance. They learn from data when it becomes available, managing

a trade-off in their decisions between exploiting accrued knowledge and exploring uncharted branches of the environment.

Counter-intuitively, the environment in our case is the workspace occupant. Since we stated our problem as finding an optimal sequence of scenes, we can express it using sequential decision-making terminology. We use Markov Decision Processes (MDP) [68] for such formulation. This will later allow us to borrow tools from Reinforcement Learning literature to learn an optimal sequence on the fly, treating the occupants as model-free and training the agent through the occupant interactions.

MDP is a discrete-time stochastic control process that extends the traditional Markov chains by adding primitives for actions and rewards. It is formally defined by a set of 5 tuple (S, A, P, R, T) :

1. S is the finite set of states, which includes a seed state s_0 .
2. A is the finite set of possible actions.
3. P is a state-transition probability function $P(s, a, s') \rightarrow p \in [0, 1]$. It denotes the probability of transitioning from state s to s' when taking action a .
4. R is a reward function $R(s, a) \rightarrow r$. It denotes that taking action a , from state s , yields a reward r .
5. $T \subset S$ is the set of terminal states, which end an episode.

Once defined, solving the MDP would be finding a policy π . This policy is essentially a function $\pi(s) \rightarrow a$, deciding what action the agent should take at any given state. The optimal policy is commonly denoted as π^* . It is the policy that, if followed, yields the highest expected sum of rewards during the length of an episode.

5.4 MODELLING ACTIONS AND STATES

The actions are the decisions the agent should make and has control over. These are fairly straightforward in our case. They represent transitioning the scenes rendered in the workspace from one to another. We accordingly denote the set of actions A to be the set of all possible scenes.

$$A = \{a \mid a \in \{\text{Neutral, Forest, Rotch, ...}\}\} \quad (8)$$

We treat our problem as a discrete decision making process. That is, we define a fixed time interval Δ in which decisions are made. Once a decision is made the resulting scene is rendered in the workspace for the remaining duration of the time interval. When the interval ends, a new decision is made and the room may transition its appearance.

Transitioning to the Neutral scene is regarded as another possible action, like any of the other scenes.

Modelling the states is a more challenging task. Ideally these should encompass any information relevant to determine the occupant's reaction to the scenes. Such information is many times classified into the occupant's context and state [66]. Both of which are, of course, infeasible to fully model. The occupant's context covers many aspects, including the tasks at hand, the tools used to complete them, particular cognitive requirements and acquired proficiency. The occupant's state similarly bundles diverse factors such as affective state and mood, current physiological conditions and even personality. Each of these factors by itself contains an almost boundless set of possible options.

We therefore offer to flip the problem and model the occupants' states by the sequence of scenes they experience. A similar approach was proposed by [52], designing an emotion-based music playlist recommendation system. The rationale behind it is that since the agent is limited by its set of actions, the best it can do is assume the occupant fully reacts to them. It should recognise that its effect on the occupant experience depends on the full sequence of its actions (or up to k steps back in history). A state $s \in S$ is thus an ordered tuple of rendered scenes ranging from 0 for the first scene to k for the last scene in the episode. Formally we define:

$$S = \{(a_1, a_2, \dots, a_i) \mid 1 \leq i \leq k; a_j \in A, \forall j \leq i\} \quad (9)$$

Though simplified, such a representation still has enough expressive power to capture some latent variables determining the occupant reactions. These will manifest as stochasticity in the reward function, i.e. taking the same action a in the same state s may not always result in the same reward. In other words, our agent recognises that rendering the same scene in the same place in the sequence depends on other factors it has no control over. The Reinforcement Learning technique our agent will apply is robust enough to handle such cases; it optimises for the expected accumulative reward, instead of optimising for the reward in a particular episode.

Moreover, latent variables that have structured temporal dependency can, by proxy, be fully expressed through the scene sequence. Consider a hypothetical case of a call centre employee whose work is divided into three well-defined phases; Within a 90 minutes work episode, the employee first answers calls for 50 minutes, then files paper work for 20 minutes and then takes a 20-minute break. If these determine the employee's reactions, the employed Reinforcement Learning technique will have the agent learn an optimal scene sequence that expresses this temporal structure. That is, even without knowing it a priori. To collect maximal reward, in the first 50 minutes the agent will learn to render scenes matching the employee's preference




Figure 33: States and actions model for the intelligent agent

for answering calls. For the next 20 minutes it will play scenes matching paperwork, and finally it will render scenes for taking a break. In this regard, the state formulation we propose is well suited for learning an occupant's unknown temporal working routine. These exist in many real-world scenarios and make a compelling argument supporting our approach.

Figure 33 summarises the states and actions model for our agent. We begin in state $s_0 = (a_0)$, representing the seed scene a_0 , that was chosen by the workspace occupant. Based on what the agent has learned in previous episodes, it selects the next scene a_1 and changes the current state to $s_1 = (a_0, a_1)$. After time interval Δ in rendered scene a_1 , the agent collects its rewards, learns from them and continues to take action a_2 . Note that these definitions of S and A induce a deterministic transition function P , which defines a single possible state s' after taking action a in state s .

Occupant interactions add some subtleties worth considering. We allow the occupant to override the agent's decision at any time using two primitives: either selecting a specific scene to transition to, or skipping the current scene and signalling the agent to choose a different one. We export these primitives to a simple web-based controller placed on the workspace desk.

Using the state-action terminology we established, when a scene is skipped the agent is simply forced to take its next action earlier; When a specific scene is chosen, the agent considers the occupant's decision as its own and proceeds regularly. Both cases are leveraged as opportunities for the agent to learn and are taken into account in our rewards model. The only question remaining is whether the agent's state should be advanced by both the overridden scene and the occupant selection or only by the occupant selection. For our agent, we implement a threshold configuration parameter $\Delta_{\text{threshold}}$ to determine this. The overridden scene introduces a state transition if it was played for more than $\Delta_{\text{threshold}}$, and it is excluded from the state transitions otherwise.

The particular choice of rendering time Δ , maximal sequence length k , and the number of possible scenes $|A|$, directly affects the size of the state space S , and in turn may significantly impact the number

of episodes needed for learning convergence. The tradeoff between their value and convergence time can be optimised for specific scenarios through simulation. Such a simulation is demonstrated in Section 5.8.

For our agent we limit the sequence length to 6 scenes rendered for 10 minutes each, resulting in an hour working episode. We allow 6 scenes to be rendered in our workspace, which sets the size of the state space to $6^6 = 46656$ states. Though large, we find this space size can be feasibly solved by many Reinforcement Learning algorithms. In future development, the state space size can be reduced by applying an N-Gram model: limiting the state to remember only the N previous scenes rather than the full sequence. This technique is commonly used in natural language processing and has also been demonstrated in music playlist generation by [12]. Another technique is to apply dimensionality reduction (which is relevant to the rewards model) to the set of scenes. For example, reducing the set of scenes to a two dimensional set such as "quiet" and "loud". These labels can be used in the sequence and as actions instead of the full set of scenes. The agent should also have a matching selection function. When the agent decides the next scene is "loud", the function will select, for instance randomly, an actual scene with a "loud" label.

The proposed state-action formulation should be revised in cases where aspects of the occupant's state and context are: (1) well modelled, (2) can be sensed, and (3) have higher influence on expected rewards. Imagine a hypothetical case where an occupant's fatigue influences the rewards significantly more than the sequence of rendered scenes. Assuming we can produce a set of sensors to effectively capture fatigue, the states model should be modified accordingly. A simple two states model representing "occupant alert" and "occupant tired" may yield higher cumulative rewards, as it will better characterise the rewards function. Alternative state models may still maintain a sequence structure to express temporal dependency. The labels within the sequence, however, should be modified to come from a set that better represents context than the set of scenes.

5.5 DESIGNING THE REWARDS FUNCTION

The reward function defines what the learning algorithm optimises for, and thus in many ways determines the agent goals. Except for minor limitations, the function can be freely designed to capture almost any measurable objective. The real-time sensor data collection infrastructure integrated into the Digital Cubicle, provides rich measurements supporting the design of a wide range of reward functions. For example, one can design a reward function, which grants rewards according to the occupant's posture as measured by the Bioharness accelerometer; The longer the occupant holds an upright pos-




Figure 34: Implicit reward function based on occupant interactions

ture the higher the reward and vice versa. Alternatively we can design a reward function based on viewing orientation: optimising for the amount of time an occupant spends resting her eyes viewing outside the computer screen, or vice versa.

For this work, we experiment with the design of two reward functions. The first function takes into account the occupants' interactions and optimises for their preferences. The second function relies on changes in physiological signals and attempts to optimise for lower stress development and higher stress restoration. For our particular implementation, we scope down the second reward function to only consider heart rate development and recovery. In later sections we exemplify how the two functions can be used either separately or intertwined.

5.5.1 Occupant Interactions

We take an implicit approach to reward for occupant interactions. We assume that if the agent's chosen scene matches the occupants' preference, they will allow it to be rendered without any interruptions. Hence, the agent is rewarded by the time its chosen scene is rendered in the workspace. We set the maximal reward to 1.0 and grant it if the scene was played all the way till the end of period Δ . The minimal reward is set to -1.0 and is granted when the scene is immediately overridden. To obtain a reward for any playing time in between, we linearly extrapolate between the minimal and maximal values.

The aforementioned reward policy is illustrated in Figure 34. Note that the reward is granted for state transitions, i.e. taking particular action a in state s . When a scene is skipped there are therefore two reward-granting transitions that the agent can learn from. The first (usually negative) reward will be associated to transitioning from scene A to skipped scene B. The second reward can either be associated with transitioning from scene B to C or from scene A to C. We choose the second reward to be associated with transition B to C only

if scene B was played for long enough to be included in the state sequence. Otherwise, we consider the second reward to be associated with transition A to C.

The case when occupants pick a scene themselves is similar to skipping a scene with regards to the overridden scene. The chosen scene, however, is regarded differently. We treat the occupants' choice as an implicit reflection of their preference, and hence add an additional positive term (+1) to the chosen scene reward.

5.5.2 *Physiology-based Rewards*

We design our physiological reward function based on the physiological change metrics we developed in Chapter 4. A different design can similarly be applied to accommodate other methods of stress estimation. We calculate the physiological rewards based on the change in the physiological signal between two, consecutive, 20 seconds time intervals. Specifically, our implementation caches and constantly updates a window worth of 40 seconds of samples. It calculates the mean of the first and second, 20 seconds windows, and subtracts them. This calculation is done asynchronously every 10 seconds, introducing some overlap between the samples. The resulting physiological change ΔZ are then fed to the reward function.

For consistency with our previous analysis, we add a preprocessing capability to our sensor collection infrastructure responsible for real-time standardisation. It takes the raw heart rate samples and converts them to their Z-score representation. The signal mean and standard deviation used for the computation are constantly updated. With this approach, the resulting signal is very attenuated at first. Empirically, however, we see that after the first 7 minutes in the workspace, in which an occupant worked through one of our test-break sessions, the signal settles to regular values.

Our reward function accumulates the values of physiological change throughout the entire time the scene is rendered. It then rewards the scene based on the likelihood of change in the desired direction. Formally it computes:

$$R(s, a) = 2\left(\frac{N_{\ominus}}{N_{\ominus} + N_{\oplus}} - 1\right) \quad (10)$$

Where N_{\ominus} and N_{\oplus} are the number of negative and positive physiological changes ΔZ respectively, that were observed while the scene a was rendered in state s . Like the preference reward function, this calculation results in a reward between -1 and 1.

5.6 SOLVING WITH Q-LEARNING

Q-learning is a model-free Reinforcement Learning technique to find an optimal policy π^* for any given finite MDP. It estimates the optimal policy through sequential steps (resembling dynamic programming), in which the environment is sampled and the estimate is updated (resembling Monte-Carlo Methods). Q-learning is model-free in the sense that it requires no knowledge on the environment's dynamics. Stochastic state transitions and rewards are handled without further adaptations. It has been proven that for any finite MDP, Q-learning will eventually converge to the optimal policy, collecting the highest expected accumulative reward [90]. Its main drawback is that, though computationally efficient and elegant, it tends to require a lot of experimental data to learn.

5.6.1 The Learning Step

Q-learning works by learning an action-value function Q that stores the expected utility for taking any given action at any given state. With each iteration, the algorithm updates the function using the following rule:

$$Q(s_t, a_t) \leftarrow (1 - \alpha_t)Q(s_t, a_t) + \alpha_t(r_{t+1} + \gamma \cdot \max_a Q(s_{t+1}, a)) \quad (11)$$

Where s , a and t are the state, action and reward at time t , respectively. α_t and γ are algorithm parameters ($0 \leq \alpha, \gamma \leq 1$), that tune the learning process.

α is called the learning rate. It is used to balance between the importance of newly acquired information and older one; When α is 1 the agent will only consider new information, and when it is 0 the agent will no longer learn. It is common practice to begin with a non-zero learning rate and decrease along the way. We use the following term to adjust the learning rate:

$$\alpha_t(s_t, a_t) = \frac{1}{1 - \text{VisitCount}(s_t, a_t)} \quad (12)$$

Where $\text{VisitCount}(s, a)$ represents the number of visits of $Q(s, a)$.

The γ parameter factors into the update the expected reward from possible future transitions. The higher the parameter is, the more the algorithm will consider rewards that can possibly be granted in the future; Low parameter values will make the agent greedier, relying more on current rewards. In our implementation, we set γ to 0.01, primarily in light of the immediate nature of physiological change and our long time interval Δ . We would like to avoid, for instance, the hypothetical case where the agent decides to overly arouse occupants, just to it can slowly have them recover in consecutive future intervals.

5.6.2 Choosing an Action

When the agent is asked to choose the next action, it can take one of two strategies: exploiting the knowledge it has already obtained, or exploring for new, possibly more useful, information. To exploit accrued knowledge, the agent should choose, in a given state s , an action a , that maximises $Q(s, a)$. This policy takes into account all previous observations, and is the optimal policy once the algorithm had converged.

The choice of an optimal exploration strategy is less definite and is the focus of much ongoing research. For our agent we choose the ϵ -greedy strategy. That is, when asked to choose an action, exploiting knowledge with probability $1 - \epsilon$, and choosing randomly from all possible actions, otherwise. We use an exponentially decaying value for ϵ , such that:

$$\epsilon_t = \epsilon_0 e^{-\lambda t} \quad (13)$$

Where ϵ_t is the value of ϵ at time t , ϵ_0 is its initial value, and λ is the decay constant. We set ϵ_0 and λ to 1 and 0.001 for our agent. This setting results in high probability for exploration during the first 100 steps.

5.7 IMPLEMENTATION

We implemented the intelligent control agent as an add-on service on top of the Digital Cubicle platform. Figure 35 illustrates the top-level software design of the agent and how it interfaces with the Digital Cubicle services. The agent was implemented in Python and executes in a single process. It is designed to work asynchronously using Twisted [24], applying a Reactor design pattern. For the lack of openly available software package, we implemented the Q-learning algorithm ourselves.

There are three types of asynchronous events that the agent handles:

1. **Timing Events** - An *EpisodeTimer* module uses the Twisted functionality to time the smart agent. It notifies the agent when a learning episode begins and then, until the end of the episode, every $\Delta = 5$ minutes when a new decision needs to be made about which scene to render.
2. **User interactions** - User interactions are sent to the agent through a web socket from a controller on the workspace desktop. These events ask the agent either to skip a scene or to change the rendered scene to a specific one. They are handled both by the *PreferenceRewards* module and the agent's main algorithm. The former translates the user interactions into rewards as described in Section 5.5.




Figure 35: Top-level design of the Intelligent control agent implementation

3. **Physiological samples** - Heart rate samples are streamed to the agent from the cubicle's *Sensor Collection Service* through a local connection. The samples are Z-scored, cached and interpreted into rewards by the *PhysiologicalRewards* module as described in Section 5.5.

The events are handled through the agent's main algorithm. The core of algorithm is described below through a short Python code snippet. The *Qlearn* module implements the Q-learning algorithm, including the ϵ -greedy exploration policy and the adjustable α parameter. A *StatesActionsManager* handles state transitions and maintains the agent's current state and last action. Once the agent decides on the next action to take, it applies it by sending a control command to the *SceneController* service in the Digital Cubicle backend.


```

def SmartAgent_advance(next_scene):
    # Collect physiological and preference rewards
    rewards = PreferenceRewards.collect() + \
              PhysiologyRewards.collect()

    from_state = StatesActionsManager.current_state()
    last_action = StatesActionsManager.last_action()
    to_state = StatesActionsManager.construct_next_state(
        from_state, last_action)

    # Learn based on the collected rewards
    Qlearn.update_q(from_state, last_action, rewards, to_state)

    # Advance the state only if the scene was played long enough
    if EpisodeTimer.scene_exceeded_delta_threshold():
        StatesActionsManager.advance_state(to_state)

    # Select the next scene, unless it was selected by the user
    if next_scene is None:
        next_scene = Qlearn.choose_action()

    # Command the workspace to transition
    StatesActionsManager.update_action(next_scene)
    SceneController.transition_workspace(next_scene)

```

5.8 EVALUATION METHOD

In this section we evaluate our smart agent design through simulations. We revisit the design principles outlined in Section 5.2 and wish to confirm these are all met by the Reinforcement Learning approach we proposed.

To anchor our simulations with real data, we apply the agent's algorithm to a simulated scenario, which is similar to the experimental protocol described in Section 4.1. That is, we configure our simulation to follow a sequential set of working contexts that imitates the sequence of a test and break periods in a session. We leverage the physiological data collected in Chapter 4 to create realistic simulations of possible physiological responses by workspace occupants. Likewise, based on participants' qualitative feedback, we derive a set of possible heuristics, simulating occupants' preferences and interactions.

5.8.1 Implementation

We design and implement a set of software components, allowing simulative execution of our smart agent. Figure 36 summarises these and illustrates how they interface with the agent modules. The com-




Figure 36: Top-level design of the simulation implementation

ponents were designed so that the vast majority of the agent’s original code is executed through our simulations. Their main objective is to simulate possible environments (in the Reinforcement Learning sense) that the agent may interact with and learn from. In our particular case, this translates into simulating the workspace occupant physiological responses and interactions.

1. Episode Simulator

This component drives the sequential episodes, from which the agent learns, and simulates changes in context throughout the episodes. Its implementation defines temporal patterns, if such exist, of context change that the agent needs to attempt to learn. We focus our simulation on simulating working episodes, which are similar to the sessions we have designed for the physiological user study. Though these may not necessarily represent a typical working session, they importantly allow us to bolster our simulation with real data collected from human participants.

The *Episode simulator* simulates the time progression of the simulation. It artificially times the agent, notifying it when it needs to make new decisions, and feeding the *Preference Reward* module with elapsed time within each scene. To fit the relatively short

episode time of a session we adjust the agent’s Δ parameter to 60 seconds and $\Delta_{\text{threshold}}$ to 5 seconds.

With every time step the *Episode simulator* also determines the occupant’s context according to how far it is into the simulated session. To simplify the session we modify its sequence to only include two possible contexts: test and break. We collapse each of the segments in the original experimental design to either one of them. We consider the waiting periods and the survey segment as breaks and the Necker Cube test as test. To shorten the sequence, we additionally reduce the minimal time for context change to a minute (rather than every 30 minutes in the original experiment). The aforementioned modifications result in a simulated session, which is broken down into a sequence of context changes as illustrated in Figure 37.

Simulation Clock:	1 Min.	1 Min.	1 Min.	1 Min.	1 Min.	1 Min.	1 Min.	1 Min.	1 Min.
Context:	Break	Test	Test	Test	Break	Break	Break	Test	Break

Figure 37: Sequence of contexts for a simulated session

We emphasise that the agent is unaware of the context as it is simulated by the *Episode simulator*. From the agent’s perspective the context can only manifest by proxy through the physiological samples and occupant interactions. Through our evaluation, we wish to see that the agent is capable of treating context as a latent variable and learn a policy, which takes it into account. We note that our *Episode simulator* changes the sequence of context deterministically. An interesting improvement to our simulation would be introducing some noise to changes in context and the time intervals in which they change. Our agent should theoretically be robust enough to handle both.

The episode simulator can execute many episodes sequentially order of magnitudes faster than the time they would actually take in reality. Once an episode is over it obtains the agent’s cumulative rewards for the episode and continues to simulate a new session from which the agent can again learn. By following the episodes’ rewards we can identify how long it takes the agent to converge to an optimal policy and what is the policy’s expected reward. It also provides a powerful tool to adjust the agent’s learning parameters, optimising them for quicker convergence.

2. Physiology simulator

This component simulates a physiological signal. With every time step of the simulation, it generates simulated physiological samples, which follow a model of some possible occupant.

This model is built out of a linear function, describing the occupant's innate physiological regulation, and a pdf describing probability of physiological changes on top of it. These are identical to the probabilistic models described in Section 4.2, which we have widely used to analyse physiological responses. In fact we use the models we constructed in our physiological study to simulate 9 different workspace occupants.

With every time step the simulator determines the elapsed time since it has last updated the physiological signal. For every 20 seconds interval within this time, it samples the pdf which models physiological change. It additionally calculates expected change from the occupant's innate reaction, by plugging in the signal's current reading to the linear function. Both changes are then added to the current reading to determine the next physiological sample. These physiological changes are fed to the smart agent's *PhysiologicalRewards* module, where the agent translates them into rewards.

The simulator's pdf describing possible physiological change, can be altered continuously during the simulation. This allows us to update the expected changes in physiology, according to scenes and contexts the occupant is exposed to. For example, when the agent decides to change the rendered scene in the workspace, say from Kites to Forest, the simulation will fetch the occupant's corresponding pdf function and will update the *Physiology Simulator* accordingly. From that moment on, generated samples will use the new pdf, and thus will take into account the new rendered scene. Likewise, the pdf is updated when the occupant's context changes, say from a test to a break.

At the beginning of each episode, the physiological signal is set to a random value between -0.5 and 1.5. This adds realistic uncertainty about the occupant's physiological reading when he or she initiates the working episode. Together with the probabilistic sampling of the pdf, It assures that the simulated physiological signal will differ during the learning episodes. It allows us to evaluate the agent's ability to handle stochastic behaviour in the physiological signals.

The *Physiology Simulator* is designed to support simulating any type of physiological signal. The samples it generates are in units of standard deviations and thus simulate physiological trends rather than actual physiological values. In the scope of this work, however, due to limitations in the *Physiological Rewards* implementation, we only load occupants' pdfs generated from heart rate changes.

3. Participant Simulator

This component is designed to fully simulate a workspace occu-

pant. It consists of three main functionalities: physiology simulation, user interaction heuristics and context dependent logic. As mentioned, occupant simulation is based on real data, collected through our user study in Chapter 4. Therefore, when a simulation begins, the participant simulator is initialised with a Subject ID. According to the Subject ID, the *Participant simulator* loads raw and processed data, which was previously obtained from the subject, that will be relevant for the simulation. This includes the subject's probabilistic models of physiological changes, Necker Cube test scores, perceptual ratings for the scenes, etc. The data is used to guide physiology simulation and the user interaction heuristics.

Physiology simulation is based on the *Physiology Simulator* detailed above. The *Participant Simulator* is responsible to initialize it at the beginning of each episode and update its pdf according to changes in scene and context. The *Episode Simulator* updates the *Participant Simulator* whenever a change in context occurs. Likewise, the *Participant Simulator* is notified on any changes in scene by the *Scene Controller Service*. With the scene context pair, the *Participant Simulator* selects the pdf that best simulates the study participant and updates the *Physiology Simulator* accordingly.

The user interaction heuristics extend the *Participant Simulator* describing occupant behaviours we wish to simulate. For example, we can implement a heuristic, in which every 40 seconds the participant randomly changes the scene in the workspace. The heuristic may be as complex as we wish and rely on participant data and current context – e.g. skipping a scene if its Compatibility rating is lower than zero with probability of 90% during test context and a probability of 40% during breaks. We develop three such heuristics for our agent's evaluation and present them in the next section. User interactions, decided upon by the heuristics, propagate to the *Preference Rewards* module and the main algorithm, where they affect learning and are applied.

5.9 RESULTS AND DISCUSSION

We design three heuristics, representing different scenarios our smart agent should handle. For each heuristic we evaluate two aspects of the agent's performance:

1. **Convergence:** Whether or not the agent converged to an optimal policy, and how many episodes did it take for it to do so?
2. **Correctness:** What is the optimal scene sequence the agent has learned and is it indeed the optimal one according to the underlying simulation models?




Figure 38: Rewards progression throughout simulations of Heuristic-1

For this section, since the physiological study data for the Sunset scene is incomplete, we remove it from the agent’s possible actions. This should not affect our evaluation in any way.

5.9.1 Heuristic 1: Learning from occupant preferences

We begin by implementing a simplistic heuristic for occupant preference. In this heuristic the occupants have a single preferred scene, which they adhere to. Whenever the agent selects a scene, which does not match their preference, the occupants will, in 90% of the cases, change the scene to their preferred one after 10 seconds. For this example, we use Subject 1’s data to model physiological response and choose Forest as the preferred scene due to its high Compatibility rating.

Figure 38 illustrates the simulation results. The graphs indicate the preference, physiology and total accumulated rewards that the agent collected for each episode of the simulation. For clarity, the results have been smoothed using a 5 sample moving average filter. We executed the simulation multiple times with different exponential decay parameters for the ϵ -greedy exploration strategy. The results from the different executions are indicated by colours.

As expected the preference rewards converge to a horizon value of 8.33. We report that the policy the algorithm converged to is one where the agent consistently selects the Forest scene. These encouraging results exemplify the algorithm’s capability to correctly learn the occupant’s preferences. Our preference rewards model correctly identified the occupant’s latent preference and steered the agent towards the theoretically optimal policy in this case.

We note that the accumulated physiological rewards (illustrated in the middle graph) are not as dominant in the total rewards as the preference rewards. This observation is the result of the physiology rewards model we have defined. The probabilistic nature of the phys-

iological signals (and accordingly their simulation) makes it unlikely for them to win full point rewards. Changing the models' full reward values can easily modify the balance between the preference and physiology rewards. We choose however to keep the balance as it is, favouring occupant preferences over physiology.

In this scenario, we obtain relatively short convergence times. Depending on the value of the decay parameter, the agent needed as few as 10 episodes to learn the correct optimal policy. The results demonstrate, however, the effects and importance of the decay parameters on convergence times. The stronger the decay is, the less time it will take the algorithm to converge. A too-short decay, on the other hand, will damage exploration and the agent will fail to converge to an optimal policy. Whenever we present simulation results from here on, we will present the results obtained from the simulation with the shortest decay parameter that still converged to the optimal policy. Each simulation we run, we run multiple times with λ values of 1, 0.1, 0.001, 0.0001 and 0.00001.

5.9.2 *Heuristic 2: Learning an optimal heart rate recovery sequence*

Next, we develop a heuristic to evaluate the agent's optimisation for physiological rewards. In this heuristic, the occupants have no interactions with the agent, letting physiology be the sole factor driving the rewards model. With this heuristic, we run simulations for each of the study participants separately. Since the underlying physiological models, which govern the simulations, differ between participants, we expect the agent to converge to different optimal policies – tailored for the participant.

Table 2 summarises the optimal policy the agent converged to, for each of the participants. For brevity, we do not present the full sequence, but rather extract per context the scene the agent converged to choosing. The correctness of the optimal policy can be verified using Figure 25. A correct policy will choose the lowest probability bar for stress development during test and the highest probability column for restoration during break. Encouragingly, all the policies the agent converged to are the correct one.

These results demonstrate the agent's excellent capability to personalise its actions. We note the high variability in the optimal policies among participants. For each participant, the agent succeeded in learning the individual underlying physiological model, and adjusting its actions accordingly. This is in spite of the stochastic nature of the physiological signal, which poses a significant challenge.

The agent's ability to learn temporal patterns within the episodes is also well demonstrated. In all cases, the agent has successfully identified the latent sequence of context changes and modified its actions along the sequence accordingly. As a result, we observe the agent

Subject	Test scene	Break scene	λ
1	Neutral	Forest	0.0001
2	Forest	Forest	0.0001
3	Rotch	Forest	0.001
4	Rotch	Kites	0.00001
5	Neutral	Shibuja	0.0001
6	Rotch	Forest	0.0001
7	Rotch	Kites	0.00001
8	Kites	Neutral	0.0001
9	Forest	Kites	0.0001

Table 2: Simulation results for Heuristic 2. For each participant the table summarises the optimal policy the agent converged to, i.e. the optimal scene for each working context. λ is the ϵ -greedy decay parameter of our exploration strategy.

choosing different scenes for different contexts, even though it has no direct knowledge about the occupant’s current context. Instead, the agent was able to identify, by proxy, the temporal pattern in which the physiological models change, and learn an optimal policy accounting for it.

Figure 39 illustrates the progression of episodic physiological rewards during the simulation for each of the participants. For clarity, the results are smoothed using a 150 samples moving average filter. Interestingly, we obtain different convergence times for different participants. For example it takes 400 episodes for the agent to converge for Subject 3 and about 4000 episodes to converge for Subject 8. Convergence for Subjects 4 and 7 is achieved outside the range of the presented graph after about 30,000 episodes. These differences appear to be related to the differences in participant reactions among the scenes. Referring to Figure 25, the choice of optimal scenes of Subject 3 is very clear; For Subject 7 the scene reactions are more similar, making the optimal scenes more difficult to identify.

We note that the horizon physiological reward differs between the study participants. For Subject 3, for instance, the expected optimal reward is about 4, whereas it is only 1.5 for Subject 9. According to our rewards model, this translates for Subject 3, into a 75% probability on average, that heart rate will decrease in any 20-second interval of an episode; A similar probability of about 60% for heart rate reduction is obtained for Subject 9. These hint towards the potential beneficial effects, that the Digital Cubicle, combined with the smart agent, may achieve.




Figure 39: Rewards progression throughout simulations of Heuristic-2 for each of the study participants

5.9.3 Heuristic 3: Combining focus and stress restoration

The final heuristic evaluates the agent’s ability to combine physiological optimisation with input from occupant interactions. In this heuristic, occupants will request the agent to skip a scene after 10 seconds in a test context whenever the scene is distracting and interferes with their ability to direct attention. Occupants will accept any scene during break contexts. We take into account the Necker scores to determine if the simulated occupant will find a scene distracting or not. Scenes with negative Necker Z-score are considered to allow focusing, and distracting otherwise. Like before, we ran simulations for each of the study participants separately.

Table 3 summarises the simulations results. These indicate that the agent has successfully converged to the correct optimal policy, for each of the participants. The policies’ correctness can be verified using Figure 25 and Figure 31 from the previous chapter. The optimal policy chooses for the test context the scene with the lowest probability of stress development, whose Necker Z-score is positive. For the break context, the optimal policy maximises for the physiology rewards.

These results successfully demonstrate our approach’s ability to combine multiple control objectives into our smart agent. We exemplify how these can be formulated and balanced through the rewards model we implemented. In our case, we favour the occupants’ preferences over desired changes in occupants’ physiology. Our approach, however, is flexible enough to balance the two differently and even incorporate additional objectives with minor changes.

Figure 40 illustrates progression of the episodic preference, physiological and total rewards during the simulation for each of the participants. We note that convergence times of this heuristic are generally longer than convergence times of the previous one. In fact, for Subject 4 the simulation did not converge for the decay constants we

Subject	Test scene	Break scene	λ
1	Rotch	Forest	0.0001
2	Rotch	Forest	0.0001
3	Rotch	Forest	0.0001
4	N/A	N/A	0.00001
5	Neutral	Shibuja	0.001
6	Rotch	Forest	0.00001
7	Neutral	Kites	0.00001
8	Kites	Neutral	0.001
9	Shibuja	Kites	0.0001

Table 3: Simulation results for Heuristic 3. For each participant the table summarises the optimal policy the agent converged to, i.e. the optimal scene for each working context. λ is the ϵ -greedy decay parameter of our exploration strategy.

normally used. It took a simulation with λ of 10^{-6} for the agent to converge to an optimal policy for her.

As with the previous heuristic, we observe different convergence times for different participants. Convergence for participants 4, 7 and 6 is not presented in the plot's range and was achieved after about 10000 steps. Interestingly, we see that the physiology and preference rewards for the same participants may take different times to converge. For Subject 5, for example, the preference reward converges after about 300 steps where as the physiology reward converges only after about 1000.

5.10 THE CHALLENGE OF CONVERGENCE TIME

Our simulation results demonstrate our agent's excellent abilities, meeting all the design principles outlined in Section 5.2. They also highlight, however, convergence time as its greatest weakness. In our particular scenario, a convergence time of 4000 episodes is equivalent to about 75 working days of using our system, assuming that the working episode is continuously repeated during a full 8 hours working days.

Convergence time in the order of $10^{4\text{act}}$ episodes makes it somewhat unlikely that the agent will converge to an optimal policy through common usage. During such long periods of times changes in the occupants' routines and even baseline physiology may occur, jeopardising the agent's ability to converge. We note that even if the agent did not yet converge, it still gradually improves its performance – benefiting the occupant as a result. The workspace, however, will not be utilised to its full potential.




Figure 40: Rewards progression throughout simulations of Heuristic-3 for each of the study participants

5.10.1 Reducing state-action space and parameter tuning

The most straightforward scheme to shorten convergence time is reducing the state-action space and fine parameter tuning. We can trade the size of the state space, thus shortening convergence, for the temporal resolution we wish our agent to learn. Naturally, learning a sequence of three two-minute contexts would be faster than learning a more fine-grained sequence of six one-minute contexts. The algorithm parameters have also been demonstrated to greatly affect convergence. For particular cases, we can exhaustively search the parameter space and tune the algorithm parameters for optimal convergence.

5.10.2 *Boost learning with simulation*

We propose an additional scheme, in which we see great promise, to shorten the agent convergence and improve its performance in this regard. This scheme is not implemented in the scope of this work. It is laid out here as a starting point for future explorations.

Assume that the agent had, for some occupant, experimental data like the ones collected during the user study. It could then follow the same steps, detailed throughout this chapter, to automatically execute a simulation of its performance, anchoring the simulation with the data it has. In this case, when the occupant is about to begin a new episode in the workspace, the agent can come synthetically more prepared. It can set its initial state to be the state after, for instance, 500 simulated episodes. If the occupant models constructed from the data are somewhat correct, the agent is likely to significantly improve its convergence time.

Once the occupant is done with the workspace, new data is now available to the agent. The agent learns from it of course through the traditional learning techniques we have previously described. However, the same data can also be utilised to update the occupant models. For example, the pdfs for physiological changes, can be recalculated to include the newly available data. In preparation for the next occupant usage, the agent can in turn run an additional simulation using the updated pdfs.

Formally, this approach can be regarded as model-based Reinforcement Learning [49]. In model-based Reinforcement Learning the algorithm attempts to additionally learn the state transition and reward function, so as to simulate arbitrary amounts of experimental data. Using this data it can then approximate a solution for the MDP in a process that can be thought of as planning through forward look-ahead search. Though computationally expensive, such an approach is expected to significantly shorten convergence time. It depends strongly on the model used to learn the reward function. We believe, that the probabilistic modelling we have developed is well suited for this task. It captures well the occupant's reactions and can be easily updated with newly available data.

CONCLUSIONS AND NEXT STEPS

Throughout this work, we have explored our vision of intelligent ambiance in the workspace. Our exploration was split into three stages:

1. **Digital Cubicle**

First, we demonstrated how ambiance could be mediated with digital media. Building the Digital Cubicle, we leveraged lighting, video projection and sound to dynamically manipulate the physical characteristics of a space. A set of tools was presented, experimenting with objective measurements as a mean to compose scenes in the Digital Cubicle. These were shown to elicit different perceptions and experiences. A flexible sensor data collection infrastructure was introduced, enabling real-time data collection, to infer occupant state and context.

We wish to build on our current design to expand our exploration in the future. Through recent collaboration, we are investigating opportunities to commercially produce a second-iteration prototype that embeds the digital components in the furniture, resulting in one integral setup. Besides aesthetics, this will allow producing more than a single prototype and for those to be transported easily. Such an improvement is important for deployment in office spaces, investigating possible effects in real-life scenarios. We hope this will also make our prototype more accessible to other researchers, developing additional physical manipulations and data collection schemes on top of it. Likewise, we hope the library of available scenes and artistic composition effects will grow.

2. **Physiological Effects**

Second, we investigated how the digital compositions affect occupant physiological responses. We presented a signal processing technique to quantify stress development and restoration responses, on top of innate regulation tendencies, and in various working contexts. Our user study indicated highly personal, yet significantly different physiological responses per participant. Similar trends, using an objective measure, were reported in occupants' capability to direct attention in the presence of the different digital compositions. Interestingly, we also found correlations between objective occupant responses (both focusing and stress restoration) and their self-reported perception of the space.

These observations corroborate our hypothesis on possible physiological effects of the Digital Cubicle and call for further investigation. As an immediate step, we wish to analyse the remaining physiological signals we have collected, including occupant electrodermal activity, facial expressions and head and viewing orientations. These may not only identify other manifestations of stress development and restorations, but also shed light on other relevant aspects such as concentration levels and fatigue. Moreover, by repeating our experiment with the same participants, we intend to investigate the consistency of our observations over time. This is important to ascertain our learnings and establish the gravity of possible longitudinal effects like habituation. In a different line of inquiry, empirical data should also be collected in working contexts, other than just tests and breaks (e.g. during creative thinking, or tinkering). This will allow gaining insight into the contextual dependency of one's physiological responses and may lead to better characterisation of the Digital Cubicle effects.

3. Intelligent Control

Fianlly, we proposed a framework to develop intelligent, adaptive control agents – leveraging Reinforcement Learning techniques to bridge real-time sensor data with ambiance mediation. Formulating the problem and implementing a prototype agent, we have demonstrated how agents can be designed to foster occupant wellbeing. We also demonstrated how agents could combine multiple objectives – heart rate recovery and preferences – in their learning and optimisation. We evaluated our agent through simulations, anchored by real human subject data. These showed our agent successfully converging to correct optimal policies, identifying latent temporal dependencies in occupant context and personalising scene selection to match individual physiological response. For some subjects, we estimated improvements from chance of up to 25 % in the probability of heart rate recovery.

Motivated by these simulations we look forward to extending our evaluation, including real human subjects. As convergence time was identified to be a substantial limiting factor, we intend to utilise our physiological simulations as a model-based bootstrapping scheme (See 5.10). We will soon be conducting a short user study, exposing the participants from the physiological study to the optimal policies the agent has learned for them. This will provide quick assessment of this scheme's efficacy. Furthermore, we plan to research other models for state, actions and rewards within our framework. For instance, we are encouraged to examine a non-sequential states model, using sensor data rather the agent's actions to infer occupant context.

Likewise, we wish to experiment with a rewards model that in which occupants can set physiological goals, rather than naively directing physiology in a predefined direction.

APPENDIX: PHYSIOLOGICAL STUDY MATERIALS

A.1 COGNITIVE TASK

The study participants were asked to answer a reading comprehension task during the test segment of the experimental protocol. The following is a sample of such task.

A.1.1 *Text (presented in the left of the screen)*

Adaptation, the decrease in responsiveness that follows continuous stimulation, is common to all sensory systems, including olfaction. With continued exposure to chronically present ambient odours, individuals' perception of odour intensity is greatly reduced. Moreover, these perceptual changes can be profound and durable. It is commonly reported that following extended absences from the odorous environment, re-exposure may still fail to elicit perception at the original intensity.

*Most **research on olfactory adaptation** examines relatively transient changes in stimulus detection or perceived intensity—rarely exceeding several hours and often less—but because olfactory adaptation can be produced with relatively short exposures, these durations are sufficient for investigating many parameters of the phenomenon. However, exposures to odours in natural environments often occur over far longer periods, and the resulting adaptations may differ qualitatively from short-term olfactory adaptation. For example, studies show that even brief periods of odourant stimulation produce transient reductions in receptors in the olfactory epithelium, a process termed **receptor fatigue**. Prolonged odour stimulation, however, could produce more long-lasting reductions in response, possibly involving structures higher in the central nervous system pathway.*

A.1.2 *Questions (presented in the right of the screen)*

Select and indicate the best answer from among the five answer choices.

According to the passage, the phenomenon of olfactory adaptation may cause individuals who are re-exposed to an odorous environment after an extended absence to:

- 1. experience a heightened perception of the odour*
- 2. perceive the odour as being less intense than it was upon first exposure*
- 3. return to their original level of perception of the odour*

4. *exhibit a decreased tolerance for the odorous environment*
5. *experience the phenomenon of adaptation in other sensory*

Select and indicate the best answer from among the five answer choices.

*The phrase **research on olfactory adaptation** appears in the middle of the passage, where it is shown **boldfaced**. The passage asserts which of the following about the exposures involved in that research?*

1. *The exposures are of long enough duration for researchers to investigate many aspects of olfactory adaptation.*
2. *The exposures have rarely consisted of re-exposures following extended absences from the odorous environment.*
3. *The exposures are intended to reproduce the relatively transient olfactory changes typical of exposures to odours in natural environments.*
4. *Those exposures of relatively short duration are often insufficient to produce the phenomenon of receptor fatigue in study subjects.*
5. *Those exposures lasting several hours produce reductions in receptors in the olfactory epithelium that are similar to the reductions caused by prolonged odour stimulation.*

Select and indicate the best answer from among the five answer choices.

*The phrase **receptor fatigue** appears near the end of the passage, where it is shown **underlined and boldfaced**. The author of the passage discusses **receptor fatigue** primarily in order to:*

1. *explain the physiological process through which long-lasting reductions in response are thought to be produced.*
2. *provide an example of a process that subjects would probably not experience during a prolonged period of odourant stimulation.*
3. *help illustrate how the information gathered from most olfactory research may not be sufficient to describe the effects of extended exposures to odours.*
4. *show how studies of short-term olfactory adaptation have only accounted for the reductions in response that follow relatively brief absences from an odorous environment.*
5. *qualify a statement about the severity and duration of the perceptual changes caused by exposure to chronically present ambient odours.*

[CORRECT ANSWERS: 2, 1, 3]

A.2 SURVEY

At the end of each experimental session the study participants were presented with the following survey. The survey was presented through the experiment's website and was filled electronically.

Please answer the following questions:

1. *Recall one of those times when you worked hard on a project that required intense and prolonged intellectual effort. Remember how it felt. You probably reached a point where you could tell that your ability to work effectively had started to decline and that you needed a break. You needed to do something during the break that would restore your ability to work effectively on the project. Put yourself in that mind set now and please rate, how good this setting would be to take a break and restore your ability to work effectively on your project.*

(Not very good) -2 | -1 | 0 | +1 | +2 (Very good)

2. *You have just finished breakfast and have only one thing on your agenda for the day. You have a project that you need to think about. Thinking deeply and thoroughly about this project is your goal. Please rate this setting on how good a place it is to accomplish your goal.*

(Not very good) -2 | -1 | 0 | +1 | +2 (Very good)

3. *How much do you like the setting? This is your own personal degree of liking for the setting, and you do not have to worry about whether you are right or wrong or whether you agree with anybody else.*

(Not very much) -2 | -1 | 0 | +1 | +2 (Very much)

4. *Sometimes even when you are near your office it can still feel like you are far away from everyday thoughts and concerns. How much does this setting provide an escape experience or a feeling of being away?*

(Not very much) -2 | -1 | 0 | +1 | +2 (Very much)

5. *Some settings have many interesting things that can draw your attention. How much does this setting easily and effortlessly engage your interest? How much does it fascinate you?*

(Not very much) -2 | -1 | 0 | +1 | +2 (Very much)

6. *Sometimes a setting can feel like a whole world of its own. How much does this setting feel like there is much to explore and discover in many directions?*

(Not very much) -2 | -1 | 0 | +1 | +2 (Very much)

7. *Some settings are confusing, have no organisation and have too much going on. Please rate how chaotic and distracting this setting feels?*

(Not very much) -2 | -1 | 0 | +1 | +2 (Very much)

BIBLIOGRAPHY

- [1] Fons Adriaensen. "A tetrahedral microphone processor for ambisonic recording." In: *Proc. of the 5th Int. Linux Audio Conference (LACo7), Berlin, Germany*. 2007, pp. 64–69.
- [2] Judith Amores Fernandez. "Essence: Olfactory Interfaces for Unconscious Influence of Mood and Cognitive Performance." MA thesis. Massachusetts Institute of Technology, 2016.
- [3] Asier Aztiria, Alberto Izaguirre, and Juan Carlos Augusto. "Learning patterns in ambient intelligence environments: a survey." In: *Artificial Intelligence Review* 34.1 (2010), pp. 35–51.
- [4] AA Bak and DE Grobbee. "A randomized study on coffee and blood pressure." In: *Journal of human hypertension* 4.3 (1990), pp. 259–264.
- [5] Robert A Baron. "The physical environment of work settings: Effects on task performance, interpersonal relations, and job satisfaction." In: *Research in organizational behavior* 16 (1994), pp. 1–1.
- [6] Yehuda Baruch. "Teleworking: benefits and pitfalls as perceived by professionals and managers." In: *New Technology, Work and Employment* 15.1 (2000), pp. 34–49.
- [7] Marc G Berman, John Jonides, and Stephen Kaplan. "The cognitive benefits of interacting with nature." In: *Psychological science* 19.12 (2008), pp. 1207–1212.
- [8] Christopher M Bishop. "Pattern recognition." In: *Machine Learning* 128 (2006).
- [9] Geoffray Bonnin and Dietmar Jannach. "Automated generation of music playlists: Survey and experiments." In: *ACM Computing Surveys (CSUR)* 47.2 (2015), p. 26.
- [10] Gregory N Bratman, J Paul Hamilton, and Gretchen C Daily. "The impacts of nature experience on human cognitive function and mental health." In: *Annals of the New York Academy of Sciences* 1249.1 (2012), pp. 118–136.
- [11] Chun-Yen Chang, William E Hammitt, Ping-Kun Chen, Lisa Machnik, and Wei-Chia Su. "Psychophysiological responses and restorative values of natural environments in Taiwan." In: *Landscape and Urban Planning* 85.2 (2008), pp. 79–84.

- [12] Chung-Yi Chi, Richard Tzong-Han Tsai, Jeng-You Lai, and Jane Yung-jen Hsu. "A reinforcement learning approach to emotion-based automatic playlist generation." In: *2010 International Conference on Technologies and Applications of Artificial Intelligence*. IEEE. 2010, pp. 60–65.
- [13] Hwan-Hee Choi, Jeroen JG Van Merriënboer, and Fred Paas. "Effects of the physical environment on cognitive load and learning: towards a new model of cognitive load." In: *Educational Psychology Review* 26.2 (2014), pp. 225–244.
- [14] Bernadine Cimprich. "Attentional fatigue and restoration in individuals with cancer." In: (1990).
- [15] Bernadine Cimprich. "Development of an intervention to restore attention in cancer patients." In: *Cancer nursing* 16.2 (1993), pp. 83–92.
- [16] Diane J Cook and Sajal K Das. "How smart are our environments? An updated look at the state of the art." In: *Pervasive and mobile computing* 3.2 (2007), pp. 53–73.
- [17] Kevin C Darr, DAVID R Bassett, BARBARA J Morgan, and D PAUL Thomas. "Effects of age and training status on heart rate recovery after peak exercise." In: *American Journal of Physiology-Heart and Circulatory Physiology* 254.2 (1988), H340–H343.
- [18] YAW De Kort, AL Meijnders, AAG Sponselee, and WA IJsselstein. "What's wrong with virtual trees? Restoring from stress in a mediated environment." In: *Journal of environmental psychology* 26.4 (2006), pp. 309–320.
- [19] Gwyneth Doherty-Sneddon, Vicki Bruce, Lesley Bonner, Sarah Longbotham, and Caroline Doyle. "Development of gaze aversion as disengagement from visual information." In: *Developmental psychology* 38.3 (2002), p. 438.
- [20] Ken Ducatel, Marc Bogdanowicz, Fabiana Scapolo, Jos Leijten, and Jean-Claude Burgelman. "Scenarios for ambient intelligence in 2010." In: *Office for official publications of the European Communities* (2001).
- [21] ETS. *The GRE Tests*. 2016. URL: <http://www.ets.org/gre/> (visited on 07/26/2016).
- [22] T Edward. *Hall, The hidden dimension*. 1966.
- [23] Gary W Evans and Rachel Stecker. "Motivational consequences of environmental stress." In: *Journal of Environmental Psychology* 24.2 (2004), pp. 143–165.
- [24] Abe Fettig. *Twisted network programming essentials*. "O'Reilly Media, Inc.", 2005.

- [25] MG Figueiro, GC Brainard, SW Lockley, VL Revell, and R White. "Light and Human Health: An Overview of the Impact of Optical Radiation on Visual, Circadian, Neuroendocrine, and Neurobehavioral Responses." In: *Illuminating Engineering Society Technical memorandum, IES TM-18-08* (2008).
- [26] Maurizio Garbarino, Matteo Lai, Dan Bender, Rosalind W Picard, and Simone Tognetti. "Empatica E3?A wearable wireless multi-sensor device for real-time computerized biofeedback and data acquisition." In: *Wireless Mobile Communication and Healthcare (Mobihealth), 2014 EAI 4th International Conference on*. IEEE. 2014, pp. 39–42.
- [27] VF Gladwell, DK Brown, Joanna L Barton, MP Tarvainen, P Kuoppa, J Pretty, JM Suddaby, and GRH Sandercock. "The effects of views of nature on autonomic control." In: *European journal of applied physiology* 112.9 (2012), pp. 3379–3386.
- [28] Wolfgang Härdle. *Smoothing techniques: with implementation in S*. Springer Science & Business Media, 2012.
- [29] Terry Hartig, Kalevi Korpela, Gary W Evans, and Tommy Gärling. "A measure of restorative quality in environments." In: *Scandinavian Housing and Planning Research* 14.4 (1997), pp. 175–194.
- [30] Terry Hartig, Gary W Evans, Larry D Jamner, Deborah S Davis, and Tommy Gärling. "Tracking restoration in natural and urban field settings." In: *Journal of environmental psychology* 23.2 (2003), pp. 109–123.
- [31] John Hattie and David Watkins. "Preferred classroom environment and approach to learning." In: *British Journal of Educational Psychology* 58.3 (1988), pp. 345–349.
- [32] Selig Hecht. "The visual discrimination of intensity and the Weber-Fechner law." In: *The Journal of general physiology* 7.2 (1924), pp. 235–267.
- [33] Arran Holmes, Hakan Duman, and Anthony Pounds-Cornish. "The iDorm: Gateway to heterogeneous networking environments." In: *International ITEA Workshop on Virtual Home Environments*. Vol. 1. 2002, pp. 20–21.
- [34] Steelcase Inc. *Acoustic Solutions*. 2016. URL: <https://www.steelcase.com/products/technology/acoustic-solutions/> (visited on 07/26/2016).
- [35] Hiroshi Ishii and Brygg Ullmer. "Tangible bits: towards seamless interfaces between people, bits and atoms." In: *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. ACM. 1997, pp. 234–241.

- [36] Hiroshi Ishii, Craig Wisneski, Scott Brave, Andrew Dahley, Matt Gorbet, Brygg Ullmer, and Paul Yarin. "ambientROOM: integrating ambient media with architectural space." In: *CHI 98 Conference Summary on Human Factors in Computing Systems*. ACM. 1998, pp. 173–174.
- [37] H 1 Ising, B Kruppa, et al. "Health effects caused by noise: evidence in the literature from the past 25 years." In: *Noise and Health* 6.22 (2004), p. 5.
- [38] Joris H Janssen, Egon L van den Broek, and Joyce HDM Westerink. "Tune in to your emotions: a robust personalized affective music player." In: *User Modeling and User-Adapted Interaction* 22.3 (2012), pp. 255–279.
- [39] Patrik N Juslin and John Sloboda. *Handbook of music and emotion: Theory, research, applications*. Oxford University Press, 2011.
- [40] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. "Reinforcement learning: A survey." In: *Journal of artificial intelligence research* 4 (1996), pp. 237–285.
- [41] Peter H Kahn, Batya Friedman, Brian Gill, Jennifer Hagman, Rachel L Severson, Nathan G Freier, Erika N Feldman, Sybil Carrère, and Anna Stolyar. "A plasma display window? The shifting baseline problem in a technologically mediated natural world." In: *Journal of Environmental Psychology* 28.2 (2008), pp. 192–199.
- [42] N Kamarulzaman, AA Saleh, SZ Hashim, H Hashim, and AA Abdul-Ghani. "An overview of the influence of physical office environments towards employee." In: *Procedia Engineering* 20 (2011), pp. 262–268.
- [43] Syahrul Nizam Kamaruzzaman and Emma Marinie Ahmad Zawawi. "Influence of Employees' Perceptions of Colour Preferences on Productivity in Malaysian Office Buildings." In: *Journal of Sustainable Development* 3.3 (2010), p. 283.
- [44] Rachel Kaplan and Stephen Kaplan. *The experience of nature: A psychological perspective*. CUP Archive, 1989.
- [45] Stephen Kaplan. "The restorative benefits of nature: Toward an integrative framework." In: *Journal of environmental psychology* 15.3 (1995), pp. 169–182.
- [46] Igor Knez. "Effects of indoor lighting on mood and cognition." In: *Journal of environmental psychology* 15.1 (1995), pp. 39–51.
- [47] John Kowal and Michelle S Fortier. "Motivational determinants of flow: Contributions from self-determination theory." In: *The Journal of Social Psychology* 139.3 (1999), pp. 355–368.

- [48] William H Kruskal and W Allen Wallis. "Use of ranks in one-criterion variance analysis." In: *Journal of the American statistical Association* 47.260 (1952), pp. 583–621.
- [49] D Kuvayev and Richard S Sutton. *Model-based reinforcement learning*. Tech. rep. Citeseer, 1997.
- [50] Li Lan, Pawel Wargocki, David Peter Wyon, and Z Lian. "Effects of thermal discomfort in an office on perceived air quality, SBS symptoms, physiological responses, and human performance." In: *Indoor Air* 21.5 (2011), pp. 376–390.
- [51] Ya-Yun Lee, Carolee J Winstein, James Gordon, Giselle M Petzinger, Elizabeth M Zelinski, and Beth E Fisher. "Context-dependent learning in people with Parkinson's disease." In: *Journal of motor behavior* (2015), pp. 1–9.
- [52] Elad Liebman, Maytal Saar-Tsechansky, and Peter Stone. "Djmc: A reinforcement-learning agent for music playlist recommendation." In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems. 2015, pp. 591–599.
- [53] Remco Magielse and Philip R Ross. "A design approach to socially adaptive lighting environments." In: *Proceedings of the 9th ACM SIGCHI Italian Chapter International Conference on Computer-Human Interaction: Facing Complexity*. ACM. 2011, pp. 171–176.
- [54] Brian Mayton. *Tidmarsh Living Observatory*. 2016. URL: <http://tidmarsh.media.mit.edu/> (visited on 07/26/2016).
- [55] Janetta Mitchell McCoy and Gary W Evans. "The potential role of the physical environment in fostering creativity." In: *Creativity Research Journal* 14.3-4 (2002), pp. 409–426.
- [56] Ravi Mehta and Rui Juliet Zhu. "Blue or red? Exploring the effect of color on cognitive task performances." In: *Science* 323.5918 (2009), pp. 1226–1229.
- [57] Ravi Mehta, Rui Juliet Zhu, and Amar Cheema. "Is noise always bad? Exploring the effects of ambient noise on creative cognition." In: *Journal of Consumer Research* 39.4 (2012), pp. 784–799.
- [58] Joan Meyers-Levy and Rui Juliet Zhu. "The influence of ceiling height: The effect of priming on the type of processing that people use." In: *Journal of Consumer Research* 34.2 (2007), pp. 174–186.
- [59] MC Mozer. *Lessons from an Adaptive Home. Smart Environments: Technologies, Protocols, and Applications*, edited by DJ Cook and SK Das. 2005.

- [60] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [61] Nuria Oliver and Fernando Flores-Mangas. "MPTrain: a mobile, music and physiology-based personal trainer." In: *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*. ACM. 2006, pp. 21–28.
- [62] Jifei Ou, Gershon Dublon, Chin-Yi Cheng, Felix Heibeck, Karl Willis, and Hiroshi Ishii. "Cillia: 3D Printed Micro-Pillar Structures for Surface Texture, Actuation and Sensing." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM. 2016, pp. 5753–5764.
- [63] Fred GWC Paas and Jeroen JG Van Merriënboer. "Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach." In: *Journal of educational psychology* 86.1 (1994), p. 122.
- [64] Margherita Pasini, Rita Berto, Margherita Brondino, Rob Hall, and Catherine Ortner. "How to measure the restorative quality of environments: The PRS-11." In: *Procedia-Social and Behavioral Sciences* 159 (2014), pp. 293–297.
- [65] Rosalind W Picard and Roalind Picard. *Affective computing*. Vol. 252. MIT press Cambridge, 1997.
- [66] Jan L Plass, Roxana Moreno, and Roland Brünken. *Cognitive load theory*. Cambridge University Press, 2010.
- [67] Thorsten Prante, Richard Stenzel, Carsten Röcker, Norbert Streititz, and Carsten Magerkurth. "Ambient agoras: Inforiver, siam, hello. wall." In: *CHI'04 Extended Abstracts on Human Factors in Computing Systems*. ACM. 2004, pp. 763–764.
- [68] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [69] Bradley J Rhodes and Pattie Maes. "Just-in-time information retrieval agents." In: *IBM Systems journal* 39.3.4 (2000), pp. 685–704.
- [70] Marc Jeffrey Rosenberg. *E-learning: Strategies for delivering knowledge in the digital age*. Vol. 3. McGraw-Hill New York, 2001.
- [71] Pierre Salame and Alan Baddeley. "Disruption of short-term memory by unattended speech: Implications for the structure of working memory." In: *Journal of verbal learning and verbal behavior* 21.2 (1982), pp. 150–164.
- [72] Jan C Schacher and Philippe Kocher. "Ambisonics spatialization tools for max/msp." In: *Omni 500* (2006), p. 1.
- [73] D Schmidt, Michael Stal, Hans Rohnert, and F Bushmann. *Patterns for Concurrent and Networked Objects, volume 2 of Pattern-Oriented Software Architecture*. 2000.

- [74] RA Schmidt. "8c Lee, TD (2005)." In: *Motor control and learning: A behavioral emphasis (4th ed.)*. Champaign, IL: Human Kinetics ().
- [75] Nandita Sharma and Tom Gedeon. "Objective measures, sensors and computational techniques for stress recognition and classification: A survey." In: *Computer methods and programs in biomedicine* 108.3 (2012), pp. 1287–1301.
- [76] Bernard W Silverman. *Density estimation for statistics and data analysis*. Vol. 26. CRC press, 1986.
- [77] Tobias Skog, Sara Ljungblad, and Lars Erik Holmquist. "Between aesthetics and utility: designing ambient information visualizations." In: *Information Visualization, 2003. INFOVIS 2003. IEEE Symposium on*. IEEE. 2003, pp. 233–240.
- [78] Steven M Smith and Edward Vela. "Environmental context-dependent memory: A review and meta-analysis." In: *Psychonomic bulletin & review* 8.2 (2001), pp. 203–220.
- [79] Andrea Faber Taylor, Frances E Kuo, and William C Sullivan. "Views of nature and self-discipline: Evidence from inner city children." In: *Journal of environmental psychology* 22.1 (2002), pp. 49–63.
- [80] Carolyn M Tennessen and Bernadine Cimprich. "Views to nature: Effects on attention." In: *Journal of environmental psychology* 15.1 (1995), pp. 77–85.
- [81] Basheer Tome. "Exoskin: pneumatically augmenting inelastic materials for texture changing interfaces." PhD thesis. Massachusetts Institute of Technology, 2015.
- [82] Mikko P Tulppo, Timo H Mäkikallio, Tapio Seppänen, Raija T Laukkanen, and Heikki V Huikuri. "Vagal modulation of heart rate during exercise: effects of age and physical fitness." In: *American Journal of Physiology-Heart and Circulatory Physiology* 274.2 (1998), H424–H429.
- [83] Endel Tulving and Donald M Thomson. "Encoding specificity and retrieval processes in episodic memory." In: *Psychological review* 80.5 (1973), p. 352.
- [84] Roger S Ulrich. "Visual landscapes and psychological well-being." In: *Landscape research* 4.1 (1979), pp. 17–23.
- [85] Roger S Ulrich. "Aesthetic and affective response to natural environment." In: *Behavior and the natural environment*. Springer, 1983, pp. 85–125.
- [86] Roger S Ulrich. "Biophilia, biophobia, and natural landscapes." In: *The biophilia hypothesis* 7 (1993).

- [87] Roger S Ulrich, Robert F Simons, Barbara D Losito, Evelyn Fiorito, Mark A Miles, and Michael Zelson. "Stress recovery during exposure to natural and urban environments." In: *Journal of environmental psychology* 11.3 (1991), pp. 201–230.
- [88] Eunjoon Um, Jan L Plass, Elizabeth O Hayward, Bruce D Homer, et al. "Emotional design in multimedia learning." In: *Journal of Educational Psychology* 104.2 (2012), p. 485.
- [89] Juan M Del Castillo Von Haucke. *Modular privacy screen assemblies*. US Patent 4,928,465. 1990.
- [90] Christopher JCH Watkins and Peter Dayan. "Q-learning." In: *Machine learning* 8.3-4 (1992), pp. 279–292.
- [91] Joseph Wilder. *Stimulus and response: The law of initial value*. Elsevier, 2014.
- [92] Glenn F Wilson and Christopher A Russell. "Real-time assessment of mental workload using psychophysiological measures and artificial neural networks." In: *Human Factors: The Journal of the Human Factors and Ergonomics Society* 45.4 (2003), pp. 635–644.
- [93] Ngai-ying Wong and David Watkins. "Self-monitoring as a mediator of person-environment fit: an investigation of Hong Kong mathematics classroom environments." In: *British Journal of Educational Psychology* 66.2 (1996), pp. 223–229.
- [94] Xuehan Xiong and Fernando De la Torre. "Supervised descent method and its applications to face alignment." In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2013, pp. 532–539.
- [95] Alison Jing Xu and Aparna Labroo. "Incandescent affect: Turning on the hot emotional system with bright light." In: *NA-Advances in Consumer Research Volume 41* (2013).
- [96] Zephyr. *Bioharness* 3. 2016. URL: <https://www.zephyranywhere.com/products/bioharness-3> (visited on 07/26/2016).
- [97] Nan Zhao. "Mediated Atmospheres for an Adaptive Built Environment." PhD thesis. Massachusetts Institute of Technology, 2016.
- [98] Nan Zhao, Matthew Aldrich, Christoph F Reinhart, and Joseph A Paradiso. "A Multidimensional Continuous Contextual Lighting Control System Using Google Glass." In: *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*. ACM. 2015, pp. 235–244.
- [99] Marjolein D van der Zwaag, Joris H Janssen, and Joyce HDM Westerink. "Directing physiology and mood through music: Validation of an affective music player." In: *IEEE Transactions on Affective Computing* 4.1 (2013), pp. 57–68.